

# Feasibility in Action and Attitude<sup>φ</sup>

G. Brennan\* and N. Southwood\*\*

## I Introduction

According to the principle of “ought-implies-can”,

(OC) A ought to X in C only if A can X in C.

There are many different versions of (OC).<sup>1</sup> The version that we are interested in interprets the claim that “A can X in C” as meaning that “it is *feasible* that A X in C”. According to what we shall call the principle of “ought-implies-feasibility”,

(OF) A ought to X in C only if it is feasible that A X in C.

(OF) is a non-trivial principle. First, it has important substantive implications. It rules out many options (including many desirable options<sup>2</sup>). If we lack relevant medical knowledge and expertise, it may not be feasible that we perform a delicate neurological operation on a patient’s right hemisphere – even if this is the only way to save her life. If we are recovering alcoholics, it may not be feasible that we have ‘just one’ glass of Sauternes with our foie gras. If certain contestable claims about human motivation are correct, it may not be feasible that we eradicate global

---

<sup>φ</sup> We have been led to direct this paper to the role of ‘feasibility’ in normative analysis because, in addition to its intrinsic interest, this issue seems to us to be a highly significant one in the conversation between economists and philosophers – a conversation in which Wlodek himself has played such a key role over the last two decades. Although we do not here engage directly with Wlodek’s written work, it should be noted that the distinction between actions and attitudes of which we make much use, is partly stimulated by Wlodek’s writing and partly by many enjoyable and profitable conversations with him.

\* Brennan is Professor of Economics in the Social and Political Theory Program at RSSS, ANU and Professor of Political Science, Duke University and Professor of Philosophy, UNC-Chapel Hill.

\*\* Southwood is Postdoctoral Fellow in the Philosophy Program, RSSS ANU.

<sup>1</sup> For a good survey of the terrain, see Copp 2002.

<sup>2</sup> It is sometimes said that “*desirability* implies feasibility”. This is patently false.

poverty, or bring about a communist Utopia. According to (OF), if these things are indeed infeasible, then they are not things that we ought to do.<sup>3</sup>

Not only does (OF) rule many options out. It also thereby indirectly rules many other options *in*. If we are such bad sports that congratulating and shaking the hand of an opponent who has just defeated us (without smashing his face in) is infeasible, it may turn out that what we ought to do instead is walk away and return when we have cooled off.<sup>4</sup> If we have carelessly missed a plane to attend a friend's wedding in New Zealand and cannot get seats on any alternative flight that will get us there on time, it may turn out that what we ought to do instead is call and apologise. Of course, it may also be true that we ought not to be a bad sport; and that we ought not to have missed the plane. But given that we *are* a bad sport or that we *have* missed the plane, (OF) will often have the consequence that certain less desirable (but feasible) options constitute, in effect, what we ought to do.<sup>5</sup>

(OF) also has important meta-ethical implications. In particular, it seems to go with certain views about the nature of normativity. For example, it seems to go with a view according to which what we ought to do is importantly constrained by the capacities and circumstances that we actually possess and the circumstances that we actually occupy. It supports a conception of normativity as, in an important sense, *real-worldly*. It also seems to go with certain other views concerning the alleged *practicality* or *action-guidingness* of the normative – for example, that in saying of an agent that she “ought to X in C”, we mean to say that it would be appropriate (given full relevant information) to *advise* her to X in C.

---

<sup>3</sup> Notice that if A's performing X is infeasible (OF) implies that it is not the case that A ought to X, not that A ought not to perform X.

<sup>4</sup> This example is due to Michael Smith (1995, p. 111).

<sup>5</sup> For a good discussion of the application of (OC) to cases of the “second best”, see Howard-Snyder 2006, pp. 236-8.

Finally, (OF) has important methodological implications. If it is right that an important class of descriptive facts – the facts that determine what is feasible – have a bearing upon what people ought to do, it suggests that as normative theorists we must take an interest in such facts. This suggests, for example, a potentially significant role for the natural and social sciences in normative theory. We have to be careful here. There is no straightforward transition from facts about whether P is the case to facts about whether P is *feasible* is the case. Nevertheless, the natural and the social sciences set themselves the ambition not just of describing the world but also of explaining its causal structure. The accounts offered in the physical and social sciences purport to explain why, given that W is the case then Z or V must also be – or that if you do T it will have consequence S. Once the possibility of ‘science’ is admitted, it carries with it the implication that not all imaginable states of the world are feasible, or at least not feasible without the intervention of extraordinary factors. Thus, evolutionary biologists and psychologists may show that certain forms of behaviour are sufficiently hard-wired to give us evidence for the claim that radical deviations from such behaviour are not feasible for the sorts of creatures we are. Political scientists may show that certain forms of government and other institutional arrangements are sufficiently unstable to give us evidence for the claim that reform along those lines is not feasible. If (OF) is correct, then such ‘facts’ are things that normative theorists cannot afford to ignore.

Many theorists, however, have expressed doubts about (OF).<sup>6</sup> In particular, it has been argued that (OF) fails in certain important cases. One response to these worries about (OF) is to give up on it. A number of philosophers have taken this path. From our perspective, this is a sizeable cost. There is, of course, room to challenge (OF). The problem, however, is that if (OF) really turns out to be unjustifiable, we may begin to doubt the relevance, to us as we actually are, of claims about what we ought to do. Granted that to tie oughts too closely to facts about feasibility may seem in at

---

<sup>6</sup> See, for example, Stocker 1971; Sinnott-Armstrong 1984; Saka 2000.

least some cases to make those oughts too easily avoided, it is equally the case that to make oughts unconstrained by facts about feasibility invites the charge that they are irrelevant.

In this paper, we pursue a rather different tack. Rather than giving up on (OF), we want to see how far one might finesse the anxieties about its application by appeal to a distinction between (a) what A ought to do – the actions A ought to perform – and (b) how A ought to think about what A does – the attitudes A ought to have towards what A does. According to some views, these come to the same thing. Our view is that this is a mistake; although actions and attitudes about them are clearly connected in almost all cases, they can and do come apart. The sorts of considerations that impact upon, feed into and constrain the one may be rather different from the sorts of considerations that impact upon, feed into and constrain the other.

This opens up the following interesting possibility. Even supposing that (OF) is correct – and hence that feasibility considerations play a decisive, constraining (perhaps trumping) role in the case of action – this leaves open the role (if any) that feasibility considerations play in the case of attitudes towards those actions. And indeed, as we shall argue, there is good reason to suppose that feasibility plays a far more subtle and complicated role in the case of the attitudes than the role that, according to (OF), they play in the case of actions. If we need a slogan: feasibility bites considerably less deeply in the case of attitudes than in the case of actions.

In principle, of course, this leaves open the possibility that the critics of (OF) are right and that feasibility doesn't bite all that deeply even in the case of actions. Our own view is that this isn't right: that (OF) is, in fact, a sound principle, and that giving up on it is a mistake. More interestingly, however, our proposal will show that many of the anxieties regarding (OF) rely upon assuming precisely the kind of

overly tight connection between what A ought to do and how A ought to think about what A does that our proposal is meant to undermine. So, at the very least, the conclusions we reach rather take the sting out of the critics' attack.

The paper is in four main sections. We begin in section II by clarifying what is at stake in (OF). In section III, we discuss some putative counterexamples to (OF), arguing that they are not persuasive. We then turn in section IV to the action/attitude distinction, arguing that some version of the autonomy of attitudes thesis is unavoidable. Finally, in section V, we present our positive proposal according to which feasibility bites less deeply in the case of attitudes than in the case of action.

## II Interpreting (OF)

Recall that, according to (OF),

(OF) A ought to X in C only if it is feasible that A X in C.

We shall take (OF) to be a putative *conceptual* claim – that is, as something like a *platitude* about “ought”, or at least as a priori derivable from such platitudes. This is not the only way to understand (OF), or something in the ballpark of (OF). One might understand (OF) – or more plausibly something rather weaker – as a *normative* principle, i.e. simply as attesting to the normative relevance of feasibility considerations in determining what makes actions such that we ought to perform them. One might, à la Sinnott-Armstrong, hold that “ought” *conversationally implies* “feasibility”.<sup>7</sup> Alternatively, we might understand OF as a metaphysically necessary

---

<sup>7</sup> Sinnott-Armstrong 1984.

claim. We mention these possibilities to set them aside. In this paper, we are only concerned with (OF) as a conceptual claim about the meaning of “ought”.

Our first task will be to offer a somewhat more precise characterisation of (OF). We pointed out that (OF) is one interpretation of (OC). But clearly (OF) itself stands in need of interpretation in respect of:

- the nature of the “ought”;
- the nature of the “X”;
- the nature of the “A”; and
- the notion of feasibility.

We will consider these in turn.

So, first, the “ought”. The “ought” in (OF) is a *normative* “ought”. Not all “oughts” are normative. For instance, there are “oughts” of expectation such as when we say that “she ought to be here by now”. There are also predictive “oughts” such as when we say, looking at the clouds, that “it ought to rain tomorrow”. These non-normative “oughts” are not covered by (OF). The only kinds of “oughts” that are covered by (OF) are normative “oughts”. At the very least, these include moral “oughts” such as “you ought to give money to charity” and prudential “oughts” such as “you ought to insure your house”. Other candidates for normative “oughts” are epistemic “oughts” (“you ought to believe that creationism is false”), “oughts” of etiquette (“you ought to RSVP to wedding invitations”), legal “oughts”, aesthetic “oughts”, “oughts” of grammar, and so on. We’re not sure which of these are normative “oughts”. Nor are we sure how best to characterise the difference between normative and non-normative “oughts”.<sup>8</sup> So we shall mainly focus on the “oughts”

---

<sup>8</sup> One common and reasonably plausible suggestion is that normative “oughts” necessarily imply claims about reasons, whereas non-normative “oughts” don’t necessarily imply claims about reasons. What makes the “ought” in “you ought to give money to charity” normative and the “ought” in “it ought to rain tomorrow” non-normative, according to this proposal, is that it is part of the content of the former that you have a reason to give money to charity, whereas it is not part of the content of the latter that anyone has any reason to do anything. But we want to finesse the nature of the normative here.

that we are sure qualify – moral and prudential oughts. We think it doesn't matter within this class, whether we are thinking of pro tanto "oughts" or all things considered "oughts". Either way, the "ought" lies in the class for which (OF) sees feasibility as a necessary condition.

Consider, next, the "X in C". (OF) says that "A ought to X in C only if it is feasible that A X in C". For the most part, we are interested here in cases where "Xing in C" refers to the performance of particular *actions* in particular circumstances. The "C" then refers to the background conditions – the circumstances – that surround A's X-ing. Notice, however, that the "Xing in C" may also plausibly refer to the possession of certain attitudes, such as belief, desire, guilt, fear, and so on. Consider belief. To believe a proposition requires certain capacities and mastery of certain concepts. Sometimes it may not be feasible that an individual believes a given proposition P. This will be true, for instance, of creatures who lack the cognitive sophistication to believe any propositions whatsoever. It will also be true of creatures who lack the relevant concepts and perhaps the background knowledge necessary to understand the proposition. And it may be true if P contains an (obvious) contradiction. If any of these conditions holds, (OF) implies that it is not the case that A ought to believe that P. And what is true of beliefs here is taken to apply in principle to other 'attitudes'. The point is that it is not just actions we are interested in. Crucially, attitudes are seen to be subject to the same (OF) test.

What then of "A". (OF) says that "A ought to X in C only if it is feasible that A X in C". To whom does this "A" refer? The answer is that it refers to all *particular agents*. It does not refer to non-agents. And it does not refer to "agents in general" or "most agents". The set of particular agents includes particular *collective* agents such as the Australian government, Wal-Mart, the Rotary Club and the ANU Philosophy department, as well as particular *individual* agents such as Wlodek Rabinowicz and Glen McGrath. Many of the most interesting questions about (OF) concern what

collective agents ought to do: whether governments ought to implement a basic income scheme; whether corporations ought to employ environmentally sustainable technologies; and so on. Of course, there are deep issues about what is required to be an agent and whether so-called “collective agents” constitute agents proper. This is not the place to try to resolve those disputes. In some (albeit deflationary) sense of the term “agent”, it makes sense to say that “corporations are agents”. That is the sense that we are interested in here, and which we intend the “A” in (OF) to include.<sup>9</sup>

Finally, what about the notion of “feasibility”? This raises perhaps the trickiest issues of all. One idea is that “feasibility” just means “logical possibility”. We do not think this is right. It is true that for an action to be feasible *requires* that it be logically and nomologically possible. Thus, it is infeasible that we square the circle. And it is infeasible that we levitate to the moon. Feasibility, however, requires more than mere possibility. Immanuel Kant, a prominent and vocal sceptic about (OF), nonetheless granted that “moral aims ... amount to duties ... [only] so long as it is not demonstrably impossible to fulfil them”.<sup>10</sup> Many actions are logically possible yet infeasible. There is nothing logically or nomologically impossible about a medical ignoramus successfully performing a neurological operation for which, as it so happens, he lacks the relevant expertise, or a notoriously bad sport overcoming

---

<sup>9</sup> There are however a couple of complications here to which we should be alert. First, it is to be emphasised that the A captured under the ‘ought’ is the *same* A as is captured under the feasibility constraint. In most cases this will be unproblematic but there is a special difficulty that arises when A is “we”. Often people will say: “we ought to X” – but really have in mind some specific “theys”. Perhaps the remark is addressed to the management of the group in question, or to some subset of the “we’s” who could not by themselves perform the X in question. We don’t think this matters greatly provided the “we” is well-defined and its identity across the (OF) formulation preserved. But “we”, unlike “I”, can be more or less encompassing, and that might allow scope for shifting identities within the (OF) formulation in some cases, and we would want to rule such shifts out. Second, from what we have said before, it should be clear that the A is defined not only in terms of agency but also in terms of “attitude-possession”. Some As that might qualify as agents might not qualify as attitude bearers, and vice versa. We do not, as in the case of agency, want to get embroiled in the question of what is entailed in qualifying as an attitude bearer – but we do want to make it clear that any entity that does qualify falls under the scope of (OF).

<sup>10</sup> Kant, 1977, p. 89.



her bad sportingness and shaking the hand of a victorious opponent without punching the opponent in the face, or the realisation of a communist ideal. Yet these things are the kinds of infeasible options that (OF) is intended to rule out.

Another idea is that “feasibility” means “sufficient probability” or “likelihood”. This may seem to do better. It is presumably extremely unlikely that a medical ignoramus succeeds in performing a delicate operation; that a notoriously bad sport shakes the hand of his victorious opponent; and that we realise a communist ideal. So this proposal gives the right answers in some important cases. In some other cases, however, it doesn’t fare so well. Suppose that one is so lazy that one is highly unlikely to go to one’s daughter’s hockey game on Saturday morning. It’s not that there is anything preventing one from doing so. Nor would it be especially difficult or costly. It’s just that one is so lazy that one will almost certainly stay in bed and read the paper instead. It would be a mistake, we take it, to claim that one’s going to watch one’s daughter’s hockey game is “infeasible”. It’s perfectly feasible. It’s just that the chances are that it won’t happen. Feasibility isn’t the same as sufficient probability.

However, there is an amendment to the sufficient probability view of feasibility that might work. The reason why the sufficient probability view will not do as it stands is that there are different explanations for why it may be unlikely that an agent performs an action; and some of them are not sufficient for it to be infeasible that the agent performs the action. The case of the father who prefers to stay at home and read his newspaper is one such case. The explanation for why it is unlikely that he will be at his daughter’s hockey game is that it is unlikely that he will *try*. If he were to try, then he would have a high probability of success. This suggests a way of understanding feasibility, not in terms of reasonable probability, but in terms of

reasonable probability of success conditional upon trying. So far as we can tell, this analysis gives the right answers in most of the troublesome cases.<sup>11</sup>

### III Counterexamples to (OF)

We have been focusing on the question of how (OF) is to be interpreted. Let us now briefly discuss some of the main putative counterexamples to (OF), thus interpreted.

The first important class of alleged counterexamples to (OF) involves some kind of past wrongdoing or negligence that bears on current feasibility. Perhaps the clearest cases are those where agents intentionally act or refrain from acting so as to ensure that their performing certain future acts is infeasible.<sup>12</sup> Walter Sinnott-Armstrong presents a fairly standard scenario:

Suppose Adams promises at noon to meet Brown at 6:00 p.m. but then goes to a movie at 5:00 p.m. Adams knows that, if he goes to the movie, he will not be able to meet Brown on time. But he goes anyway, simply because he wants to see the movie. The theater is 65 minutes from the meeting place, so by 5:00 p.m. it is too late for Adams to keep his promise. Consequently, if 'ought' entailed 'can', it would not be true at 5:00 that Adams ought to meet Brown. Similarly, if Adams is still at the theater at 6:00, he cannot then meet Brown on time. Consequently, if 'ought' entailed 'can', it would not be true that Adams ought to meet Brown.<sup>13</sup>

---

<sup>11</sup> Perhaps there are further counterexamples that will require additional amendments. Perhaps we shall be forced to conclude that no interesting analysis of "feasibility" can be given, that it is simply a brute primitive. Nonetheless, we shall assume that it is correct in what follows.

<sup>12</sup> Other cases include those where agents *knowingly* but not intentionally, or *negligently* but not knowingly act or refrain from acting in such a way that it is infeasible that they perform certain future acts.

<sup>13</sup> Sinnott-Armstrong 1984, p. 252.

Sinnott-Armstrong takes these implications to be sufficiently “counterintuitive” that “we must give up on the claim that ‘ought’ entails ‘can’”.<sup>14</sup> By “can” Sinnott-Armstrong means “possibility”. Since “feasibility” entails “possibility”, it follows that we must also give up on (OF). Sinnott-Armstrong’s reason for thinking that (both at 5:00 p.m. and at 6:00 p.m.) Adams ought to meet Brown at 6:00 p.m. is presumably the familiar one that to deny this would be to let Adams off the hook way too easily. If at 4:55 p.m. Adams knows that he ought to meet Brown at 6:00 p.m. (in virtue of having promised), and the only thing that has changed between 4:55 p.m. and 5:00 p.m. is that he has knowingly ensured that it is now infeasible that he meet Brown at 6:00 p.m., then to hold that he no longer ought to do so would seem to be a spectacular case of rewarding moral wrongdoing.

It might be thought that there is a possible amendment to (OF) that is capable of responding to such cases. As it stands, (OF) is most naturally interpreted as

(OF)’ A ought, at time  $t$ , to  $X$  in  $C$  only if it is feasible, at  $t$ , that  $A X$  in  $C$ .

Suppose, however, that instead we interpret (OF) as

(OF)’’ A ought, at time  $t$ , to  $X$  in  $C$  only if there is a time  $t^*$  such that, at  $t^*$ , it is feasible that (at time  $t$ )  $A X$  in  $C$ .

This amendment is suggested by Frances Howard-Snyder for (OC).<sup>15</sup> The advantage of (OF)’’ is that it is not susceptible to Sinnott-Armstrong’s counterexample since there are times at which it is feasible that Adams meets Brown, namely all the times up to and including (but not after) 4:55 p.m. So (OF)’’ does not imply that at 5:00 pm and at 6:00 p.m. it is not the case that Adams ought to meet Brown at 6:00 pm. And similarly for any analogous case.

---

<sup>14</sup> Sinnott-Armstrong 1984, p. 252.

<sup>15</sup> Howard-Snyder 2006, p. 235.

Even so, we do not find this amendment plausible. It seems to us to be an important feature of (OF) that it can explain certain changes in what we ought to do that result from changes in what is feasible. Thus, for example, consider a revised version of the Adams/Brown story in which, at 4:55 p.m., just as he is leaving to meet Brown, Adams slips and breaks his leg, making it infeasible that he meets Brown at 6:00 p.m. Under these circumstances, it seems right to say that it is no longer true that he ought to meet Brown at 6:00 p.m. However, (OF)'' does not give this result, since both before and after breaking his leg, there was a time, namely before breaking his leg, when it was feasible that Adams meet Brown at 6:00 p.m. The proponent of (OF) should therefore not respond to Sinnott-Armstrong's argument in this way.

It seems to us that the proponent of (OF) should respond instead by insisting that there is nothing "counterintuitive" about denying that, say, at 5:45 p.m., Adams ought to meet Brown at 6:00 p.m., once that claim is properly understood. To say that "Adams ought to meet Brown at 6:00 p.m." is ambiguous between

- (1) It ought to be the case that Adams meets Brown at 6:00 p.m; and
- (2) Adams ought to *do* something, i.e. go and meet Brown at 6:00 p.m.

Denying (1) at 5:45 p.m. may or may not strike us as "counterintuitive". However, the proponent of (OF) is not committed either way. (OF) is not meant to cover instances of "what ought to be the case". What about (2)? The proponent of (OF) is indeed committed to denying that, at 5:45 p.m., Adams ought to go and meet Brown at 6:00 p.m. Is this "counterintuitive"? On the contrary, it strikes us as pure commonsense. Suppose that, at 5:45 p.m., Adams were asking you for advice about what he ought to do. How would you advise him? You'd probably start by giving him a good scolding and advise him not to act in the future as he has done today. Then you'd advise him to do something else. You might advise A to call B on his

mobile phone and explain that A is not going to make the appointment, apologise profusely but indicate that B should not wait for A. Or perhaps A should go to the meeting place anyway in the hope that B is still there and apologise. Or you might advise A to go around to B's house the next day with a bottle of wine and a box of chocolates. What you plainly would not advise A to do is to go and meet B at 6:00 p.m.<sup>16</sup>

A second class of counterexamples do not involve any past wrongdoing whatsoever on the part of the agent. Here is a case due to Michael Stocker:

Consider a man who correctly believes that ... he ought to pick up the broken glass on a path taken by barefoot children. If some of the children lasso him and tie him to a tree, he can certainly complain that they are making it impossible for him to do what he ought to do ... But if 'ought' does imply 'can', whatever one is prevented from doing cannot be what one ought to do ... But it is a commonplace that we are often so prevented.<sup>17</sup>

It may be tempting to accuse Stocker of trading upon the discrepancy between the nuanced and complex obligations that we actually have and the less nuanced and more straightforward way that we sometimes *talk* about those obligations. Thus, we might say that "you ought not to lie". But unless we are Immanuel Kant we do not intend to be taken so literally as to rule out lying to a would-be murderer about the whereabouts of his would-be victim. Rather, we mean to refer to the fact that we have an obligation not to lie in certain circumstances; where the circumstances in which we find ourselves are relevantly different, the obligation does not apply. Similarly, it might be said that when we say that Stocker's man "ought to pick up the

---

<sup>16</sup> One might make the point more vividly by considering a version of the example where, for no other reason than to avoid meeting Brown, Adams kills himself at 3:00 p.m. Under these circumstances we would surely not say, at 4:00 p.m., that "Adams ought to go and meet Brown at 6:00 p.m.", though we might say that "it ought to be the case that Adams meets Brown at 6:00 p.m.".

<sup>17</sup> Stocker 1971, p. 312.

broken glass on the path taken by barefoot children” we are speaking sloppily of a far more nuanced obligation that he has in certain circumstances. And presumably being tied to a tree represents a prime candidate for a departure from the circumstances in which the obligation applies.

The problem with this response is that it does not allow us to vindicate Stocker’s intuition that the children are “making it impossible for him to do what he ought to do”. For, by hypothesis, it is not the case that he has an obligation to pick up the glass when he is in the circumstance of being tied up to a tree. But, in that case, the children cannot be making it impossible for him to do what he ought to do. At most, they are making it impossible for him to do what he would have had an obligation to do, had they not made it impossible. If this is coherent, which is not obvious, it does not seem to capture accurately the actual phenomenology of the case. If we think that phenomenological accuracy is an important way to resolve such disputes, then some other response seems required.

A better response, it seems to us, is to point out that the obligation that the man has to pick up the glass is a *dissolved pro-tanto* obligation. The test for whether an obligation is dissolved is whether the agent who fails to act in accordance with it owes a justification and apology to those who have been relevantly adversely affected. If the agent owes a justification and apology, then this means that the obligation has not been dissolved (though it may or may not have been outweighed). By contrast, if the agent does not owe a justification and apology, then this means that the obligation has been dissolved. Suppose, then, that while Stocker’s hero is tied up to the tree a small boy cuts his foot on glass left on the path. Does the man owe a justification and apology – say, to the boy or his parents? Surely not. It may be thoughtful to offer them an explanation. But clearly no justification is called for. Being tied up against one’s will uncontroversially dissolves the man’s obligation to pick up the glass from the path.

This means that Stocker's case represents no counterexample to (OF). As we saw above, (OF) only applies to non-outweighed, non-dissolved obligations. Moreover, it means that the proponent of (OF) can allow that the man legitimately complain that the children make it impossible for him to do what he ought to do, where this is interpreted as meaning that they make it impossible for him to act in accordance with his *pro tanto* duty to pick up the glass from the path.<sup>18</sup>

In these examples, the anxiety does not seem to be that the (OF) principle suggests the wrong course of action. The recommendation that A actually meet B at 6:00, or that Stocker's hero actually pick up the glass while tied to the tree – actions that are by stipulation impossible – cannot be read as action-guiding. No-one seems to have any anxiety about the claim that A should avoid doing the impossible. The anxiety is rather that the (OF) principle invites a means of the agent's finessing moral responsibility or that it fails to provide a vocabulary for intuitions about moral responsibility that we would like to be able to register.

We think that these anxieties are reasonable enough; but are not appropriately handled as an attack on (OF). We think they are better handled in another way. The idea is that issues of moral responsibility are properly thought of in terms of A's *attitudes* to the actions she takes (and fails to take). Agent motivations are, of course, influenced by such attitudes – but attitudes have a more extensive life. We can retain (OF) as an action-guiding rule, and accommodate issues about appropriate attitudes

---

<sup>18</sup> One can easily construct analogous cases where the obligation is not dissolved but outweighed. Thus, suppose that the man lives in a war-zone, where, in addition to glass on the path taken by the children, there are also sometimes unexploded shells. Given that he is the only one with the relevant expertise, he believes that he has a *pro tanto* obligation to pick up the shells. Where he must choose between picking up the glass and picking up the shells, he believes correctly that the obligation to pick up the shells outweighs the duty to pick up the glass. Now suppose that the children tie him to the tree so that he is unable to pick up either the glass or the shells. It may still be right that he can complain that the children are making it impossible for him to do what he ought to do, namely pick up the glass, but only where this is understood as an *outweighed* *pro-tanto* obligation.

in a more direct way. That at least is our proposal. We attempt to explore it in the remainder of this paper.

## IV Oughts and Attitudes

The key to our proposal involves focusing attention away from how feasibility considerations impact upon what we ought to *do* and onto how they impact upon how we ought to *think* about what we do – away, that is, from the *actions* that we ought to perform and to avoid performing and onto the *attitudes* we ought to have concerning the actions we perform and don't perform.

In drawing this distinction it will be useful to offer a number of more general observations about attitudes and actions. In particular, our discussion depends on two key claims:

- 1) Having one's attitudes in order matters as well as what one does; attitudes are subject to normative requirements just as actions are.
- 2) What we ought to do and how we ought to think about what we do sometimes come apart.

Let us take each of these claims in turn.

### 1) Attitudes are subject to genuine oughts

Normative analysis deals in part with the actions we ought to perform: whether we ought to save the one or the five, whether we ought to give to charity, whether we ought to take out insurance, and so on. However, actions are not the only things that are subject to normative appraisal. So, too, are our attitudes: our beliefs and desires, motives, dispositions and character traits, and so on. We think that we ought to believe in accordance with the evidence and hence, perhaps, that creationism is false



(epistemic “oughts”) and not to believe logical contradictions (logical “oughts”). We think that we ought to experience guilt and remorse when we have acted wrongly, and pity and compassion when we see others suffering (moral “oughts”). We think that we ought not to be paralytically fearful of strangers or monumentally angry when the bus is running two and a half minutes late (rational “oughts”).<sup>19</sup>

We think these things in part for action-related reasons. Believing what is in fact the case, and having morally approvable desires, are usually justified on the grounds that they lead to the right actions. But we do not think that this exhausts the reasons for having right attitudes. Rather, we submit that having the right attitudes is something of worth *in its own right*. In other words, we want to reject any simple-minded ‘action-reductionism’ in relation to attitudes. In fact, we want to reject even more complex action-reductionist moves – say, dispositional<sup>20</sup> or esteem-related variants<sup>21</sup>. We do not deny that these more complicated versions of ‘action-reductionism’ make the thesis more plausible. But we think getting one’s attitudes ‘right’ makes claims on one’s moral attention beyond the reach even of these considerations.<sup>22</sup>

---

<sup>19</sup> Some theorists have denied that attitudes are subject to genuine “oughts” on the grounds that most attitudes are not under our voluntary control. We cannot – so the line goes – believe, or desire, or experience guilt or remorse or moral satisfaction *at will*. In order for genuine “oughts” to apply to attitudes, those attitudes would have to be “sufficiently under voluntary control to render [deontic] concepts ... applicable to them” (Alston 1989, p. 118). This is not the place to attempt to unpack all the errors we see this claim as containing. An attempt at that exercise has been made by one of us in a different collaboration elsewhere (see Chuard and Southwood 2006). But in our view the “ought-implies-voluntary-control” principle on which this line depends is itself highly questionable as a general conceptual claim. And further we deny that agents’ attitudes lie as far beyond their influence as the critics suggests.

<sup>20</sup> The dispositional variant is that A ought to have a given attitude not because it yields the right action on this specific occasion but because it gives right actions in general over the long haul. Additional arguments would be required to indicate why attitudes were such that they could not be more finely grained to deal with each case on its merits – but there are certainly arguments available that have this effect. [Pettit & Brennan (1986) for example].

<sup>21</sup> The argument here is that one ought to have the “right” attitudes because people respond to those attitudes in their quest for esteem – and their actions designed to avoid disesteem. Brennan & Pettit (2004).

<sup>22</sup> In particular, for example, we are sceptical of examples in which ‘appropriate’ attitudes towards, say, two paintings are to be induced from response to counterfactual scenarios in which buildings are burning and only one painting can be saved. One way of putting our point would be to observe that

## 2) The Autonomy of Attitudes

The idea that attitudes have intrinsic value means that attitudes are morally relevant in situations in which there is no action at all at stake – attitudes say towards the weather; or how beautiful the bride looks in her new white dress; or the truth (or otherwise) about some matter of purely intellectual interest. But it also means that when action *is* at stake, there is more than the action to be considered. This is the sense in which how what we ought to do might come apart from what we *think about* what we ought to do. Our attitudes to our actions do double duty as it were. They influence our actions – but they might in principle have independent value and this independent value might invoke different considerations.

To take a simple example, consider one of the standard cases of ‘moral dilemmas’ – say, the famous case of ‘Jim and the Indians’ offered by Bernard Williams (or the analogous case described by John Fowles in *The Magus*). The characteristic feature of these tragic choice examples is that A is confronted with options both of which involve doing something that is abhorrent – killing one Indian, or bringing about the death of ten, in Jim’s case. The thought here is that what Jim *ought to do* may be different from what Jim *ought to feel* about what he does. Though Jim chooses the right action, it may be appropriate that Jim feels guilty about what he does. Failure to feel guilt in either case might just be ‘inappropriate’: anyone who did not feel some measure of guilt (or something like it) would reveal herself to be a moral monster. And this mismatch between action and attitude raises the logical possibility that the course that involves the lesser guilt might nevertheless be the worse course of action. We do not of course assert that this is the case in Jim’s dilemma or any

---

‘feasibility constraints’ change as we shift from an afternoon’s gallery visit to a conflagration, and we should not expect that ‘appropriate attitudes’ would be insensitive to such changes.

other specific case. We simply claim that this is a logical possibility that cannot be ruled out.

A related set of examples crops up when the best action that A can take requires conditions of imperfect information. If A believed what is in fact the case he would be led to an action that would be worse. In such cases what is the appropriate attitude for A to have towards the false proposition on the basis of which he acts is, at best, ambiguous. It is on one reading inappropriate for A to have an attitude of belief to what is false: but A is led to the better action by virtue of that attitude. Whether the false belief is in that event justified *tout court* depends on the size of the loss involved if A chose the wrong action – but in cases where the action stakes are large, the false belief might be better. Equally, where the action stakes are small, the issue of right attitude becomes decisive.

A specific case of ‘false belief’ arises in relation to what we might term the aspirational quality of moral judgement. Suppose X is some moral ideal. Suppose it is in fact infeasible. Suppose though that believing X to be feasible induces agents to work harder to achieve the value embodied in the ideal: people get closer to X under the false belief that X is feasible, than under the correct belief that X is infeasible. It is, we might suppose, just a fact about the moral psychology of ordinary mortals that ideals have a special appeal to us and engage in us a special motivating force. So, for example, what we give to the suffering poor in the Third World is a function of our beliefs about the feasibility of eliminating poverty. And if it is a better thing to give more, then having the associated false belief is a good thing. At least, it is a good thing on action-related grounds. Of course, in some cases (economists are inclined to think lots of cases) holding false beliefs will lead one to make serious (moral) mistakes – aspiring to the impossible can induce a kind of moral hernia. But there is no reason to assume that all cases are like this – and indeed it seems very unlikely that they will be.

John Broome tells the story of the sick man whose rate of recovery is a positive function of his degree of optimism. If he believes he will get better in  $n$  days, he will actually get better in  $n + 1$  days ( $n > 1$ , say). So if he believes he will get better tomorrow, he will get better the day after tomorrow and that is the best he can do. False beliefs are to his advantage.

We do not intend here to be saying anything too tendentious here. Indeed the thought is platitudinous. The team that believes it can win keeps on trying. The idealist is less likely to give up in the face of adversity than the realist – and *a fortiori* than the hardened cynic. The focus within moral philosophy on ‘ideals’ – salient attractors, with magnetic qualities – can presumably be explained in part<sup>23</sup> by this kind of reasoning.

## V Attitudes and Feasibility

A clarification is called for. The focus of interest here is the attitudes we ought to have towards actions that are infeasible, not the feasibility of the attitudes themselves.

Consider (OFa):

(Ofa) A ought to have attitude  $\varphi$  towards A's Xing in C only if it is feasible that A have attitude  $\varphi$  towards A's Xing in C.

We certainly have no gripe with (OFa), but it is not the focus of our attention here. For (OFa) says nothing about whether A's X-ing in C is feasible or not. For all that

---

<sup>23</sup> The other possibility is of course that moral philosophers tend to get captured by the attractions of their own theories.

(OFa) says, it is entirely open-ended whether X's infeasible has any implications at all for the attitudes A ought to have towards her X-ing.

Further, consider (OFb):

(OFb) B ought to have attitude  $\phi$  towards A's not X-ing in C only if it is feasible for B to have that attitude towards A's not X-ing in C.

Equally, (OFb) says nothing about whether A's not X-ing is attributable to X being infeasible for A in C or not. (OFb) does not rule out the "oughtness" of *our* having contempt for A's failing to X on the grounds that X is infeasible for him; just as (OFa) doesn't rule out the "oughtness" of A having contempt for his *own* failure to X merely on the grounds that X is infeasible in C.

We shall want to come back to the question of how to think about feasibility in attitude space but we should note that the wedge driven between attitude and action permits the relevant players to maintain the right kinds of attitudes in cases where X is infeasible. So if X is infeasible because A culpably omits to take prior action that A ought to have taken, then A's properly feeling guilty about that omission is not ruled out. And if X is undesirable, and infeasible only because of some prior action that A undertook, then A might take moral satisfaction from his not-Xing. And if A feels deep regret – even guilt -- because he was not able to prevent playing children from cutting their feet on broken glass (through no fault of his own, as we might put it), then that is entirely permissible. Perhaps we as external observers will regard any guilt as misplaced; or perhaps we observers will both regard that guilt as appropriate and also sympathise with A that he feels so wretched. The infeasibility of X does not rule out any of these possibilities.

Of course, one might feel that the separation of attitude from action in this way is excessively liberal. But we are not saying (and do not believe) that action and proper attitude toward action are entirely disconnected. We are simply making the more modest claim that attitude toward action and infeasibility of that action are not hard-wired in the way that some interpretations of the “ought implies can” nostrum would suppose. The (OF) rule properly interpreted does not require any such logical connection. And like some critics we think that, if it did, that would be bad news in some cases.

So what should our attitude be to failure to do X when X is infeasible? Well, that is a complicated normative question. But it is a question not about action so much as about attitude: a matter, that is, of getting our attitudes ‘right’. In cases where those attitudes have feedback effects on actions, then the issue of right attitudes will require consideration of the actions to which those attitudes give rise – actions say by the attitude-holder and by third parties who are esteem-seekers (which may well be pretty much everyone). But consideration of the desirability or otherwise of those actions will not exhaust the normative assessment of attitudes. And where those attitudes are themselves normative, it will be important for its own sake to get them right or ‘apt’ (to use Gibbard’s terminology<sup>24</sup>).

Finally a few brief thoughts about feasibility in relation to attitudes. In the first place, we think there is a difference between actions and attitudes in this respect. When we make choices over actions, there is always an opportunity cost – there is always something that is not done by virtue of doing what one does. X rules out Y – you cannot have both. If you could have both, A would not have to make a *choice*. In this sense, choice is comparative – and moral choices no less than non-moral ones. It is for this reason that economists are always inclined to insist on choice calculus being *comparative*: “compared to what?” is a natural economic question. And according

---

<sup>24</sup> Gibbard 1990.

to John Broome (and us, for that matter), this is one of the things that philosophers ought to learn from economists. This is an important part of what motivates Broome to think that moral method should dispense with “goodness” in favour of “betterness”.<sup>25</sup> When it comes to action-guiding considerations, we think that this is right.

But the situation is not so clear in relation to attitudes. When we are invited to assess a piece of music or a poem, we do not immediately think “compared to what?”. We are inclined to speak of the affective response that the music or the poem draws out of us. So we say of it that it was “noble” or “prosaic” or “compelling”. Doubtless we do so on the basis of ‘standards’ that are gleaned from long exposure to other pieces of poetry or other musical works – but there is no natural alternative, no status quo, against which the work in question is to be assessed. Comparators might assist in coming to a fine-grained judgement – but they do not pop out from niggardly nature in quite the way that action-alternatives do. We might say that in action choice certain options (*X* and its alternative *Y*) are ‘infeasible’. It is that infeasibility that induces us to count the cost of action in the way that economists insist we should. But to insist that an evaluation of Schubert’s String Quintet must proceed against the backdrop of a similar evaluation of, say, the Debussy Quartet strikes one as a kind of comparativist fetish. There is nothing to indicate that our attitude to both cannot be one of total admiration and deep affection: perhaps both have that same capacity to speak beauty to something in the core of us. If they do, there is nothing naturally competitive about the emotional response that each of them draws – they stand as similar instances of the same quality and the insistence that we say of them which is the better seems bizarre and perhaps obnoxious. To hold the works in mind together may be to augment our response to each: it is not intrinsically rival, in the manner that action-choice is.

---

<sup>25</sup> See Broome 1999, ch. 10.

If this is so, attitude-guiding has a different quality from action-guiding. In the latter case, goodness may be “fully reducible to betterness” – or at least betterness may be the ‘main game’. But in the former, we can develop an attitude that is appropriate in itself – doubtless, more apt than other attitudes we might have – but not requiring direct reference to what is feasible in the way that action-contexts demand. Of course, it may be the case that there are certain attitudes that we simply cannot entertain – ones that we cannot hold however apt, or however conducive to desirable action, they may be. We do not need to deny that feasibility constraints might bite in the case of attitudes to argue that they bite in the case of actions in a different and more insistent way.

## **VI Conclusions**

What is the upshot of all this?

If (OF) is right, then feasibility considerations have a kind of trumping power in the case of determining what actions we ought to perform. However, in the case of how we *think* about what we do, feasibility considerations play a more complicated role. In particular, nothing in (OF) compels us to regard infeasibility as performing an exculpatory role in relation to action “failures”.

This observation goes some way perhaps towards taking the sting out of certain familiar worries concerning the feasibility requirement. More generally, the distinction between attitude-guiding and action-guiding rules seems to us to be of considerable significance in its own right. And we suspect that that distinction may be of special significance in facilitating the conversation between economists and moral philosophers – a conversation in which anxieties about feasibility considerations play a central role.



## References

- Alston, W. (1989). The Deontological Conception of Epistemic Justification. Pp. 115-52 in *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press.
- Brennan, G. and P. Pettit (2004). *The Economy of Esteem*. Oxford: Oxford University Press.
- Broome, J. *Ethics Out of Economics*. 1999. Cambridge: Cambridge University Press.
- Chuard, P. and Southwood, N. (2006). Doxastic Norms Without Voluntary Control. Unpublished paper.
- Copp, D. (2002). "Ought" Implies "Can", Blameworthiness, and Alternate Possibilities. Pp. 265 in D. Widerker and M. McKenna (eds), *Freedom, Responsibility, and Agency: Essays on the Importance of Alternative Possibilities*. Aldershot, U.K.: Ashgate Press.
- Goodin, R. (1982). *Political Theory and Public Policy*. Chicago and London: Chicago University Press.
- Howard-Snyder, F. (2006). "Cannot" Implies "Not Ought". *Philosophical Studies*, 130, 233-46.
- Kant, I. (1977). On the Common Saying: This May Be True in Theory But It Does Not Apply In Practice in H. Reiss (ed), *Kant's Political Writings*. Cambridge: Cambridge University Press.
- Pettit, P. and Brennan, G. (1986). Restrictive Consequentialism. *Australasian Journal of Philosophy*, 64 (4), 438-55.
- Saka, P. "Ought" Does Not Imply "Can". *American Philosophical Quarterly*, 37, 93-105.
- Sinnott-Armstrong, W. (1984). "Ought" Conversationally Implies "Can". *The Philosophical Review*, 93, 249-61.
- Smith, M. (1995). Internal Reasons. *Philosophy and Phenomenological Research*, 55, 109-31.
- Stocker (1971). "Ought" and "Can". *Australasian Journal of Philosophy*, 39, 303-16.