



The role of all things considered judgements in practical deliberation

Edmund Henden

To cite this article: Edmund Henden (2006) The role of all things considered judgements in practical deliberation, *Philosophical Explorations*, 9:3, 295-308, DOI: [10.1080/13869790600815814](https://doi.org/10.1080/13869790600815814)

To link to this article: <https://doi.org/10.1080/13869790600815814>



Published online: 01 Dec 2008.



[Submit your article to this journal](#)



Article views: 155



[View related articles](#)



Citing articles: 1 [View citing articles](#)

THE ROLE OF ALL THINGS CONSIDERED JUDGEMENTS IN PRACTICAL DELIBERATION

Edmund Henden

Suppose an agent has made a judgement of the form, 'all things considered, it would be better for me to do a rather than b (or any range of alternatives to doing a)' where a and b stand for particular actions. If she does not act upon her judgement in these circumstances would that be a failure of rationality on her part? In this paper I consider two different interpretations of all things considered judgements which give different answers to this question, one suggested by Donald Davidson, the other by Paul Grice and Judith Baker. I argue that neither interpretation is adequate. However, a third interpretation that combines features of the Grice/Baker view with the Davidsonian view is possible. In the final section of the paper I defend this interpretation against two objections.

KEYWORDS practical irrationality; defeasibility; second-order reasons; Davidson

Introduction

We sometimes make judgements of the form, 'all things considered, it would be better for me to do *a* rather than *b* (or any range of alternatives to doing *a*)' where *a* and *b* stand for particular actions. Let's call such judgements 'all things considered judgements' or 'ATC judgements' for short. Suppose an agent has reached an ATC judgement in her practical reasoning, and now is the time to act. The agent has not changed her mind and no new information has been made available to her. If she does not act upon her ATC judgement in these circumstances would that be a failure of rationality on her part?

In this paper I examine and criticize two different interpretations of ATC judgements which seem to give different answers to this question. While one is due to Donald Davidson who argues that ATC judgements are conditional in the same sense as probability judgements are thought to be conditional, the other is due to Paul Grice and Judith Baker, who in a well-known criticism of Davidson's account of weakness of will argue that ATC judgements cannot be conditional in this sense, that in fact they must be *unconditional* judgements. The discussion between Davidson and Grice/Baker is relevant to a larger issue that has drawn some attention among philosophers recently. One general view in the study of rationality has been that a set of practical reasons that is recognized as such by a rational agent must prompt her to action on pain of practical irrationality, unless it is countered by stronger or equally strong opposing reasons. This is a view that has been defended by many authors, and which is implicit in Grice and Baker's account of

ATC judgements as unconditional.¹ Thus, on their view, an agent who makes an ATC judgement has, by definition, an unopposed and conclusive reason for action, one that defeats all opposing reasons, and an agent who recognizes it as such must act on it, on pain of practical irrationality. Read in this context, Grice and Baker's objection to Davidson's account of weakness can plausibly be taken as an argument for the more general view. Davidson rejects this objection but fails, in my opinion, to give any convincing reasons for this response. In fact, in some respects Grice and Baker's interpretation of ATC judgements *does* seem superior to Davidson's own. Does this mean that we should accept their view that ATC judgements are unconditional? If we do, we shall also have to accept that ATC judgements are rationally compelling in the sense that not acting on them implies practical irrationality. In this paper I re-examine the discussion between Davidson and Grice/Baker in the light of recent contributions in the study of practical rationality. I argue that although Grice and Baker's interpretation of ATC judgements in some respects is superior to Davidson's own, there is a way of combining it with the view that they are conditional judgements. However, this strategy comes with a cost: one has to abandon Davidson's *principle of continence*, according to which it is always the case that one should do what one judges is best all things considered. The structure of the paper is as follows. In the first section I present Davidson's view of ATC judgements, discuss Grice and Baker's objection and consider some possible responses to this objection, including one that has been suggested by Davidson himself. In the second section I suggest an alternative response to this objection, one that combines features of Grice and Baker's interpretation of ATC judgements with the view that they are conditional judgements. In the final section I defend this interpretation against two possible objections.

Davidson's View of All Things Considered Judgements

The central importance of the notion of ATC judgements appears especially clearly in connection with weakness of the will. The weak-willed agent is claimed to freely and deliberately perform a particular action *a* against her judgement that some incompatible action *b* would be better *all things considered*. The reason why this description has seemed paradoxical to many is that free, deliberate action traditionally has been thought to correspond to the agent's ATC judgement that performing that action is better than performing any alternative to it, a view that appears to rule out the possibility of a mismatch between the agent's action and her ATC judgement. Davidson accepts the view that free, deliberate action always corresponds to a better-judgement, but argues that the paradox can be avoided if we distinguish between conditional and unconditional better-judgements (see Davidson 1980, 21–42). On his view, the weak agent makes the conditional ATC judgement that doing *a* is better than doing *b* relative to the total set of relevant reasons available to her, but instead of moving on to make this judgement *unconditionally*, she makes the unconditional judgement that doing *b* is better than doing *a sans phrase*, thereby violating the principle of continence, which says that one should always do what one concludes is best relative to one's total set of available relevant reasons. Because the weak agent's ATC judgement is conditional on the total set of relevant reasons available to her, while the judgement corresponding to her intentional action is unconditional, she is not entertaining a logical contradiction.

Before considering an objection to this view of ATC judgements, let me say something more about the terms 'conditional' and 'unconditional'. I should emphasize that in

what follows I will be concerned with these terms only as they are understood by Davidson. This note of caution is important, because the conditional–unconditional distinction raises various complex issues which cannot be settled here.² However, as Davidson uses these terms, they are intended to apply to the form or logical grammar of evaluative judgements. The idea is, very briefly, that the content of an ATC judgement has a ‘conditional form’. However, since these conditional contents do not satisfy the material implication constraint $p \rightarrow q$, they are not conditionals in the truth functional sense of classical logic. Their form is still conditional in the general (and softer) sense of conditional probability statements; while a conditional probability statement expresses the judgement that the probability of p happening rather than q is *conditional* on (the evidence) e , a conditional ATC judgement expresses that the desirability of doing a rather than doing b is *conditional* on (the reasons) r . Thus, there is a nice analogy between theoretical and practical reasoning (see Davidson 1980, 37).

What does it mean that a judgement is ‘unconditional’? On the Davidsonian view, it means that it is detached from the relativization to the agent’s total set of available relevant reasons. The idea is that in order to act on her ATC judgement, an agent needs to make a further move in her reasoning to an unconditional judgement that, if expressed in words, would be of the form ‘doing a is better than doing b ’, implying that a is better not just in this and that respect, but better *absolutely* or *all-out*. This move is required according to Davidson because, ‘Reasoning that stops at conditional judgements . . . is practical only in its subject, not in its issue’ (1980, 39). Although a crucial claim, it is not one Davidson explains very well. In the next section I shall suggest a way to interpret it that may shed some light on the dispute between Davidson and Grice/Baker.

But how is the move to an unconditional ‘all-out’ judgement to be understood? Davidson is not explicit about this, but perhaps the theoretical case may provide us with a clue. Suppose I believe that all the evidence I possess shows that it is highly probable that smoking causes cancer. From this belief, I may draw the conclusion that it *is true* that smoking causes cancer or, simply, that smoking causes cancer. Now, on one possible view of ‘evidence’, this conclusion does not just express the proposition that the total set of relevant evidence available *to me* shows that it is highly probable that smoking causes cancer. It expresses the belief that the total set of relevant evidence *available* shows that smoking causes cancer. In drawing this conclusion, I eliminate the relativization to my evidence. One suggestion then could be that the move from a conditional ATC judgement to an unconditional ‘all-out’ judgement could be understood in a similar way: in moving from the conditional ATC judgement that a certain action is better than the alternatives to the unconditional judgement that that action is better *sans phrase*, I am moving from the judgement that a is better than b relative to all the relevant reasons that are available to me to the judgement that a is better than b relative to all the relevant reasons available. In other words, just as one eliminates the relativization to one’s evidence when one makes the parallel move in theoretical reasoning (one simply concludes that p is true or is the case), so one eliminates the relativization to one’s reasons when one makes the analogous move in practical reasoning (one concludes that doing a is better than doing b).

Paul Grice and Judith Baker have put forward an important objection to the view that ATC judgements are conditional judgements (1985, 27–49). Let us call it ‘the unconditionality objection’. According to this objection, the most plausible view of what it *means* to say that doing a is better than doing b all things considered is that doing a is better than doing b given the fact that on one’s evidence doing a is better than doing b and that one’s

judgement that this is so is 'optimal'.³ By 'optimal' Grice and Baker mean that one has taken into account all the relevant reasons in one's possession and fulfilled whatever call there was upon one at the time in question to maximize one's possession of relevant reasons (1985, 38). The question they raise is whether ATC judgements, interpreted in this way, are entitled to be classified as conditional in the same sense as probability judgements are thought to be conditional. To be conditional in this sense, they claim, these judgements must be defeasible, that is, they must be such that they may be upset without falsification of any of their premisses.⁴ But consider a case in which the optimality condition obtains. The agent has, in effect, taken into account *all* the relevant reasons in her possession and fulfilled whatever call there was upon her at the time to maximize her possession of relevant reasons; no new information is available to her and she has no reason to change her mind about her judgement. An agent who judges that her judgement that *a* is better than *b* is 'optimal' in this sense must, it seems, also judge that *a* is better than *b sans phrase*, on pain of logical incoherence (see Grice and Baker 1985, 46–47). But if this is correct, how can her ATC judgement be *defeasible*? On the contrary, if she were to consider additional reasons in these circumstances which were to upset her ATC judgement, that would have entailed that this judgement in fact was *not* optimal for her. Grice and Baker conclude that if an ATC judgement is optimal for the agent, it simply cannot be defeasible. And if it is not defeasible, neither can it be conditional in the same sense as conditional probability judgements. If Grice and Baker are correct that ATC judgements for this reason must be unconditional, it follows that an agent who recognizes her judgement as an ATC judgement but fails to act accordingly must be practically irrational.

Of course, if ATC judgements are *unconditional* judgements, the Davidsonian account of weakness is in difficulties. The reason is this: suppose the agent judges that relative to her total set of available relevant reasons, doing *a* is better than doing *b*. If she judges that this judgement is optimal for her, it entails the conclusion that doing *a* is better than doing *b*. However, according to Davidson, the weak-willed agent moves from her ATC judgement to the conclusion that doing *b* is better than doing *a*. But if this is correct, it seems to follow that she is guilty of a very serious *logical* mistake in her reasoning: she is like Lewis Carroll's tortoise who is refusing to move from 'given the fact that *p* and that if *p*, *q*' to '*q*' (Grice and Baker 1985, 40). Since it seems very unlikely that the weak-willed agent is making a logical mistake of this magnitude, Grice and Baker take this to show that something is wrong with Davidson's view. The culprit, as they see it, is the idea that an agent has to make *a further move* in her reasoning after she has reached her ATC judgement. Thus, while on Davidson's view the weak agent is practically irrational, not simply by failing to *act* on her ATC judgement, but more importantly, by failing to *draw* the correct conclusion from it, on Grice and Baker's view she is practically irrational not by failing to draw any conclusion, but by failing to act on the conclusion she has already drawn, i.e. the conclusion expressed by her ATC judgement. Of course, if Grice and Baker are correct about this, we should abandon Davidson's view that free, deliberate action always corresponds to an unconditional judgement that performing that action is better than performing any alternative to it. However, Grice and Baker are happy to accept this consequence. In their view it is far easier to attribute to people a failure to act as they fully judge they should act than to attribute to them a failure to judge what they fully judge they should judge (1985, 41).

Let me begin by considering one possible response to the unconditionality objection, which is to reject the claim that ATC judgements have to be 'optimal' for the agent in the

sense suggested by Grice and Baker.⁵ Using Davidson's notation, let '*pf*' represent a *prima facie* operator, '*r*' the agent's set of reasons, and let what's inside the parenthesis be the content of the agent's judgement.⁶ Then on one account, *S* judges that doing *a* is better than doing *b* all things considered if and only if *S* judges that *pf* (*a* is better than *b*, *r*) and *r* is the total set of relevant reasons now available to *S*. This judgement seems to be defeasible in the required sense; it may be defeated by the agent's belief that there are available relevant reasons she has not considered. However, the trouble is that it also, therefore, is consistent with the agent believing that *r* is only a part of her total set of available relevant reasons, in which case she would not be judging that doing *a* is better than doing *b* all things considered even if the right-hand side of the biconditional should obtain.⁷

Perhaps this problem can be avoided with a simple modification. Suppose *S* judges that doing *a* is better than doing *b* all things considered if and only if *S* judges that *pf* (*a* is better than *b*, *r* & *r* is the total set of relevant reasons now available to *S*). In contrast with the first interpretation, this interpretation rules out the possibility that *S* believes that *r* is only a part of the relevant reasons available to her. However, now the difficulty seems to be that *S* may believe that *r* is the total set of relevant reasons that is now available to her but that this set is inadequate for judging that doing *a* is better than doing *b* all things considered; perhaps she believes that she should make further inquiries or collect more information than what is now available to her. But in that case she would not judge that doing *a* is better than doing *b* all things considered.

This brings us back to the interpretation of ATC judgements suggested by Grice and Baker. According to this interpretation, the agent judges that her judgement is *optimal*. This appears to take care of the counterexamples above since it is part of the optimality condition that the agent must have fulfilled whatever call there was upon her to maximize her possession of relevant reasons, something that rules out that she believes her set of reasons is only part of the total set of relevant reasons or that she should make more inquiries or collect even more information. It seems fair to conclude that the 'optimality interpretation' of ATC judgements is more promising than the alternatives offered above. So, does that mean that we have to accept that ATC judgements are unconditional, as Grice and Baker claim in their objection to Davidson? If we do, we shall also have to accept that making such judgements will prompt an agent to act, on pain of practical irrationality.

In the next section I argue that there are several good reasons why we should reject this view of ATC judgements, even if we accept Grice and Baker's 'optimality' interpretation of them.

All Things Considered Judgements Reconsidered

According to the unconditionality objection, conditional ATC judgements must be defeasible. Since ATC judgements on the most plausible interpretation appear not to be defeasible, it follows that they cannot be conditional in the same sense as probability judgements. So, they must be unconditional. However, the unconditionality objection is flawed. There are good reasons for thinking that ATC judgements in fact *are* defeasible. It follows that even if judgements, in order to be conditional in the required sense, must be defeasible, there are no good reasons for thinking that ATC judgements are not conditional. The unconditionality objection trades on a simple paradox for anyone who wants to argue that ATC judgements are conditional: if *S*'s move from a conditional judgement to an unconditional judgement is defeasible, the conditional judgement cannot have

been a proper ATC judgement. If, on the other hand, the judgement *S* starts with *is* a proper ATC judgement, the move to the corresponding unconditional judgement cannot have been defeasible, in which case it cannot have been a proper conditional judgement. In neither case, therefore, can a proper ATC judgement be both conditional *and* defeasible.

An assumption of the unconditionality objection seems to be that *had S's* conditional judgement been a proper ATC judgement, it could not have been the case that *S* had failed to take into account any reasons that were available to her. However, without any such failure, what reasons could *S* possibly have that could defeat her proper ATC judgement? Presented in this way, everything seems to depend on the notion of 'failure' assumed in the argument. The claim appears to be that only a failure to take into account available reasons could explain the defeasibility of ATC judgements. However, because such a failure is ruled out by the only plausible construal of ATC judgements (the 'optimality construal'), it follows that they cannot be defeasible.

In fact, the impression that only a failure to take into account available reasons could explain the defeasibility of ATC judgements is reinforced by Davidson's own appeal to the principle of continence as a universal principle of practical rationality. According to Davidson, the move from a conditional ATC judgement to an unconditional 'all-out' judgement is governed by this principle, according to which one always should do what one judges is best all things considered. Davidson claims this principle is similar to the principle of total evidence in inductive reasoning, according to which one should always give credence to the hypothesis supported by one's total set of available relevant evidence (1980, 41). If this is the correct view of the step from conditional ATC judgements to unconditional 'all-out' judgements, the consequence seems to be that an agent who fails to act on her ATC judgement must be violating the principle of continence and therefore be practically irrational. Thus, if an agent has available a relevant reason not to act on her ATC judgement, that can only mean one of two things: either she has not taken this reason into account, in which case she cannot have made a proper ATC judgement, or she has taken it into account, but it has been outweighed by her other reasons, in which case acting on it would be a violation of the principle of continence (she would not be doing what she judges is best all things considered).

In order to answer Grice and Baker's unconditionality objection I think a Davidsonian need to abandon the view that practical reasoning is governed by a principle of continence of the sort Davidson suggests. Before explaining how this solves the problem, let me say something about the principle of continence. There have been objections to the view that practical reasoning is governed by such a principle. For example, it has plausibly been argued that the principle of continence makes practical reasoning too demanding since it seems unreasonable that one is rationally required to dredge up all one's reasons every time one is deciding to do something (see e.g., McCann 1998, 221). However, I will suggest another objection. The objection is that it seems plausible that there can be circumstances in which one is rationally justified in not acting on one's ATC judgement. What kind of circumstances could this be? Let me first say something about practical reasons in general. Joseph Raz (1978, 1990) distinguishes between first- and second-order reasons. First-order reasons, he identifies as 'relations between facts and persons', e.g., the fact that it is raining out is a reason for me to take an umbrella (Raz 1990, 22). The example he provides is the fact that it is raining out is a reason for me to take an umbrella. According to Raz, such reasons have a dimension of 'strength' or 'weight' such that when they conflict, one may outweigh the other. However, there are also second-order reasons. A second-order

reason, according to Raz, is any reason to *act for a reason* or to refrain from *acting for some reason*. Such second-order reasons do not strictly speaking conflict with first-order reasons. Instead, if the first-order reasons fall within their scope, they 'exclude' them in virtue of their higher order. For example, a second-order reason to refrain from acting for some reason (Raz calls such reasons 'exclusionary') functions so as to cancel out the practical force of certain reasons, i.e. when an agent has such a reason, the strength of her first-order reasons is irrelevant.

There are, of course, familiar difficulties with the idea of second-order reasons. How do such reasons get their greater normative force? Is it simply in virtue of their being second order? Why should anyone accept that certain considerations exclude the normative force of other considerations which would otherwise enjoy normative force? I cannot settle these complex issues here.⁸ Raz's view is that conflicts between first- and second-order reasons are resolved not by strength of the competing reasons but by a general principle of practical reasoning according to which one should always act for an undefeated reason for action. For the purposes of this paper I shall assume that something like this view is correct.⁹ My concern will be to investigate whether this kind of view of practical reasons can shed some light on the role of ATC judgements in practical deliberation, and more specifically on the dispute between Davidson and Grice/Baker. So, let us assume we can draw a distinction between practical reasons in the way Raz proposes. Returning to Grice and Baker's optimality interpretation of ATC judgements, this seems to give us two possible ways of understanding the phrase 'being optimal for *S*'. On a weak reading, 'being optimal' means that the reasons for and against doing *a* or *b* that the agent has considered, satisfy certain standards, i.e. that they are the total relevant set available, that the agent has done whatever can reasonably be required of him in order to maximize this set, that he believes he has done so, and so on. Perhaps it sounds odd to talk about 'the weak reading' here. What could possibly be *stronger* than including the total set of relevant available reasons for and against doing *a* or *b*, as part of the optimality condition? But a stronger reading is possible. We only need to include *S*'s reasons to decide to act on his first-order reasons. In other words, 'being optimal for *S*' could involve not only the belief that the total set of relevant available reasons for and against doing *a* or *b* satisfies certain standards, but also the belief that the reasons in this set are *sufficient to act* on, i.e. enough to justify acting. It seems plausible that making a decision means that one has such a belief. A decision seems to *settle* matters in one's mind (see Bratman 1987; Raz 1978). Thus, having such a belief seems to rule out, for example, that one regards oneself as having an exclusionary reason to refrain from acting on one's reasons. This suggests that a minimal requirement on deciding to act is that one regards oneself as having no exclusionary reason to refrain from acting on one's first-order reasons.¹⁰

With these two readings of the optimality condition in mind, let us consider a case in which an agent makes an evaluative judgement about what it would be best to do. Suppose Smith is looking for a good way to invest his money and is told by a friend of a possible investment.¹¹ The friend gives Smith a pile of documents that contains all the relevant available information and tells him that he has to make a decision that same evening otherwise the offer will be withdrawn. The fact that the investment gives a good chance of pay-off may be a reason for Smith to judge that accepting the offer would be best, while the fact that it carries a small risk may be a reason for him to judge that declining it would be best. These reasons will be of the first order since they are relations between himself and certain facts, i.e. the fact that the investment offers a chance of pay-off, that

it carries a small risk, and so on. Suppose Smith considers all the relevant reasons for and against these two alternatives and on that basis judges that accepting the offer would be best. If we assume that his judgement that this is so is optimal for him in the weak sense, we can conclude that Smith has made an ATC judgement in favour of accepting the offer. According to the principle of continence it follows that he will be practically irrational if he does not accept the offer.

Before considering the plausibility of this claim, let us ask how this case should be described if we instead assume a *strong* reading of the optimality condition. Since Smith may not yet have formed the belief that the reasons he has considered are sufficient to act on, he has not yet made an ATC judgement according to the strong reading. Now, let us add the following information to the story: suppose Smith just before his friend told him about the investment had a couple of glasses of wine. Although he may feel only faintly intoxicated, he may still believe this is enough to distrust his own judgement. In this case he has a second-order or 'exclusionary' reason in Raz's sense to refrain from acting on his first-order reasons. Since an ATC judgement, according to the strong reading, must include the belief that his set of reasons is sufficient to act on, and Smith has an 'exclusionary' second-order reason to refrain from acting on his set of reasons, it follows that Smith does not make an ATC judgement in this case. Rather, he is suspending judgement on whether it would be best to accept or decline the offer.

Note first that Grice and Baker's unconditionality objection assumes a strong reading of the optimality condition. That is because it assumes that once the agent has reached his ATC judgement there is no further move for him to make in his reasoning and no further relevant reasons available to him that could possibly defeat his ATC judgement. His ATC judgement is indefeasible and, therefore, unconditional. The only possible explanation, on the strong reading, of why the agent would not act on his ATC judgement in these circumstances would be either through some failure of rationality (as in weakness) or through some failure to take into account some relevant available reasons. If the first were the case, the agent would, of course, be practically irrational. If the second were the case, he would not have made a proper ATC judgement.

Let me start by considering the case assuming a weak reading of the optimality condition. The reason Smith would be practically irrational if he did not accept the offer according to this interpretation is that he would be violating the principle of continence. But it seems wrong that Smith would be practically irrational if he did not accept the offer. In fact, it seems plausible that in this case Smith can *rationaly refuse* to either accept or decline the offer since he has a second-order reason not to act on his first-order reasons. The fact that he feels faintly intoxicated is such a second-order reason. Thus, it is not itself a reason for him either to judge that declining the offer is best or that accepting it is best. Rather, the fact that he feels faintly intoxicated is a reason for him to refrain from acting on the ATC judgement he has already made; it is an 'exclusionary' reason in Raz's sense, that is, a reason to refuse *deciding* to accept or decline the offer. Suppose the friend replies that this violates the principle of continence. In response, Smith may insist that he can rationally refuse to make a decision because he has a reason not to act upon the merits of the case, a reason which is not taken into account by the principle. Although he does not regard his mental state as a reason for judging either that accepting the offer would be best or that declining it would be best, he does regard it as a reason not to act on his best assessment of the reasons bearing on whether it would be best to accept or decline the offer. In general there seem to be many kinds of situations in which an agent, due to

stress, temptation, pressure of time or the influence of strong emotions, may be rationally justified in not acting on his ATC judgement construed in the weak sense.¹² Such cases seem to be counterexamples to the claim that practical reasoning is always guided by a principle of continence of the sort Davidson suggests.

So, what if we instead assume the strong reading of the optimality condition? Assuming this reading, Smith's second-order reason is claimed to be a reason for him to rationally refuse to make an ATC judgement in this case; it is a reason for him to suspend judgement altogether. But it seems wrong that Smith does not make an ATC judgement. Smith considers all the relevant reasons available to him and judges that accepting the offer would be best. The fact that he feels faintly intoxicated is not itself a reason for him to either judge that accepting the offer is best or that declining it is best; it does not make him change the comparative assessment of the reasons he already recognizes, e.g. recalculate their strength. In fact, it does not affect the balance of his reasons at all. Thus, Smith may maintain that based on the total set of relevant reasons available to him bearing on whether it would be best to accept or decline the offer, accepting it would be best. Still, he has a second-order entitlement to discount the practical force of his set of reasons in this case, that is, to refuse *deciding* to accept or decline the offer. This second-order entitlement does not conflict with the reasons he has considered; it does not outweigh them in virtue of its strength. Rather, it excludes them in virtue of its higher order. So, why should it not be included in the considerations which led him to his ATC judgement? The answer is that an ATC judgement, by definition, is a judgement about whether doing *a* is better than doing *b*, while a second-order reason is not about the desirability of doing *a* or *b* at all. It is about the desirability of acting on *one's reasons for judging* that doing *a* is better than doing *b*. Thus, in a sense it is about *the mind*, while the reasons making up one's ATC judgement are about *actions*, or *the world*.¹³ It seems plausible that in order to make a decision to act based on deliberation, a rational agent will have to take into account such second-order reasons. However, it is not clear why she should have to take them into account in order to judge which of her alternatives is best relative to her reasons. In other words, it seems wrong to include second-order reasons in ATC judgements *precisely because* they are second order; they are not reasons for or against doing *a* or *b*. The mistake of the strong reading of the optimality condition, therefore, is to conflate ATC judgements with decisions to act, while these should be kept apart.

Now, if the distinction between first- and second-order reasons is valid, we should accept the weak reading of the optimality condition, but reject the principle of continence as a universal principle of practical rationality. If we do, the reply to Grice and Baker's unconditionality objection will be pretty straightforward: if an agent does not act on her ATC judgement, it need be through no failure of rationality nor through any failure to take into account relevant available reasons. The agent may simply take her ATC judgement to have been defeated by a second-order exclusionary reason. In other words, on the weak interpretation, ATC judgements *are* defeasible and, therefore, conditional in the same sense as conditional probability judgements. This interpretation allows that it can be perfectly rational not to act on one's ATC judgement even if it is optimal in Grice and Baker's sense. It also gives us a way to understand Davidson's claim that 'reasoning that stops at conditional judgements . . . is practical only in subject, not in issue' (1980, 39). This can now be understood in terms of a gap between the agent's recognition of the reasons which make up her conditional judgement, on the one hand, and her practical sensitivity or responsiveness to these reasons, on the other. Thus, although Smith recognizes that

his reasons for accepting the offer defeat his reasons for declining, it need not motivate him to act accordingly since he has a second-order entitlement to discount these reasons in this case. He simply takes himself to lack sufficient reason to act in accordance with his ATC judgement. Unfortunately, Davidson does not explain how such a gap between recognition and motivation can be rational, and it seems difficult to do so as long as the principle of continence is maintained as a universal principle of rationality.¹⁴ One consequence of the weak interpretation of ATC judgements, therefore, may be a need to modify the Davidsonian framework to allow for rational behaviour that is not in accordance with the principle of continence. If this is the correct way to go for a Davidsonian, weakness of will should be seen as a violation of the principle that one should always act on an undefeated reason for action rather than a violation of the principle of continence. In the final section of this paper I shall consider some objections to the view that we should prefer the weak reading of the optimality condition.

Some Objections

The first objection I want to consider is that there is no basis for regarding the weak ATC judgement as a *genuine* 'all things considered' judgement. Clearly, if we have failed to consider what we would ordinarily regard as *all* the relevant reasons, it cannot be correct to characterize the judgement as an 'all things considered' judgement. After all, why should weakly optimal judgements be regarded as all things considered judgements if they do not even take into account the most important thing, i.e. whether *one should act*?

I think this objection fails to appreciate that there is an important distinction to be made between ATC judgements about what is best and decisions to act. I have tried to bring out this distinction by distinguishing between the kinds of reasons involved, but perhaps another way could be by noting that there appear to be two ways in which practical deliberation might fail. First, a deliberator might fail to reach a judgement that performing some action is best simply because her reasons are unclear, indeterminate or both sides seems equally compelling or incommensurable. Still, it seems plausible that she can *decide* to perform that action. Second, a deliberator might fail to reach a decision to perform some action even if she has judged that doing it is best, because she can remain *unsettled* about what to do. Gary Watson has for similar reasons recently argued that when things go as intended, practical deliberation involves making up one's mind *twice*; while making up one's mind about what is best to do is coming to a judgement *that* such and such is the thing to do, making up one's mind about what to do is deciding *to do* such and such (2003, 176). In general there seems to be clear distinguishing marks between ATC judgements and decisions to act. Decisions are not *ipso facto* judgements; judgements are true or false, decisions are right or wrong. ATC judgements also seem to have a conditional form, while decisions to act appear to involve asserting the consequent of conditionals. All these considerations seem to support the view that there is a distinction to be made between ATC judgements and decisions to act.

A second objection to the view that we should prefer the weak interpretation of ATC judgements could be that there is no particular reason to believe that the absence of second-order defeaters explicate the meaning of the phrase 'sufficient reason to act'. There could be many other plausible interpretations of this phrase. For example, in many cases we might consider that there is sufficient reason to do *a* if there is simply

most reason to do *a* and no reason not to do *a*. There is no appeal made here to the absence of second-order defeaters.

First of all, my goal has not been to explicate the meaning of the phrase 'sufficient reason to act'. A systematic analysis of decision might have to provide such an explication. My more modest goal has been to suggest what seems to me *one* necessary condition of making rational decisions in the context of practical deliberation. Whatever 'sufficient reason' may *mean*, it seems very plausible that an agent who takes himself to have such a reason regards himself as having no reason to disregard this reason. Second, however, many authors do in fact argue that decision involves a reflexive element (see, e.g., Bratman 1996; Pink 1996, 5). This may suggest a similar distinction as I have argued for in this paper between ATC judgements and decisions, since ATC judgements do not appear to be reflexive in this sense. Thus, a rational agent can weigh considerations provided by conflicting desires in deliberation and on that basis judge one course of action good or best, without reflectively endorsing the considered desires on which this judgement is based. Whether decision based on deliberation *is* some form of 'reflective endorsement' of one's reasons for judging one course of action best, is a further question that I have not addressed in this paper.¹⁵

To sum up, I have argued that in cases which involve practical deliberation, one's reasons for judging an action best are about actions or the world, while one's reasons for deciding to act are in part about one's *reasons* for judging that action best, i.e. in part about one's own mind. We should, therefore, distinguish ATC judgements which involve only the first kind of reasons from decisions to act, which also involve the second kind. Thus, on this view, ATC judgements do not require action, even in the absence of strong enough opposing reasons to defeat them. If this is correct, there is available a plausible interpretation of ATC judgements according to which they are conditional judgements, practical only in subject, not in issue.

ACKNOWLEDGEMENTS

I would like to thank Bill Child and David Charles for helpful comments on earlier drafts of this paper. Thanks also to two anonymous referees for *Philosophical Explorations*.

NOTES

1. For this way of characterizing the general view, I am indebted to Greenspan (2005). One version of this view is sometimes called 'judgement internalism' and is often associated with Kant. According to this view, a judgement of value provides motivation for a rational agent to act accordingly. Kant applies this view to the moral sphere, where he claims that those who recognize the ability of their maxim to serve as a law for all rational beings are imperfectly rational if they fail to act on it. See Kant (1997). But the view may also apply more broadly to non-moral value judgements. Thus, in setting up his sceptical argument against the possibility of weakness of will, Davidson assumes what he calls 'a mild form of internalism' according to which 'a judgement of value must be reflected in wants (or desires or motives)' (1980, 26). For further details of different kinds of internalism, see Brink (1989, 37–43), Darwall (1997) and Koorsgaard (1986).
2. Other authors who have done important work relevant to this distinction include Edgington (1995), Lewis (1986) and Stalnaker (1994), to name a few.

3. In their own words, to facilitate exposition, Grice and Baker use the form 'x should do *a*' rather than 'doing *a* is better than doing *b*'. Thus, they claim that 'the kind of judgement whose status is in question is one such as "given the fact that on my evidence I should do A, and that my judgement that this is so is optimal, I should judge that I should do A"' (see Grice and Baker 1985, 38). However, for reasons I hope will become clear, I shall stick to the Davidsonian schema and use the relational predicate 'is better than'.
4. It is important to note that the sense of 'conditional' in question here is *not* the truth functional sense of classical logic, that is, not the sense in which a judgement of the form 'if *p* then *q*' is a conditional judgement. Clearly, the latter judgement need not be a defeasible conditional. The sense of 'conditional' Grice and Baker have in mind is the sense in which a probability judgement is *conditional on the evidence*. Being a conditional judgement in this 'softer' sense is usually taken to imply that it can be upset without falsification of any of its premises. It is this view of 'conditional' I shall assume throughout this paper (see Grice and Baker 1985, 39).
5. Davidson's own response to the objection seems to be along these lines (see Davidson 1985, 201–6). Thus, he argues that the objection is harmless because it assumes one particular understanding of ATC judgements, while in fact his own distinction between conditional and unconditional judgements is consistent with many different construals of ATC judgements, some of which (presumably) do not give rise to troubles of the kind Grice and Baker describe. However, Davidson does not provide any plausible alternative account of ATC judgements to back up this claim. Both of the two construals I discuss in the text can, I think, be found in Davidson. Unfortunately, neither of them seem to work against the unconditionality objection.
6. For those unfamiliar with Davidson's notation, some remarks of clarification are in order. Davidson introduces the '*pf*' as a *prima facie* sentential operator. *Prima facie* in English means something like 'in so far as'. The idea is that the '*pf*' makes a complex sentence out of two sentences, one stating that some action *a* is better than some action *b*, the other stating in which respects *a* is better than *b*, that is, relative to what reasons *a* is better than *b*.
7. See also Grice and Baker (1985, 35) for a similar objection to this construal.
8. Raz's distinction has generated much discussion. For some recent contributions, see Pettit et al. (2004).
9. For a recent defence and extension of Raz's distinction, see Bratman (2004). Bratman distinguishes between ordinary reasons and what he calls 'framework' reasons, i.e. 'self-governing policies' of discounting or bracketing certain reasons. Dancy argues in a similar vein for a distinction between what he calls 'enticing' and 'peremptory' positive reasons, the first of which contribute to making actions worth performing, but not to requiring them (see Dancy 2004).
10. Of course, this is not to say that it is not possible to decide to do something without believing that one has sufficient reason to do it. Sometimes we clearly decide to do things even if we believe we have very little reason to do it. Weakness of will is *one* example. The point is that in the cases where we *do* rationally decide to perform some action on the basis of practical deliberation, deciding to perform that action may be seen to involve the belief that our set of reasons is, in the circumstances, sufficient to act on. When I speak of 'decision', I shall have in mind only rational decision in the context of practical deliberation.

11. This example is a variant of a similar example suggested by Raz (1978, 129). However, Raz's concern is not ATC judgements or the principle of continence.
12. Raz also argues that promises and rules can be second-order exclusionary reasons in this sense (see Raz 1978). This claim might be more difficult to defend. For a recent critical discussion of Raz's use of the concept of exclusionary reasons to explain the authority of law, see Mian (2002).
13. I borrow this way of putting the distinction between first- and second-order reasons from Greenspan (2005).
14. It can be added here that the distinction between recognition and motivation also puts pressure on Davidson's mild form of internalism since it shows that any conceptual connection between judging something best and being motivated to do it, must be a defeasible one (see Note 1). In fact, there seems to be a growing awareness of the importance of this distinction in the recent literature. For example, Fisher and Ravizza draws a similar distinction between what they call 'receptivity' and 'reactivity' as forms of 'reasons-responsiveness' (1998, 69). Greenspan defends this distinction in her paper, 'Adequate reason' (2005).
15. For a discussion relevant to this issue, see Frankfurt (2005).

REFERENCES

- BRATMAN, M. 1987. *Intention, plans, and practical reason*. Cambridge, Mass.: Harvard University Press.
- . 1996. Identification, decision, and treating as a reason. *Philosophical Topics* 24: 1–18.
- . 2004. Shared valuing and frameworks for practical reasoning. In *Reason and value: Themes from the moral philosophy of Joseph Raz*, edited by P. Pettit, S. Scheffler, M. Smith, and R. J. Wallace. Oxford: Oxford University Press.
- BRINK, D. 1989. *Moral realism and the foundations of ethics*. Cambridge: Cambridge University Press.
- DANCY, J. 2004. Enticing reasons. In *Reason and value: Themes from the moral philosophy of Joseph Raz*, edited by P. Pettit, S. Scheffler, M. Smith, and R. J. Wallace. Oxford: Oxford University Press.
- DARWALL, S. 1997. Reasons, motives, and the demands of morality: An introduction. In *Moral discourse and practice*, edited by S. Darwall. Oxford: Oxford University Press.
- DAVIDSON, D. 1980. How is weakness of the will possible? In *Essays on actions & events*. Oxford: Clarendon Press.
- . 1985. Reply to Paul Grice and Judith Baker. In *Essays on Davidson, actions and events*, edited by B. Vermazen and M. B. Hintikka. Oxford: Oxford University Press.
- EDGINGTON, D. 1995. Low on conditional probability. *Mind* 105 (420): 617–30.
- FISHER, J. M., and M. RAVIZZA. 1998. *Responsibility and control: A theory of moral responsibility*. Cambridge: Cambridge University Press.
- FRANKFURT, H. 2005. Identification and wholeheartedness. In *The importance of what we care about*. Cambridge: Cambridge University Press.
- GREENSPAN, P. S. 2005. Adequate reason. Unpublished paper, University of Maryland.
- GRICE, H. P., and J. BAKER. 1985. Davidson on 'Weakness of the will'. In *Essays on Davidson, actions and events*, edited by B. Vermazen and M. B. Hintikka. Oxford: Oxford University Press.
- KANT, I. 1997. *Groundwork of the metaphysics of morals*. Translated and edited by M. Gregor. Cambridge: Cambridge University Press.
- KORSGAARD, C. M. 1986. Scepticism about practical reason. *Journal of Philosophy* 83: 5–25.

- LEWIS, D. 1986. Probabilities of conditionals and conditional probabilities. *Philosophical Review* 95: 581–89.
- MCCANN, H. J. 1998. Practical rationality and weakness of the will. In *The works of agency, on human action, will, and freedom*. Ithaca, N.Y.: Cornell University Press.
- MIAN, E. 2002. The curious case of exclusionary reasons. *Canadian Journal of Law and Jurisprudence* 15 (1): 99–124.
- PETITT, P., S. SCHEFFLER, M. SMITH, and R. J. WALLACE, eds. 2004. *Reason and value: Themes from the moral philosophy of Joseph Raz*. Oxford: Oxford University Press.
- PINK, T. 1996. *The psychology of freedom*. Cambridge: Cambridge University Press.
- RAZ, J. 1978. Reasons for actions, decisions and norms. In *Practical reasoning*, edited by J. Raz. Oxford: Oxford University Press.
- . 1990. *Practical reasons and norms*. Princeton, N.J.: Princeton University Press.
- STALNAKER, R. 1994. Conditionals as random variables. In *Probability and conditionals*, edited by Ellery Eells. Cambridge: Cambridge University Press.
- WATSON, G. 2003. The work of the will. In *Weakness of will and practical irrationality*, edited by S. Stroud and C. Tappolet. Oxford: Oxford University Press.

Edmund Henden, Department of Philosophy, University of Oslo, P.O. Box 1020 Blindern, 0315 Oslo, Norway. E-mail: edmund.henden@ifikk.uio.no