

Instrumental Rationality

Markos Valaris

Abstract: Does rationality require us to take the means to our ends? Intuitively, it seems clear that it does. And yet it has proven difficult to explain why this should be so: after all, if one is pursuing an end that one has decisive reason not to pursue, the balance of reasons will presumably speak against one's taking the means necessary to bring that end about. In this paper I propose a novel account of the instrumental requirement which addresses this problem. On the view I develop, the instrumental requirement is normative not because agents have reasons to comply with it, but because it is a *normative standard* intrinsic to intentional action—i.e., it is a standard that partly spells out what it is to exercise one's agency well.

1. Introduction

Our thought and conduct are, insofar as we are rational and sufficiently well-informed, shaped by instrumental facts—by facts about what the means to our ends are. My aim in this paper is to propose an account of how instrumental facts get their normative grip on our thought and conduct.

Suppose that Fred is making an omelet. Moreover, suppose that Fred knows that you cannot make an omelet without breaking some eggs. How do these facts bear on Fred's actions? It is certainly not necessary that Fred will break any eggs: Fred might be too engrossed in his favorite television show to do so. Nevertheless, if Fred does not break any eggs, it is natural to say that, other things being equal, he is being *practically irrational*. Fred's conduct seems to leave him open to a distinctive sort of criticism, precisely in that he fails to take the means necessary to his end.

It is important that the criticism that attaches to instrumental failures appears to be distinctive. Suppose that Mary pursues an end that she has decisive reasons not to pursue. For example, suppose that one of Mary's ends is to smoke a cigarette. Suppose, moreover, that she believes that in order to smoke a cigarette she needs to buy a pack of cigarettes. Given all this, if Mary does not buy a pack of cigarettes, she would seem to be open to criticism in just the same way as Fred. The important point here is that this does not appear to be due to Mary's acting in a way that is inconsistent with what she has most reason to do: by hypothesis, Mary has decisive reasons not to smoke, and so—presumably—the balance of reasons does not speak in favor of her

buying a pack of cigarettes. Examples like this suggest that instrumental failures carry a sort of blame that is independent of the reasons one might have to pursue the corresponding ends.¹

In this paper I am going to argue that this appearance is correct. As I will argue, the instrumental requirement is a *normative standard* that governs intentional action, and which derives from some fundamental facts about the relation between a non-basic action and the means necessary to complete it successfully. What normative standards are, and how their normativity differs from that of reasons, are among the main questions I will discuss in what follows. In a nutshell, however, the idea is this. Reasons, in the most general sense of the term, are considerations that count (either objectively, or just in the agent's own eyes) in favor of something—an action, or an attitude like a belief or intention. Normative standards, by contrast, are standards that determine what it is to exercise a capacity or to perform an activity *well*. Crucially, the normativity of normative standards seems irreducible to that of reasons, as one may be criticized as performing badly in an activity that one did not have most reason (either objectively or subjectively) to engage in at all.² It is clear that all sorts of activities are governed by normative standards, in this sense. One of the claims I will defend in this paper is that intentional action *as such* is also governed by a normative standard: intentional action is a normatively constituted phenomenon.³ The instrumental requirement is normative not because agents always have reasons to comply with it, but because it partly spells out what it is to exercise one's agency well.⁴

It will be helpful to situate my proposal in relation to other recent views on the topic. Recent philosophers have tended to assume that, if the instrumental requirement is a genuine normative requirement, then its normativity must be explained in terms of the normativity of reasons (whether of the objective or the subjective variety). On such views, we must either show that, contrary to appearances, agents always have most reason to avoid violating the instrumental requirement, or we must accept that the instrumental requirement is a myth. My view is that we should question the assumption shared by both these views: we can give an account of the normativity of the instrumental requirement without trying to reduce it to that of reasons.⁵

The structure of the paper is as follows. In Section 2 I argue that recent attempts to explain the normativity of the instrumental requirement in terms of the normativity of reasons fail. If we are to avoid skepticism about the instrumental requirement, therefore, we need an approach that does not assume that its normativity is that of reasons. In Sections 3 and 4 I develop my positive proposal. More specifically, I argue that facts about the metaphysics of processes in general, and of intentional action in particular, show that if one is doing *A*, but one does not take the means necessary to doing *A*, one has not exercised one's agency well. This, as I argue, is the root of the instrumental requirement. Finally, I will close this paper in Section 5 by showing how my account fits in with a plausible phenomenology of instrumental failures.

2. Do We Have Reasons to Take the Means to Our Ends?

The instrumental requirement is one example of what are sometimes called ‘rational requirements’. Rational requirements in this technical sense are constraints of *subjective* rationality, or, in other words, constraints on our thought and conduct which hold even in abstraction from the facts about our situation that constitute the reasons for our attitudes or our actions. As we already saw, for example, it seems natural to say that rationality requires of Mary that she buy cigarettes if she intends to smoke, even though she has decisive reason not to smoke. The question, then, is the following: what—if not the facts that constitute the reasons for our attitudes and actions—explains why we are subject to such requirements? My goal in this section is to argue that recent attempts to answer this question do not seem promising.

According to one popular approach, rational requirements characteristically prohibit *combinations* of attitudes (or attitudes and actions), which are intuitively incoherent. Perhaps then, it is suggested, we are subject to these requirements because we have reasons to avoid such incoherent combinations which are independent of the reasons we might have for or against any particular attitude or action, considered in isolation. This is the *wide-scope* strategy, originally proposed by John Broome (1999).⁶ For example, having incoherent degrees of belief—that is, degrees of belief that violate the axioms of probability theory—can be shown to have bad consequences for one’s decision-making: it leaves one open to accepting bets which one is guaranteed to lose. Thus, regardless of how justified any one of one’s degrees of belief might be considered in isolation, their *combination* is objectionable. If this is correct, then we have a *standing reason* to avoid incoherent combinations of degrees of belief; and we might take this as the ground for a requirement of rationality that instructs us to avoid such combinations. The hope is that this explanation might serve as a model for an account of rational requirements in general.

There is a straightforward difficulty this approach faces. On this approach, rational requirements can have normative force only if the combinations they prohibit can be independently demonstrated to be objectionable—i.e., such that we have reason to avoid them. It is thus a real question whether this approach will be able to account for all—or even *any*—of the rational requirements we intuitively feel that there are. (Indeed, pressing this question is one of the strategies that skeptics about rational requirements pursue.) In particular, as we will see, it is not at all clear that we always have reason to avoid combining an intention with a failure to take (or to intend to take) the means necessary to carry it out.

Suppose, once again, that Mary intends to smoke a cigarette, although she has decisive reasons not to. Suppose also that Mary knows that, unless she buys a pack of cigarettes, she cannot smoke. Now, given that Mary has decisive reasons not to smoke, it seems plausible that she also has very good reasons not to buy cigarettes. Nevertheless, the instrumental requirement tells us at least this about Mary: Mary should avoid combining her intention to smoke with lacking an

intention to buy cigarettes.⁷ On the present approach, this must be a consequence of the fact that Mary has reasons to avoid this combination, which are independent of any reasons she has for or against each of the elements combined. But what such reasons might Mary have?

Of course, Mary has reasons to avoid the combination to the extent that it involves the intention to smoke, which we have assumed Mary has decisive reasons to avoid. But, for present purposes, this is neither here nor there. The suggestion we are evaluating is that combining an intention with a failure to take (or to intend to take) the means to carry it out is objectionable in a way that each of its elements, considered in isolation, is not. But why should this combination be especially objectionable?

One might point to the fact that the combination violates the instrumental requirement, and claim that it is objectionable on this score. But, surely, this cannot be our starting point here: what we are looking for is an account of the normativity of the instrumental requirement; thus we cannot assume at the outset that combinations that violate it are objectionable, on just those grounds. What we need is an argument to the effect that Mary has *independent* reasons to avoid violating the instrumental requirement, i.e., reasons whose existence do not presuppose that the instrumental requirement is a genuine normative requirement.

Michael Bratman has suggested that it is a constitutive fact about agency that agents have a reason to govern their own lives. Moreover, self-governance is undermined by violations of the instrumental requirement (2009: 431):

If I intend end E but I do not now intend known necessary means [. . .] there is no clear answer to the question, 'Where do I stand?' with respect to E. With respect to this end, there is as yet no relevant fact of the matter about where I stand.

Since self-governance requires taking a stand with respect to the options that are open to one, Bratman's thought is that the reason agents have for self-governance—other things being equal—generates a reason to avoid violations of the instrumental requirement.

On this view, if Mary intends to smoke but does not intend to buy cigarettes, then she has not taken a stand with respect to the question whether to smoke or not: in this respect, she has failed to govern herself. This is what, on Bratman's view, makes her combining an intention to smoke with a failure to intend to buy cigarettes objectionable, in a way that neither of the elements combined would be, considered on its own.

Thus Bratman's proposal helps with some cases. However, it is unlikely to help in general, for it is unlikely that the relevant self-governance-derived reasons will always be available. We can see this if we add to the description of our case, by stipulating that Mary *knows that* she ought not to smoke a cigarette. Mary judges her intention to smoke a cigarette to be unreasonable. This is by no means a far-fetched scenario: Mary, like many smokers, intends to smoke even against her better judgment. Mary's situation is a case of *akrasia*. In this case,

would Mary's failure to take the means to her end undermine her claim to self-governance? It is hard to see why it would: Mary is *already* conflicted; it is *already* unclear where she stands with respect to the question whether to smoke or not. It is not at all obvious that violating the instrumental requirement makes her situation any worse with respect to self-governance.

To see this, suppose that Mary asks herself the following two questions: 'Should I smoke?'; and, 'Will I smoke?'. Intuitively, it would seem that the more aligned Mary's answers to these questions are, the less conflicted Mary is. By hypothesis, Mary answers the first question firmly in the negative. Thus the firmer Mary is in answering the second question in the negative too, the less conflicted she is. But, given that Mary believes that buying cigarettes is necessary if she is to smoke, her failure to intend to buy cigarettes would tend to push her in the direction of answering the second question in the negative. Thus it seems that by violating the instrumental requirement Mary becomes less conflicted, not more.⁸ If this is correct, then Mary has *no reason at all* based on self-governance to avoid violating the instrumental requirement. But if this is right, we still have no explanation for why Mary is subject to the instrumental requirement.⁹

The above discussion assumed that the reasons Mary has to avoid violating the instrumental requirement are detectable from an *objective* standpoint, or in light of the considerations that really matter for what Mary should do. But some proponents of the view that the instrumental requirement derives from reasons would not want their thesis to be understood in this way. They intend their thesis to be understood as a claim about *subjective* reasons—i.e., reasons derived from the agent's own subjective assessment of her situation, incomplete and error-filled as it might be. Thus, the suggestion might be that, although from an objective point of view Mary's combining an intention to smoke with not intending to buy a pack of cigarettes is not especially objectionable, from her own, subjective point of view it turns out to be so.

But I think even this is not correct. Mary, by hypothesis, judges her intention to smoke to be unreasonable and regrettable. Thus, presumably, any combination that involves this intention would seem objectionable to her on that score. But, again, this is neither here nor there: on the view in question, we have to show that in Mary's eyes the combination is objectionable in a way that the elements combined are not. But why should that be so?

Perhaps it might be suggested that Mary has a subjective reason to avoid this combination precisely because she recognizes that it would violate the instrumental requirement, and she judges it to be objectionable on these grounds. This suggestion, however, is not satisfying. It would be very surprising if the *only* grounds we could give for the instrumental requirement consisted in the brute fact that agents judge that violating it is objectionable. Even if it is true that agents judge this, we need to ask why they do so. Why should it be that, from the agent's point of view, violations of the instrumental requirement necessarily appear objectionable? What we need is an argument to the effect that, from Mary's point of view, there is something *independently* wrong with combining her

intention to smoke with a failure to intend to buy cigarettes—that is, something wrong which does not simply consist in the fact that this combination violates the instrumental requirement.

Mark Schroeder (2008) has attempted to defend the idea that agents have independent subjective reasons to avoid violating the instrumental requirement by appealing to the idea that intending to do *A* entails thinking that one ought to do *A*. As he argues, if that is true, then it would seem to follow that one also thinks that one ought to take the means that one considers necessary to doing *A*. This, then, according to Schroeder puts one under a subjective requirement to take those means (Schroeder's, therefore, is a narrow-scope view, unlike the views discussed so far).

To evaluate this proposal, let us first step back for a moment. Quite generally, accounts of the instrumental requirement that are based on reasons generated by intentions face a difficult dilemma. If Mary's intention to smoke provides her merely with a *pro tanto* subjective reason for buying cigarettes, then, plausibly, this reason is outweighed (even in her own eyes) by the considerations that speak against smoking. In this case, the account fails to explain the *stringency* of the instrumental requirement—it fails to explain why something would seem to be seriously wrong with Mary if she intended to smoke but not to buy cigarettes. Schroeder (2008: 232–3) himself argues that the instrumental requirement is stringent in this sense, and so he is responsive to this problem. But how is the problem to be avoided? The problem could be avoided if we assumed that intentions generate reasons that are always *decisive*, from the agent's point of view. This is why, in Schroeder's terms, intentions generate subjective 'oughts'. But this seems false. Agents can surely intend to do what they judge they ought not to do: Mary's case of intending to smoke while judging that she ought not to smoke provides a clear (and all too familiar) example. But judging that one ought not to do *A* plausibly entails that, subjectively, one ought not to do *A*. Should we then say that Mary, subjectively, ought both to smoke and not to smoke? It seems better to reject the claim that intentions generate subjective 'oughts', or subjective reasons that are always decisive.¹⁰

I should note here that I am not opposed to the thought that intending has normative implications. My positive proposal, as we will see, is based on the related idea that intentional action has normative implications. The problem with the proposal we have been discussing here has to do with trying to cast these normative implications in the mold of subjective reasons. Now, certainly, doing *A* intentionally (or even just intending to do *A*) involves being *committed* to doing *A*, in some sense: it involves having made up one's mind to do *A*. But it is not at all clear that it must involve one's thinking that one ought to do *A*, or even that one has any good reason to do *A*. My account, as we will see, avoids such implications.

Before leaving this section let me try to say in a more general way what the cases I envision involve, and why they cause trouble for the view that the instrumental requirement is based on reasons. The cases we have been examining are cases of *akrasia*. Such cases involve agents who are in the following

combination of states: they intend to do *A* even though they have, and *know* that they have, decisive reason *not* to do *A*. In such cases, as we saw, it seems plausible that the balance of reasons—either subjective or objective—does not speak in favor of taking the means to doing *A*. Moreover, I also argued that it has not been shown that, in such cases, the agent has independent reason to avoid combining the intention to do *A* with failing to take (or to intend to take) the means to doing *A*. But then, it seems that akratic agents do not have any independent reason to avoid violating the instrumental requirement. Unless we are to endorse skepticism about the instrumental requirement, therefore, we need a different approach—one that does not take the normativity of the instrumental requirement to be that of reasons.

3. A Standard that Governs Intentional Action

The alternative suggestion I will develop is that the instrumental requirement is a normative standard that governs intentional action. The instrumental requirement is not normative because we have reasons to avoid the combinations that it prohibits; it is normative because it partly spells out what it is to exercise one's capacity for intentional action well.¹¹

Even before getting down to details, we can see how this distinction might help us avoid the problems raised in Section 2. Consider an activity such as playing a particular musical piece—say, Beethoven's violin concerto in D. Such an activity is, it seems, intrinsically governed by standards that determine what it is to perform it *well*. Most obviously, a performance of Beethoven's violin concerto cannot be a good one unless it consists of the right notes, played in the right order. But, further, a performance can be judged better or worse depending on a host of other factors, including how well it conforms to the intentions of the composer, how 'expressive' it is, and so on. These are all standards of evaluation which seem to flow from the nature of the activity, playing Beethoven's violin concerto in D. They are, in my terminology, normative standards that govern this activity.¹²

Now, the crucial point for our purposes is that the judgment that one fails to meet the standards of a good musical performance—and thus that one performs badly—is independent of the reasons one might have for that behavior. Suppose, for example, that Mary is trying to amuse you by intentionally botching her performance of the violin part of Beethoven's violin concerto in D, and has good reasons for doing so. Even in this case, Mary's performance is open to evaluation in the light of the standards that govern performances of Beethoven's violin concerto in D: her performance is a *bad* one. Mary's performance is open to criticism, but not because she is acting against the balance of reasons. Mary is open to criticism because she is violating the normative standards that govern an activity that she is engaged in. A crucial part of my argument in what follows will be that, in a somewhat similar way, the instrumental requirement captures a normative standard that governs intentional action. If so, then there is a clear

sense in which agents who violate the instrumental requirement fail to exercise their agency well, even if they are not acting against the balance of reasons.¹³

That, in broad outline, is how I propose to avoid the problem raised in Section 2. But notice that in order for this to work we must also argue for a significant difference between the instrumental requirement and more ordinary normative standards, such as the standards that govern musical performances. As we already saw, even though Mary has a decisive reason to play badly, her performance may still be evaluated in the light of those standards, and so she may be criticized as having performed badly. But it also seems that, in the circumstances, we could not legitimately have expected of her to give any weight to the standards that govern performances of Beethoven's violin concerto (and not because there are obstacles in her way). The instrumental requirement, however, is different in this respect: agents *are* expected to take the means to their ends, regardless of the reasons they have for or against pursuing those ends. The instrumental requirement is always binding, in a way that the standards that govern musical performances are not.

I believe that the instrumental requirement really is different from other normative standards in this way. After all, if I am right, the instrumental requirement is a standard that governs intentional action *as such*. Thus, an agent that violates the instrumental requirement fails, on that occasion, *as an agent*. This is not to deny that one may have most reason to act in ways that violate the instrumental requirement: as we saw in the discussion of Section 2, that seems possible. But in doing so one fails as an agent, and this is clearly a failure that goes much deeper than one's failure as a performer of a musical piece. (How can one fail as an agent, and yet be doing what one has most reason to do? I will return to instrumental failures, and to the puzzling cases of Section 2, in Section 5.)

In the present section I will begin to develop my positive proposal, by arguing that something along the lines of the following principle captures a normative standard that governs intentional action:¹⁴

- (1) For any intentional action *A*, if one is doing *A* then one is engaged in a process which, unless something goes wrong, will result in one's having done *A*.

Suppose that Mary is walking across the street. Clearly, there is no logical guarantee that it will ever be the case that she has crossed the street: she might get hit by a truck half-way across, for example. Nevertheless, there is clearly a sense in which this is a case of failure: cases in which Mary is now crossing the street, and yet never ends up having crossed the street are cases in which something *goes wrong* with her crossing. Actions, we might say, are by nature the sort of thing that tends to get done. This is the idea that (1) aims to capture.

Now, as we shall see, principle (1) will require some defense as well as a qualification. Before getting to this, however, let me try to say, in outline, how (1) bears on our topic. Principle (1) holds both of basic intentional actions—i.e., intentional actions that one does not perform *by* performing other intentional

actions—and non-basic ones. But it follows from (1) that, if one is performing a non-basic action, then one is engaged in a process which, if all goes well, will result in one's *having performed* that action. Thus, unless something goes wrong, one *does* take all the necessary means to successfully complete the action one is now engaged in. Thus, suppose that Fred is making an omelet but never gets around to breaking any eggs, with the result that the omelet never gets made. In such a case, Fred's omelet-making is a botched exercise of his capacity for intentional action, precisely because he fails to take means necessary to see it through (of course, this might or might not be his fault, in any particular case; but that, as I will explain in Section 4, is a separate question). This, as I will argue, is the basis of the instrumental requirement.

So what is it to be engaged in a course of action that has as its normal outcome having done *A*? The answer will differ for different choices of *A*, obviously. But, generally speaking, we can say that if one is intentionally doing *A*, then (at least in a core of central cases¹⁵) one must be *committed to* doing *A*, at least in the sense that one has made up one's mind to do *A*. This entails that, if *A* is a non-basic action, then one must be, to some extent, disposed to recognize facts about the means to doing *A*, and to act on them. Our capacity for intentional action is, in part, a capacity to carry out projects of arbitrary complexity by taking appropriate means. These points will play a central role in the discussion of instrumental failure, in Section 5.

Let us now take a closer look at principle (1). Why should we accept that (1) is a constitutive principle about intentional action? The relation between one's doing *A* and one's having done *A* is an instance of a more general relation, namely the relation between a *process* and its *completion*—or, in the formal mode, between a progressive sentence of the form '*a* is *V-ing*' and its perfective correlate, a sentence of the form '*a* has *V-ed*' (my terminology here follows Gendler Szabó (2004)). Consider the process of a fruit's ripening. This fact about the fruit does not entail that, at any point in the future, the fruit will have ripened: all sorts of things might happen to prevent this state of affairs from obtaining. Nevertheless, there is a sense of normality such that the fact that the fruit is now ripening serves to pick out certain future courses of events as in that sense *normal*. Given a normal course of events, the fruit will have ripened in a few days' time. The relevant sense of normality is not statistical. Even if the presence of hungry birds in the area makes it highly unlikely that either this particular fruit or any of the others in the orchard will ever ripen, there is still a clear sense in which the normal course of events involves the state of affairs of the fruit's having ripened (we would speak here of *biological* norms, which determine what counts as the normal process of ripening for the fruit of a given species). Equivalently, a course of events that does not involve the state of affairs of the fruit's having ripened is one in which the fruit's ripening somehow *fails*; it is a course of events in which something *goes wrong* (at least from the point of view of the fruit's ripening). This general structural fact about the relation between a process and its completion is what principle (1) brings to bear on intentional action in particular.

Still, one might be tempted to object to (1), on the grounds that sometimes we succeed in doing things by *luck*—and thus not through a normal course of events. For example, suppose that Fred just won the New York City marathon. It seems that we can fill in the details so as to make it clear that Fred was never engaged in any process that had as a normal outcome his winning the marathon: Fred is not a great runner, let's say, and he has not put in all that much training; he won just because all the runners who were ahead of him mysteriously collapsed just before crossing the finishing line. Such cases, it might be thought, are counterexamples to (1), as they involve agents who *V-ed* intentionally, but were never engaged in a process whose normal outcome was their having *V-ed*.

This, however, would be a mistake. We need to distinguish between the things that one can be said to *have done* and the things that one can be said to *be doing*.¹⁶ Fred won the marathon; but there was never a time when Fred *was engaged in the process of* winning the marathon. Winning, unlike *running*, or even *trying to win*, is not a process, and so not something one can be doing. But winning may be the *culmination* of a process, and so something one may have done. But principle (1) applies only to things that one can be doing. Thus examples like this do not pose problems for (1).

Nevertheless, we still need to qualify (1) before it is ultimately acceptable. As the above example of the fruit's ripening already suggests, the fact that the sense of normality invoked here is not statistical entails that (1) is not violated by cases where one is engaged in a process that one has a very slim chance of completing. However, in its unqualified form principle (1) does entail that one cannot be doing *A* unless it is at least *possible* for one to have done *A* at some time in the future. It is not clear that this is an acceptable consequence: it seems that we are sometimes doing things that it is impossible for us to ever have done.

Consider a medieval architect who is, at a given time, building a cathedral (the example is borrowed from Gendler Szabó (2004: 38–9)). Suppose, further, that the building of the cathedral takes two or three hundred years, and the labor of many successive architects, to complete. Thus there is no course of events which involves our architect's ever having built the cathedral. Clearly, our architect is right now engaged in no process which in a normal course of events will result in his having built the cathedral. Since, by hypothesis, our architect is right now building the cathedral, principle (1) fails.

Gendler Szabó uses this example in the course of arguing against semantic accounts that seek to reduce the progressive to the perfective. Since my project here is not the reduction of the progressive, however, I—unlike the semanticists Gendler Szabó criticizes—am within my rights to simply qualify principle (1). So let us replace (1) with the following:

- (2) If one is intentionally doing *A*, and it is possible for one to complete *A*, then one is engaged in a process which, unless something goes wrong, will result in one's having done *A*.

Cases like that of the architect are clearly not counterexamples to (2). Note, moreover, that we can bring agents which are engaged in actions that are

impossible to complete within the purview of (2)—and thus within the purview of the instrumental requirement—indirectly. The idea is this. If one is doing *A*, but one has no chance of completing *A*, then it seems very plausible that one must be doing *something else*, which one *can* complete, and which one conceives of as a means to the completion of *A*. Our architect, for example, presumably counts as building the cathedral only because he is engaged in some smaller construction project, which he can complete and which is part of the project of building the cathedral. Principle (2) then applies to this other thing that he is doing.

My central claim in this section is that (2) is a normative standard that governs intentional action. But the idea that (2) captures a genuinely *normative* standard might cause alarm. As we have seen, (2) derives from the application of a general, structural fact about processes to the case of intentional action in particular. This general fact, however, holds in cases where normativity is not at issue. I earlier used the example of a fruit's ripening; but the same general structural fact holds of processes as clearly free of the normative as the sun's rising or a stone's falling to the ground. Why should we think, then, that this general fact about processes should have normative implications in the case of, e.g., Fred and his omelet-making?

In general, any course of events which involves something's *V-ing* but not its having *V-ed* is a case of an *unsuccessful V-ing*, a *V-ing* which somehow went wrong. In other words, taking a process to be a *V-ing* is taking it to be subject to a particular standard of success or failure. Crucially, then, different choices of *V* will induce different sorts of standard, although all of them will share the same general shape. So, for example, the fruit's ripening is a *biological* process. Accordingly, our general fact about processes will, in this case, induce a corresponding biological standard. If the fruit's ripening never reaches completion, then—other things being equal—the fruit, or the plant on which it grows, is thereby shown to be *defective*, in a biological sense (perhaps it has been infected with some disease, for example). Concepts of biological normality and defect apply here, although obviously they have no application in the case of the sun's rising or a stone's falling to the ground.

Now, in the special case where *V* is an intentional action, the standards induced by our general fact about processes are again different. The stone's falling to the ground and the fruit's ripening are obviously not intentional actions. It follows that, regardless of the outcome, it makes no sense to ask whether the stone or the fruit were *responsible* for it or not. Questions of blame, for example, seem out of place in these cases. Fred's omelet-making, by contrast, is an intentional action. There is a clear sense in which Fred's success or failure at omelet-making is up to him: in particular, it is—other things being equal—up to Fred whether he will take, or fail to take, the means necessary for his omelet-making to succeed. It makes sense to ask whether Fred was *responsible for* its success or failure (sometimes, of course, the answer to this question will be negative—Fred may have been felled by a stroke as he was about to break the eggs). Accordingly, questions of blame are not out of place here. If Fred fails to

complete his omelet making, then—other things being equal—the failure may be blamed on him. Thus, the standards induced by our general fact about processes in this case are entwined with concepts of responsibility, blame and the like. They are *normative* standards.¹⁷

4. The Instrumental Requirement as a Normative Standard

In the previous section I argued that intentional action is a normatively constituted phenomenon, which is governed by the normative standard expressed in (2). In the present section I want to argue further that (2) is the source of our intuitions about the instrumental requirement.

The basic idea is already familiar. Suppose that one is engaged in a non-basic intentional action *A*, and *M* is a means necessary to doing *A*. Since doing *M* is necessary for doing *A*, every course of events in which one does not do *M* is a course of events in which one ends up not having done *A*. Thus, according to (2), it is a course of events in which one's doing *A* has gone wrong. It is a course of events in which one has not exercised one's agency well. Thus every course of events that involves one's failing to take the means necessary to an action one is engaged in is a course of events in which one fails to exercise one's agency well. This, I believe, is the source of the instrumental requirement.

Note that we have here an answer to the puzzle that akratic agents, such as those discussed in Section 2, raise. Suppose that one is doing *A*, even though one has decisive reason not to be doing *A*. Still, according to the account just given, if one fails to take means necessary to complete *A*, one is not exercising one's agency well. We will discuss such failures in more detail in Section 5.

Thus my account is, in its basics, very simple. However, complications arise. One way in which my approach might seem surprising has to do with the fact that (2) is formulated as a principle about *intentional actions*, while the instrumental requirement is often formulated as a principle about *intentions*. Roughly, recent authors take the instrumental requirement to require one to *intend to take* the means to what one *intends to do*, rather than to take the means to what one is intentionally doing.

To some extent, the appearance of novelty here is illusory: although intentions are not mentioned in (2), they are clearly implied. At least in the core cases of intentional action,¹⁸ if one is *intentionally doing* something then one *intends to be doing* that thing (this is a progressive intention, or an intention in action). For another, the following seems true:

(3) If one is required to do *M*, then one is required to intend to do *M*.¹⁹

What lies behind (3) is the thought that intending is the servant of doing. Intentions for the future, in particular, are nothing but our way of controlling our future actions. Assuming this is true, (2) entails that if one is intentionally doing *A* one is required to *intend to take* the means necessary to doing *A*.

This takes us closer to the usual, intention-oriented formulations of the instrumental requirement, but it does not take us all the way there. This is because, according to those formulations, having an intention to do *A* appears to have normative consequences for what one does *before* one has started doing *A*. For example, Mary's intention to smoke has normative consequences for her even if she is not yet smoking. But this gap too can be bridged. Perhaps one might argue—with Thompson (2008: 138–46)—that intending to do *A* is not distinct from doing *A*: intending to do *A* is something like an *initial stage* of doing *A*. If that is true, then the intention-oriented formulations of the instrumental requirement would be simply special cases of (2). But we don't need to commit to such a bold metaphysical claim as this. We can simply say that the intention-oriented formulations of the instrumental requirement are natural *extensions* of (2). From (2), we can derive the normative relation one has to the means to doing *A* while in the course of doing *A*. Intention-oriented formulations of the instrumental requirement move beyond this, by asserting that one stands in *the same* normative relation to the means to doing *A* even while one merely intends to do *A*. Once again, this would be justified by the thought that intending is the servant of doing.²⁰

On a different front, many authors qualify the instrumental requirement by adding a belief condition: on such views, what matters is not whether doing something *really is* a necessary means to doing *A*, but only whether one *believes that* it is. So how does my approach compare with these views? There are two ways in which this belief condition might make a difference in the application of the instrumental requirement. I will consider each in turn.

In one sort of case, an agent might fail to take means necessary to an action she is intentionally performing because she does not believe that it is necessary. Formulations that include the belief condition would not consider such a case a violation of the instrumental requirement. According to (2), by contrast, this is a case where something goes wrong with the agent's doing *A*; the agent has not exercised her agency well on this occasion.

I believe that the verdict of (2) is intuitively correct. Something *has* gone wrong with the agent's doing *A*. Of course, depending on circumstances, she might have an *excuse* for this failure—for example, her failure might be excused if figuring out the means to her end could not have been expected of her in the circumstances. But I do not think that the possibility of excusable failure needs to be built into our formulation of the instrumental requirement.

Notice that appealing to the conceptual apparatus of blame and excuse is not an unnecessary complication, peculiar to my account: any account of the instrumental requirement would need to appeal to such considerations anyway. Suppose that Fred knows that one cannot make an omelet without breaking eggs, but that, just as he is about to break the eggs, he is felled by a stroke. No belief condition is going to save Fred from violating the instrumental requirement in this case. Nevertheless, it is clear that Fred is not to be blamed for violating the instrumental requirement; Fred has an iron-clad excuse. Quite generally, the task of identifying a *requirement* and the task of identifying the

conditions under which one may be *blamed for* violating that requirement are distinct.²¹

We have been discussing cases of ignorance about the means to one's end. The other type of case that might seem to call for a belief condition involves mistaken beliefs about the means to one's end. Suppose that the agent believes that doing *M* is a necessary means to doing *A*, but she is wrong about that. Does the instrumental requirement nevertheless require of her to do *M*? According to the view I propose here, it does not: since there are (by hypothesis) courses of events which do not include the agent's doing *M* but include her having done *A*, her not doing *M* does not by itself constitute any kind of failure on her part.

Again, I think this result is intuitively acceptable. On my view, if there is a sense in which the agent in such a case is required to do *M* it is a derivative one, which needs separate explanation. In general, the following principle seems plausible:

- (4) If the fact that *p* would put one under a requirement to do *A*, the fact that one believes that *p* (perhaps falsely) puts one under a different, and derivative, requirement to do *A*.²²

To the extent that anything along these lines is correct, the type of case presently under consideration does not provide a reason for adding a belief condition to the instrumental requirement.

Finally, one might be suspicious of my approach on more inchoate, but perhaps deeper, grounds. The instrumental requirement is a requirement of *subjective* rationality. And one might think that subjective rationality must be a purely internal matter, in the sense that it must be exclusively a matter of what the agent *believes*, rather than facts about the world. My approach would seem to violate this restriction, by bringing in facts about means-ends relations.

I think we should resist this line of thought, however. Although there is a sense in which subjective rationality is an internal matter, I doubt that the sense in which this is so disqualifies means-ends relations from playing a role as determinants of subjective rationality. The requirements of subjective rationality spell out standards for the proper functioning of our rational faculties, regardless of the quality of their inputs. This suggests, as I already mentioned in Section 2, that the requirements of subjective rationality should hold even in abstraction from the worldly facts that constitute the reasons for our actions and our propositional attitudes. This is the sense in which subjective rationality is an internal matter. But this conception of interiority provides no support for the idea that the requirements of rationality should also abstract from the *real relations among* one's attitudes or one's actions. In formulating the requirement to avoid incoherent beliefs, for example, it seems that we should be allowed to appeal to facts about when a given set of beliefs *really is* incoherent. Similarly, I suggest, in discussing the instrumental requirement we should be allowed to appeal to real means-ends relations among actions.

5. Instrumental Failures

The challenge for any account of the instrumental requirement is to explain why rationality seems to require of agents to take the means to their ends, even when they pursue ends that they have decisive reason not to pursue. I have tried to meet this challenge by arguing that intentional action is a normatively constituted phenomenon, and that the instrumental requirement partly spells out what it is to exercise one's capacity for intentional action well. Thus, agents who violate the instrumental requirement *ipso facto* do not exercise their agency well. But how does all this fit into our ordinary understanding of our own agency and its failures? From our point of view as agents, what do violations of the instrumental requirement look like? I will close this paper by making some remarks about how my account fits in with a plausible ground-level phenomenology of instrumental failures.

Our capacity for intentional action is a complex capacity. It is the capacity, first, to choose a course of action, and then to execute it by taking appropriate means. This capacity can go wrong in a lot of different ways. Most obviously, it can go wrong in the very first step: we can knowingly choose to do what we have most reason not to do. But there are other ways for our capacity for intentional action to go wrong too.

Once we have set out on a project, we are committed to bringing that project to completion. It is not possible to be doing *A* intentionally without giving some weight in one's deliberations to the question whether one's actions will result in one's ever having done *A*.²³ This claim is not intended as an empirical generalization about agents: it is a metaphysical claim, about what it is to act intentionally. If Fred is puttering about in the kitchen but does not care about whether he will ever have made an omelet, then, whatever else he might be doing, he is not intentionally making an omelet. Thus, if Fred really is intentionally making an omelet, then he must be exhibiting some kind of sensitivity to facts about the means he needs to take to that end: he must be to some extent disposed to recognize and be motivated by such facts. Exhibiting this sensitivity is part of what it is for him to exercise his capacity for intentional action in this case.²⁴ This is simply the phenomenological side of (2): it is simply a description from the inside of what it is like to be engaged in a course of action which has as its normal outcome having made an omelet.

Note that we need not assume that the agent's commitment to bring her project to completion involves her thinking that her project is one she has most reason to pursue, or that she ought to pursue (that, as we saw, was Schroeder's (2008) strategy). As was emphasized in Section 2, agents are subject to the instrumental requirement even when they pursue ends that they know they ought not to pursue. But even if one is pursuing a project one knows that one ought not to pursue, one is still engaged in a course of action that will result in one's having completed that project, unless something goes wrong. One must, therefore, still be exhibiting the relevant sort of sensitivity to facts about appropriate means. Of course, in such a case one need not think of the means

as things that one has *reason* to do, just as things that are conducive to what one *is doing*.²⁵

This kind of sensitivity, however, does not entail that the agent really *will* take the means to her end. Familiarly, all sorts of things can go wrong: the agent might be distracted, lazy, paralyzed with fear, or simply absent-minded. These are all phenomena that we appeal to in explaining phenomena of *akrasia*, or acting against one's better judgment. But they can equally well serve to explain cases where agents (knowingly) fail to take the means to their ends. Whatever the explanation in a specific case, however, the crucial point is that these are all cases in which the agent fails to exercise her agency well, just by failing to complete an action she is engaged in.

Let us finally return to the puzzling cases of Section 2. The agents discussed in Section 2 are puzzling, because they have *most reason* to fail to exercise their agency well in this way. But on the picture of agency sketched here this is not such a surprising result. Our capacity for intentional action involves both a capacity to choose what projects to commit to, and a capacity to carry out our chosen projects by taking appropriate means. Accordingly, a course of action might be evaluated in (at least) two different respects. It is not too surprising that a course of action might be all right in one of those two respects, but not in the other.²⁶

Markos Valaris
 School of History and Philosophy
 University of New South Wales
 Australia
 m.valaris@unsw.edu.au

NOTES

¹ Should we say of Mary that she both *ought* to (or *should*) buy cigarettes, and that she *ought not* to (or *should not*) buy cigarettes? I do not think that we need to settle this for the purposes of this paper. My point can be made simply by noting that, apparently, *rationality requires* of agents to take the means to their ends, even when they have decisive reason not to be pursuing those ends in the first place. Whether we need to posit an ambiguity in 'ought' or 'should' to capture this point is not an issue I will discuss in this paper.

² Strictly, all that this and related arguments I discuss later in this paper show is that no *straightforward* reduction of normative standards to reasons is possible. But I have no general proof that a unified theory of practical normativity, based entirely on reasons, is out of the question. Accordingly, my claims to the effect that the normativity of the instrumental requirement is not that of reasons should be similarly qualified: my arguments in this paper show that the normativity of the instrumental requirement does not *straightforwardly* reduce to the normativity of reasons. I will leave this qualification implicit in what follows.

³ Judith Thomson (2008: 25) asserts that the kind 'act' (presumably she means *intentional* act) is not a 'goodness-fixing kind', i.e., in my terminology, that it is not the case

that intentional action as such is governed by normative standards. Thomson does not provide an argument for this claim; I hope that what I have to say in this paper will show that it is mistaken.

⁴ I use the term 'normativity' broadly, to include both *evaluations* and *directives* or *deliberatives* (for the distinction, see Wiggins 1998: 95). It seems clear that deep conceptual connections exist between evaluations and directives. The claim that *V-ing* is a good thing for one to do in a given situation surely has *some* bearing on what one should do, in that situation. The details are bound to be complex, and I cannot take them up here.

⁵ For some recent attempts to argue that, contrary to appearances, we always have reasons to avoid violating the instrumental requirement, see Schroeder 2008 and Bratman 2009 (note that Schroeder discusses *subjective* reasons; as I will argue below, this does not make a significant difference for present purposes). For the view that there is *no* distinctive requirement of instrumental rationality, see Raz 2005a, Kolodny 2005; 2008, Broome 2005; 2008. Bratman 1987 suggests a similar view, but he revisits the issue in later work (see especially Bratman 2009). For the so-called *cognitivist* view, according to which the instrumental requirement is a requirement of *theoretical* rationality, see Harman 1999, Wallace 2001, Setiya 2007. I do not discuss cognitivism directly in this paper; since cognitivists deny that the instrumental requirement is *distinctive* (it is simply an instance of a more general requirement of theoretical rationality) they might be grouped together with the skeptics. Note, however, that Setiya 2007 is an exception to the tendency to assume that the normativity of the instrumental requirement must be that of reasons: he takes seriously the possibility that it might be a requirement of *good functioning*.

⁶ Thus, according to wide-scope views, rational requirements require of one to avoid incoherent combinations of mental states (or mental states and actions) but are indifferent as to *how* one goes about doing so. According to narrow scope views, by contrast, this is not so: rational requirements require of one to adopt or avoid a specific mental state or course of action. Note that Broome has more recently (2005; 2008) expressed doubts as to whether we really have reasons to conform to the constraints of rationality.

⁷ It tells us *at least* this, but it might tell us more. On a narrow-scope view, it would tell us more: it would tell us, specifically, that Mary is required to *buy cigarettes*. I will discuss a narrow-scope proposal later on.

⁸ Of course, if Mary were to intend to buy cigarettes, then in the circumstances envisaged she would tend to answer the non-normative question '*Will I smoke?*' firmly in the affirmative. But this would not represent any net gain in self-governance, for it would simply exacerbate the conflict between her normative outlook and what she is intentionally out to do.

⁹ Bratman, in response to an argument by Setiya (2007), suggests that agents have a self-governance-derived reason to avoid violations of the instrumental requirement *only if* self-governance is *attainable* by them. Perhaps it might be argued that, given Mary's conflicted starting point, self-governance is unattainable for her in this situation, and so Bratman's account does not apply. There are two things to say in response to this suggestion. First, it is not clear that it is true to Bratman's intent. The cases Bratman's (2009: 428–9) condition is meant to rule out involve *unmodifiable* intentions. But there is no reason to think that Mary's intention to smoke is unmodifiable in the relevant sense (although, of course, it has proven to be somewhat resilient in the face of her normative judgment). Mary need suffer from no more than ordinary akrasia. It is not obvious that Bratman intends to exclude such cases from consideration. In fact—and this is the second point—Bratman *should not* want to exclude such cases from consideration. It seems clear

that akratic agents are subject to the instrumental requirement. If Bratman's account is silent about such cases, that is not a point in its favor.

¹⁰ Schroeder (2008: 238–9) has a brief discussion of akrasia. It is not clear from this discussion how he would respond to the difficulties raised in the text.

¹¹ Although Raz 2005a uses some similar-sounding language in discussing the instrumental requirement, it is important to note that the view I will develop is radically different from his. As Raz 2005b makes clear, his view is *not* that the instrumental requirement itself captures a standard of proper functioning of our practical faculty. His claim is only that instrumental failures are often evidence of *other* defects in our practical faculty, such as dithering. This is the force of arguing that instrumental rationality is a myth.

¹² The example is inspired by Gideon Rosen's 2001 discussion. Rosen discusses only the *correctness* of the performance, which is just one aspect of its goodness.

¹³ One might suggest that we could reduce normative standards to *conditional* claims about reasons, i.e., claims about the reasons one would have in specified circumstances. It is hard to see how this is supposed to work, even in the case of the standards that govern musical performances. Of course, it seems true that the standards describe how one has reason to perform, if one has reason to perform *well*. But, clearly, to say this is to *presuppose* the very standards we were supposed to explain.

¹⁴ This fact about intentional action is given center stage in Michael Thompson's 'Naive Action Theory' (2008: 85–146). The influence of this work on what follows is pervasive.

¹⁵ The cases I have in mind here are those demarcated by Anscombe's (1957) question 'why?' test, i.e., those for which requests for reasons are appropriate. Perhaps the notion of intentional action extends more widely than that. If I am walking to the store, and I know that in walking to the store I am wearing down the soles of my shoes, then perhaps there is a sense in which I am *intentionally* wearing down the soles of my shoes. Anscombe's question 'why?', however, would seem to have no application here; relatedly, no commitment is involved in such actions. Such cases are murky. In particular, one does not seem to be under any independent requirement to take the means necessary to see them through. It seems that one is under such a requirement only because (and to the extent that) those means are also means to doing *something else*, which is a core case of intentional action. (For example, I am under a requirement to take means to wear down my shoes only to the extent that these are means to walking to the store.) For this reason, I will ignore them here.

¹⁶ Philosophers of language often appeal to a four-fold distinction among verbs and verb-phrases (see Parsons 1990): those that express *states* (e.g., 'see', 'know'), those that express *achievements* (e.g., 'win the marathon', 'pass the test'), those that express *accomplishments* (e.g., 'cross the street', 'run a marathon') and those that express *activities* (e.g., 'take a walk', 'run'). Stative verbs don't express processes at all; thus, a fortiori, they don't express processes that one can be engaged in intentionally. For our purposes, the important distinction is that between achievement verbs, on the one hand, and accomplishment and activity verbs, on the other. Intuitively, achievement verbs express the *culmination* of a process, not the process itself. Thus, although they express things that one may *have been done* intentionally, they don't express things that one may *be doing* intentionally.

¹⁷ Some philosophers have no difficulty speaking of normativity even where the concepts of responsibility, blame and the like have no application (see, e.g., Thomson 2008). Such philosophers would certainly have no trouble attributing normative features

to biological processes, and even to processes that involve only non-animate stuff. I do not wish to insist that the application of concepts such as responsibility and blame is a *necessary* condition for normativity; my argument requires only that it be a *sufficient* one.

¹⁸ Remember that the cases I have in mind here are those to which Anscombe's question 'why?' has application.

¹⁹ Shah (2008) defends a stronger (but still plausible) principle, namely that all the normative features of intentions for the future are derived from the normative features of the actions they are intentions to perform. Both principles might need to be qualified in view of the possibility of actions which one is both required to perform and *will* perform, whether one intends to or not: in such cases, it would seem, one is under no normative requirement to intend to do *A*.

²⁰ One might worry about this extension, on the grounds that sometimes we may intend to do things that are not processes, and therefore not things that we can ever be *doing* intentionally. For example, Mary may intend to *win* the marathon, although winning is not the sort of thing that one may be doing intentionally. Winning is an *achievement*—i.e., the culmination of a process, rather than a process. We can get around this difficulty by noting the following: necessarily, if one intends to do *A* where *A* is an achievement, one also intends to do *B* where *B* is a process of which *A* is a possible culmination. For example, if Fred intends to win the marathon, it must also be the case that he intends to, e.g., *try to win* the marathon, or to *run* in the marathon. Otherwise what he has is an *idle wish*, not an intention.

²¹ Some authors also add the condition that one must not believe that the event necessary for one's end will happen anyway. For example, on such a view Fred would not be under a requirement to break any eggs if he believed that Mary would break the eggs for him. Again, I do not find the reasons for this condition compelling. Fred might have a good *excuse* for failing to break the eggs, if his belief that Mary would break the eggs is itself blameless. But we don't need to build this fact into the formulation of the instrumental requirement itself.

²² See Schroeder 2008: 230 for a similar principle.

²³ This might be what Raz (2005a: 17) is driving at, when he writes: 'Necessarily anyone who intends to perform an action or tries to perform it intends to perform it successfully or tries to perform it successfully.' On one natural reading of Raz's statement, it would seem to be false: I may intend to perform a musical piece badly, by the standards specific to performances of that piece. But perhaps Raz's remark should be taken to mean only that one cannot intend to do anything without thereby intending to ever have done that thing.

²⁴ This is a central thesis of Thompson 2008.

²⁵ Perhaps some might object here, by doubting that thoughts of this form are suited to play the right sort of role in deliberation: only thoughts about *reasons*, one might insist, can play the relevant role. It seems correct that my view requires rejecting this restriction, but I do not think this is a problem for my view.

²⁶ I presented an early draft of this paper at the Engaging McDowell conference in Sydney in July 2010, and benefitted from comments by Nishi Shah, Michaelis Michael and John McDowell. More recent versions of the paper were presented at seminars at the University of New South Wales and at the Australian National University. Members of the audience on these occasions raised a number of penetrating questions. I am especially indebted to comments by John Bengson and John Maier. I am also grateful to an anonymous referee for this journal, whose comments helped me greatly in clarifying my views.

REFERENCES

- Anscombe, G. E. M. (1957), *Intention*. Cambridge, MA: Harvard University Press.
- Bratman, M. (1987), *Intention, Plans and Practical Reason*. Cambridge, MA: Harvard University Press.
- (2009), 'Intention, Practical Rationality and Self-Governance', *Ethics*, 119: 411–43.
- Broome, J. (1999), 'Normative Requirements', *Ratio*, 12: 398–419.
- (2005), 'Does Rationality Give us Reasons?' *Philosophical Issues*, 15: 321–37.
- (2008), 'Is Rationality Normative?' *Disputatio*, 11: 153–71.
- Gendler Szabó, Z. (2004), 'On the Progressive and the Perfective', *Noûs*, 38: 29–59.
- Harman, G. (1999), 'Practical Reasoning', in G. Harman, *Reasoning, Meaning and Mind*, 46–75. New York: Oxford University Press.
- Kolodny, N. (2005), 'Why Be Rational?' *Mind*, 114: 509–63.
- (2008), 'The Myth of Practical Consistency', *European Journal of Philosophy*, 16: 366–402.
- Parsons, T. (1990), *Events in the Semantics of English*. Cambridge, MA: MIT Press.
- Raz, J. (2005a), 'The Myth of Instrumental Rationality', *Journal of Ethics and Social Philosophy*, 1: 2–28.
- (2005b), 'The Myth of Instrumental Rationality: A Reprise', *Journal of Ethics and Social Philosophy: Symposium*, 1: 1–19.
- Rosen, G. (2001), 'Brandom on Modality, Meaning and Intentionality', *Philosophy and Phenomenological Research*, 63: 611–23.
- Schroeder, M. (2008), 'Means-End Coherence, Stringency, and Subjective Reasons', *Philosophical Studies*, 143: 223–48.
- Setiya, K. (2007), 'Cognitivism About Instrumental Reason', *Ethics*, 117: 649–73.
- Shah, N. (2008), 'How Action Governs Intention', *Philosophers' Imprint*, 8: 1–19.
- Thomson, J. (2008), *Normativity*. Chicago: Open Court.
- Thompson, M. (2008), *Life and Action*. Cambridge, MA: Harvard University Press.
- Wallace, R. J. (2001), 'Normativity, Commitment and Instrumental Reason', *Philosophers' Imprint*, 1: 1–26.
- Wiggins, D. (1998), 'Truth, Invention and the Meaning of Life', in D. Wiggins, *Needs, Values, Truth*, 3rd edn. New York: Oxford University Press.