

CHAPTER THREE: THE SCOPE OF INSTRUMENTAL REASONING

3.0 Introduction

As argued in the last chapter, decision theory does not provide an adequate account of practical reasoning or even of instrumental practical reasoning. But instrumental reasoning itself deserves further attention on at least three counts.¹ First, it is very much a part of our ordinary experience that we adjust means to ends and regard various kinds of failures to do so, at least when the issue is clear-cut from the agent's standpoint, as irrational.² We say, for example, that people must regularly change motor oil if they want their automobiles to work properly or that they are mistaken to rely on the purchase of

¹ I do not believe that the instrumental exhausts practical reasoning, but I shall not venture far beyond it for the two reasons that much more, and much more that is interesting, can be done with instrumental reasoning than is commonly thought and for the practical (and instrumental!) reason that limits must be set somewhere to the scope of the current project if I am to finish it.

² Perhaps it would be better to say that we regard failures to adjust means to end (again, from the standpoint of the agent) as subject to criticism on the count of their rationality. A charge of irrationality is the extreme case of such criticism, but there may be failures of rationality that would be better described as non-rational or as less rational than some alternative, rather than as irrational.

lottery tickets for their retirement plans. Instrumental reasoning has the advantage of familiarity.

Second, instrumental reasoning seems non-committal, or perhaps better, minimalist, in what it requires us to assume or presuppose in the way of value theory or any necessary ontological underpinnings. Even those who are most skeptical of normative discourse find it acceptable to grade actions as being more or less appropriate means to given objectives and do not find that doing so requires the invocation of mysterious non-natural properties or cognitive powers.³ In instrumental reasoning, we get ‘oughts’ that do not seem to have a problematic relation to what is the case.⁴ Instrumental reasoning has the advantage of being metaphysically and epistemologically undemanding.

Third, as Nozick says (which may be partly explained by the second point),

The notion of instrumental rationality is a powerful and natural one.

Although broader descriptions of rationality have been offered, every such description that purports to be complete includes instrumental rationality within it. Instrumental rationality is within the intersection of all theories of rationality (and perhaps nothing else is). In this sense, instrumental rationality is the default theory, the theory that all discussants of rationality can take for granted, whatever else they think. (1993, 133)

³ See, for example, Mackie 1977, 27-30.

⁴ I do not think this means that the conclusions we get are not genuinely normative or that their normativity is or must be reducible to something else. Failures of practical rationality are not to be identified with having made one or another sort of theoretical mistake.

Because it is the default position, the one that can be presumed to be shared, whatever else may be controversial, any results that can be reached in terms of instrumental reason should be acceptable to everyone. Instrumental reasoning has the advantage that it speaks to everyone.

Instrumental reasoning, then, has a number of attractive properties.⁵ That, of course, is hardly enough to show that it is suited to play any large role in moral theory. To reach a judgment on that question – on whether and how instrumental reasoning is suited to play a large role in moral theory – part of what is needed is a more detailed examination of instrumental reasoning itself. What I shall do is to sketch the features that we ordinarily take to be involved in instrumental reasoning, beginning with the paradigmatic cases in which it is an external means to some single objective⁶ that is under consideration. I will extend this with brief attention to cases in which more than one objective bears upon the selection of means, and then by considering how well the features elicited also characterize a less paradigmatic type of case that may be – and I shall argue should be – assimilated to instrumental reasoning. I shall then say something about the normative standing of instrumental reasoning in relation to the characterization I have given.

⁵ A further feature that I find attractive is that, so far as the possession of the ends or objectives (in the light of which means are to be adjusted) is conceived in terms of motivating states of the agent, instrumental reasoning appears to satisfy the plausible internalist requirement that genuine reasons be motivating for rational agents. (For some doubts on this point, see Hampton 1998, Chapter 2.) Now, however, debates about internalism and externalism with respect to reasons have become increasingly intricate and the positions of both internalists and externalists ramified to the point that it is obscure what, if any, is the connection to the issues that originally prompted making the distinction (see, e.g., Darwall 1997 and Audi 1997). For present purposes, rather than become embroiled in that discussion, I set the issue aside.

⁶ The terminology is explained below.

3.1 Features of Instrumental Reasoning

Instrumental reasoning is most clearly involved when the reasoning focuses upon the relation of external means to some single objective. An objective, when considered in relation to an external means, is some action or state of affairs that can be specified independently of that means,⁷ and that external means is something selected, adopted or performed because of its (expected) causal contribution to the realization of the objective. An objective may or may not itself be an end – something sought, aimed at or performed for its own sake – but that distinction will not concern me here, so I shall in the remainder of this chapter sometimes refer to objectives simply as *ends*, to external means simply as *means*, and to the relation between such an objective and the corresponding external means, when achieving the objective is the only relevant consideration, as the *simple instrumental paradigm*.

3.11 The Simple Instrumental Paradigm

To consider the simple instrumental paradigm, it is useful to employ a typology to distinguish the possible sorts of cases that satisfy its conditions. Four sorts are distinguishable, and they can be represented in this way:

⁷ The objective need not be, and generally is not in fact, completely specified. There is normally a range of states of affairs that would count as the realization of the objective. The objective may be to have steak for dinner, but there are lots of different things that would count as having steak for dinner, and the agent normally will not have in mind such details as the exact size or cut of the steak or the precise microsecond at which dinner will commence.

	Sufficient	Not Sufficient
Necessary	Type I	Type II
Not Necessary	Type III	Type IV

The simplest imaginable case is one in which some means is causally⁸ both necessary and sufficient to bring about or realize the end. In the second, the means is necessary but not sufficient. In the third, the means is sufficient but not necessary. In the fourth, the means is neither necessary nor sufficient. Obviously, these logically exhaust the possibilities.^{9, 10}

Cases of the first two types can be considered together because, so long as the end is not itself in question, but the means is necessary to realize it, the sufficiency or insufficiency of the means makes no difference to the points with which I am concerned here.¹¹ Consider a case in which there is some end, *E*, at which an agent aims, and some

⁸ I shall not continue to qualify the means as *causally* (rather than, say, logically) necessary or sufficient for an end, but the qualification should be assumed.

⁹ In all of these, I am assuming, again, that no other considerations than the relation of the proposed means to their respective single objectives are relevant.

¹⁰ For my purposes, the necessity or sufficiency of the means with regard to the end (or the lack of either) is to be assessed from the standpoint of the agent's knowledge or beliefs. Whether she is correct to have those beliefs is a further question from which I am abstracting. Additionally, since we are considering *causal* necessity or sufficiency, it will often not be the case that assignment of an example to one of the four classes is clear-cut, even abstracting from the agent's standpoint. In particular, it may be difficult to ascertain that some means really is necessary in the sense that no other means, including ones not considered, would work or that it is sufficient in the sense that, once the means is adopted, nothing could derail the expected realization of the end. Similarly, it may be difficult to ascertain whether some means that is assumed either not to be necessary or not to be sufficient really is not.

¹¹ There are two ways in which some means may be necessary but not sufficient to realize an end. First, it may be that a means, *M*₁, is not sufficient unless some other means, *M*₂, or some set of other means, *M*₂, ..., *M*_{*n*}, is also employed. Accordingly, when this is the case, there is some *set* of means in the agent's power which, taken together, is sufficient to realize the relevant objective, though no proper subset is sufficient. If that is so, this is just a more complicated form of Type I case; there is something the agent can do, namely, take all the means together, that is both necessary and sufficient to achieve the objective. Second and more interestingly, it may be that some means is necessary but not sufficient because the

means, *M*, within his power, which is necessary to bring about *E*. The agent is aware of this and, being rational, selects *M* in preference to any alternative action, *N*, that may be possible in the circumstances.¹²

Consider now cases of Type III, where some means is sufficient but not necessary for the achievement of the end. There are two ways this might be so.¹³ First, it may be that there is some non-zero probability that the end will be achieved, even if the means is not employed. However, there will also be some non-zero probability that the end will not be achieved if the means is not employed, since otherwise, there would be no sense to speaking of something as *means* to the end: to identify something as a means is to

realization of the end depends in part upon factors beyond the agent's control, such as concurrent actions by others or the presence of other causal factors. These in turn may be factors about which the agent can do something, though, if the case is to remain distinguishable from the first possibility, and therefore in turn from cases of Type I, what the agent can do can at most raise the likelihood of the needed concurrence. (As Fred Miller pointed out to me, means that are necessary but not sufficient to bring about the realization of some end may not, by themselves, even increase the likelihood that the end will be realized. For example, if the end were to win a lottery, it would be necessary though not sufficient to go somewhere that lottery tickets are sold, but that does not by itself increase the likelihood of winning; one must also purchase a ticket.)

¹² At this point, there is a complication. In a particular situation, there may be no alternative to *M* open to the agent: *M* may be not only necessary to bring about *E*, but also the only thing the agent could do in any case, regardless of its bearing upon *E*. (This may be a case in which the action is over-determined. Had the agent not selected it – on his own, so to speak – some other factor would have intervened to result in his performing it anyhow. See Frankfurt 1969.) If so, then it would be unclear whether *M* was adopted as a means to *E* or not, since the agent would have performed *M* in those circumstances, even if *E* had not been his objective.

The problem is that the agent's deliberation and action are supposed to be rational, but "rational" takes its meaning in part from the contrasting case or cases in which deliberation or action is non-rational, irrational, or less rational. So, if the agent would have performed *M* whether or not *E* had been his objective – that is, if there is no contrasting case – in virtue of what is it or could it be true that *E* is his objective and that *M* is adopted as a means to *E*?

What is needed here is the truth of a counter-factual: If there were other options, alternatives to *M* that did not alter the type of case (so *M* is still necessary for *E*), then the agent would, if rational, select *M* rather than any of the alternatives. (And, by contraposition, if the agent would not or might not select *M*, then either *E* is not his end – at least not the only one that is relevant – or else he is not rational.)

¹³ The two might also be combined, but there are no distinctive points to make about the combination.

distinguish it as being in some way better than alternatives¹⁴ in relation to an end. So, when there is some non-zero probability that the end will not be achieved without employing the means, the employment of the means raises the probability that the end will be realized – in this case, from some value less than unity to unity. If it did not increase the probability, the agent would have no reason (in terms of that end) to favor the employment of the means.¹⁵

Second, the agent might have available to him more than one option that is sufficient to realize the objective. There is available to him, say, M_1 which is sufficient to realize the objective, E , and M_2 which is also sufficient for E . Since M_1 and M_2 are each separately sufficient for E , neither is necessary. Here, there are two important points. One is that when there is some set of means, each of which is separately sufficient for the end, then, though none of them is necessary, the agent still has a reason in terms of the end for selecting one of them. It is not necessary that he select some particular one, but it is (rationally) necessary that he select one over any alternative that is not a means to the end. The other point is that when the agent is making a selection from among a set of separately sufficient means, *if* there is a reason for selecting one (or a member of some subset) over the others, then that reason must have other sources. There must, e.g., be some other end in terms of which one or more of the available means is judged less costly

¹⁴ At least, something taken to be a means must be thought to be better than alternatives counterfactually available, as in note 12. In the present case, even that is not available for the relevant range of counterfactuals, for we are assuming that the end would be realized regardless of any selection of means.

¹⁵ This does not imply that he has a reason *against* employing the means. He might be indifferent whether the means is employed or not.

or more desirable than any other option among the set of available means.¹⁶

Much of what has just been said can, with appropriate modifications, be extended to cases of Type IV, in which available means are neither necessary nor sufficient to achieve the end. Again, it may be that the end could be realized whether or not some relevant means is employed, and again, there may be some set of available means from which a selection must be made. Still, this is perhaps the most interesting type of case covered by the simple instrumental paradigm because, in considering how it is rational to act in the selection of some means appropriate to the end, we encounter a feature pervasive in ordinary instrumental reasoning, the relevance of the probability that some selected means will result in achieving the given objective. Probabilities have, of course, figured in the earlier discussion of Type III cases, but then, their only role was to discriminate between some actions that do not involve employment of a means a given objective and others that do. Considerations of probability did not there bear upon selection *among* means. In none of the other three types of case did the sufficiency or insufficiency of the means for the end provide any basis for selecting among candidate means. It did not in the first and second, because there, the means was supposed to be necessary for the end, whether or not it was sufficient. It did not in the third because all of the candidate means were supposed to be sufficient.

What we have here is much more interesting. One way in which means can be neither necessary nor sufficient to achieve an end is when they stand in a probabilistic

¹⁶ I do not mean to imply that nothing but some other end could be the basis for selection among a set of means.

relation to the end. A means, M_1 , may be more likely to bring about an end than some alternative, M_2 . It appears that we can use the end, E , to judge that M_1 should be selected rather than M_2 , even though neither of them is either necessary or sufficient for the realization of E .¹⁷

There is here an extraordinarily interesting question as to why, *for the single case*, it is reasonable to act on the basis of what is most likely to happen. We can say, of course, that the agent facing a long run, or an indefinitely extended series of cases of the same type, will do better if, in each of them, she predicates her decision upon what is most likely to happen (though there are complications even here). There are also cases in which the reasonable thing to do is to take both the more and the less probable outcomes into account in deciding upon action – for example, by purchasing insurance against something less probable (and worse). But there are also cases in which such hedging of one's bets is not an option. Suppose the agent is offered, only once, a forced choice between a pair of gambles, G_1 and G_2 , with the same (desired) prize associated with each. If she selects G_1 , she will probably receive the prize; if she selects G_2 , she will probably not receive the prize. We all believe she should select G_1 , but why? She may not receive the prize if she does and may receive it if she does not. To say that the reason is that it is more likely that she will receive the prize by selecting G_1 is just to reiterate our belief that she should guide her actions by the probabilities. It does not really explain why she should. (Remember that she is not facing a long run or many cases of the same type.) It might be proposed

¹⁷ There is a hybrid case intermediate between Types III and IV. There may be some means that is sufficient for the production of the end (as in Type III) and also some means not sufficient for the end but which makes it more likely that the end will be realized than if means are not employed (as in Type IV).

that what she is really aiming at, her real objective, is the best chance of achieving or bringing about some state of affairs – namely, the one in which she receives the prize – but that seems to misdescribe the phenomenology. It is because she cares about the state of affairs, because getting the prize is her objective, that she cares about the probabilities bearing upon it; her concern with the probabilities is only derivative. I shall not pursue this further, but will take it for granted that we are correct to assume what we all *do* assume, that it is rational – indeed, rationally required rather than just that it is not irrational – to predicate one’s actions upon probabilistically expected outcomes.¹⁸

Assuming that we are correct to regard a means, M_1 , as better than another, M_2 , when M_1 is more likely than M_2 to result in the realization of the objective, E , we can describe that by saying that the sufficiency of a means for an end is a matter of degree, with the degree given by the comparative probabilities, and that a rational agent will (*ceteris paribus*) select the more over the less sufficient means.¹⁹

¹⁸ This issue first struck me about 1994. Peirce raised the issue in “The Doctrine of Chances” (1957/1878, 64ff.)

¹⁹ It is plausible that if this is so, there is a significant further constraint on the correctness of instrumental reasoning. I have said nothing so far about the source of the probability judgments upon which an agent would have to rely, but whatever their source, sets of probabilities can fail to be coherent. Even without assigning definite numerical values, we can see, for example, that it is not possible that A is more likely than B , B more likely than C , and C more likely than A . So, if an agent should rely upon probability judgments in deciding what to do, he is less likely to achieve his objectives if the probability judgments upon which he relies are not coherent. Therefore, he has a reason, in terms of his objectives, for making his probability judgments coherent. This is of more than theoretical importance because people are not, in general, very good at assessing probabilistic reasoning. We not only make mistakes (which might be explained by carelessness, the difficulty of the assessment, or inadequate time), but *systematic* mistakes. For discussion of many of these, see Dawes 1988.

3.12 The Normative Control Conditions

Two significant points emerge from the foregoing discussion of the cases covered by the simple instrumental paradigm. First, when a means, M , is necessary to the achievement of an end, E , since we have stipulatively denied the relevance of any other considerations, then the agent, if rational, must select M from among his various options. The end, E , serves as a principle of selection from among the agent's options (M , N , O , etc.) that, given the circumstances, uniquely picks out M . When a means, M_l , is not necessary to the realization of the objective, E , but is in some other way contributory to it, then the agent must, if rational, select either M_l or some alternative, M_n , that does at least as well at contributing to the realization of the objective and must select one of these in preference to any option, N , that does not contribute or does not contribute as well to achieving the objective.²⁰ Options can be sorted, in terms of E , as better or worse (when M is necessary to realizing E , into one that is good as contrasted with all others, which are not).

Second, normative force flows from the end to the means, and not *vice versa*. There is reason for selecting M in terms of E , but none, either for or against E , in terms of M . Of course, M may be an objective which anchors further instrumental reasoning about what contributes to achieving or bringing it about, but that does not alter the point: M may

²⁰ This may require some qualification, in connection with issues about maximization and satisficing. In particular, it will be reasonable to adopt a means, M_m , which does not contribute as well to the objective, E , as some alternative, M_n , when it is not well-defined what is an optimum with respect to E . M_n , for example, may promise to make me wealthier than M_m , but it may be not be well-defined what an optimum of wealth for me is, and therefore possibly not well-defined whether M_m or M_n would move me closer to the optimum.

have normative force relative to some further means, M' , but the normative or reason-giving force still flows uni-directionally from end or objective to means.

Now, it might be thought that this is just an artifact of the stipulation imposed in the initial description, that there are no other relevant considerations. After all, in ordinary thought, we do consider the acceptability of ends in the light of means, do sometimes judge that an end is not worth having if it can only be had by means we are unwilling to adopt. A person may decline an opportunity to make money by cheating. On a larger scale, a research institution may accept stringent limits on experimentation on human subjects in the study of a disease, where such experimentation might credibly promise to advance the search for a cure.

This is all true, but I do not think it alters the basic point for two reasons. First, in at least some cases, putting matters in this way misdescribes what is going on. The objective may be, e.g., not just to make money, but to make money honestly. Isolating “making money” as the end to which various means, such as cheating, may be considered, may misrepresent the agent’s actual end:²¹

Adopting something as a goal is not just a matter of attaching a positive value to its accomplishment and counting this in favor of any action that would promote it (unless this is overridden by considerations coming from elsewhere). When we “adopt a goal” we normally give that goal a particular status in our lives and in our practical thinking, such as the status

²¹ This has obvious connections to questions about constitutive means, which will be further addressed below.

of a long-term career objective, or of a whim, or of something that we want to do sometime on a vacation. That is to say, the intentions that constitute adopting the goal specify the kinds of occasions on which it is to be pursued, the ways it is to be pursued, and so on. So the limitations indicated by the qualification that other things must be equal include conditions determined by our understanding of the goal and the way it is a goal for us, not just limitations imposed by other values that might “override” it. (Scanlon 1998, 86)

If this is the sort of case under consideration, then it is not really an exception to the thesis that normative force flows uni-directionally from end to means. Only misdescription of the end makes it appear that it is being judged unacceptable in the light of the means. But, of course, there are other cases. An analysis of the above type cannot work when some restriction on how an end is to be pursued is not (once it is correctly described) part of the end itself. So, there is a second point: the stipulation that there are no other relevant considerations *is* doing substantive work. What it is doing, however, is serving to call attention to the fact that *if* some fact about or feature of the means makes a difference to the acceptability or rationality of action in the service of the objective, the difference that it makes must have its sources elsewhere. That is, features of the means may make a difference, but not *just* insofar as it is a means; insofar as *M* is a means (and nothing but a means) to the realization of *E*, *E* provides or may provide a reason for (or against) the

selection of *M* but not *vice versa*.

These two points, that an end serves as a principle of selection from among options available to the agent²² and that reason-giving force flows uni-directionally from end to means, I shall label together as *the normative control of means by ends*. Such normative control, I believe, is quite generally characteristic of instrumental reasoning. Further, I suggest that the normative control of means by ends is not only characteristic, but both necessary and sufficient for an instance of practical reasoning to count as instrumental.²³

This may appear truistic, but I know of no exceptions, neither of any clear cases of instrumental reasoning that, on one hand, fail to satisfy either of the conditions of normative control nor, on the other, of any further condition that seems essential for a tract of reasoning to count as instrumental.²⁴ If this is correct, then we can make use of the twin conditions of the normative control of means by ends as a marker to recognize instrumental reasoning in less familiar settings. In the next section, I shall briefly discuss what I take to be uncontroversial extensions of the simple instrumental paradigm.

3.2 Extensions of the Simple Instrumental Paradigm

My main concern to this point has been with instrumental reasoning in which only a single end or objective bears upon action, and there has been only glancing or

²² The options may, as noted above, only be counter-factually available if the agent has only a single option, that option being in fact a means to the objective.

²³ Such reasoning may, of course, have further features in virtue of which it also counts as some other form of practical reasoning. Some tract of deliberation may be more than instrumental without ceasing to be instrumental.

²⁴ On one level, I hope my characterization will appear truistic, for I do not intend to offer any further argument for it than to point to its presence in examples.

parenthetical reference, by way of *ceteris paribus* clauses or mention of other factors that would make a difference, were they present, to the possibility that more than one objective may be relevant to what it is reasonable to do. This has been, of course, deliberately artificial, aiming at eliciting the basic features of instrumental reasoning from the simplest available cases, but in much, perhaps most, reasoning that is uncontroversially instrumental, more than one end or objective is relevant.

3.21 Economizing

Perhaps the most familiar form this takes is what may be called *cost* or, perhaps better, *economizing*. I do not mean simply monetary cost, though that provides a useful example. When an agent has determined to pursue some objective, one question relevant to the selection of means is which will interfere least with other objectives. If he faces, say, a pair of options with regard to means that are equally good at promoting the objective, then if one of the options uses fewer resources than the other, resources which could otherwise be put to use in the service of some other objective, then the agent has reason to select that option.

Generally, the employment of some means to an objective requires that the agent give up something – time, effort or other resources – that could have been employed differently or on behalf of different objectives. What must be given up, which could have been employed differently, is the cost of pursuing that objective. Suppose that an agent has determined upon the pursuit of some objective, E_i , and that the only relevant

alternative to E_1 – the only other objective that might be pursued with the same resources – is E_2 . Suppose also that E_1 can be obtained or promoted equally well by two different employments of those resources, M_1 and M_2 . Then, if M_1 and M_2 can be ranked against one another so that, after one of them is employed on behalf of E_1 , what remains of the agent's resources is either more or else less satisfactory for the pursuit of E_2 than if the other had been employed, the agent has reason to select the employment of resources that is more satisfactory in terms of E_2 . Though M_1 and M_2 are equally good in terms of E_1 , one of them is better than the other in terms of E_2 .²⁵

To return from the abstract to the familiar, a person may have a set of options that are equally good at achieving some goal but one of which is cheaper than others. If so, she has reason to select a cheaper over a more expensive option. The reason this can work is that monetary cost is a useful proxy, for many cases, for what must be given up to achieve the goal. Had the money not been spent in achieving that goal, it would have been available for others. Or, had less been spent in achieving that goal, more would have been available for others.²⁶

²⁵ A variation upon this pattern may occur in the time-ordering of pursuits. Of a pair of objectives, E_1 and E_2 , a person might be better placed to achieve or promote both if she pursues E_1 first and E_2 second. Some resource needed for the pursuit of E_1 might become unavailable if E_2 is pursued first, but not *vice versa*. Another possibility is that the pursuit of E_1 before E_2 is better (or worse) in terms of some other objective, E_3 . Time-ordering of pursuits is not, however, always best assimilated to economizing. For some cases, it may be better regarded as an instance of constitutive reasoning (to be discussed more fully below).

²⁶ Of course, not everything has a market price, and so, not everything can be compared to everything else in terms of relative market prices, but that does not affect the point that it is normally true that monetary assets could be employed differently in the service of some other objective. Goals that cannot be pursued or advanced through monetary means are, for that very reason, not in competition (along a monetary dimension) with goals that can be pursued through monetary means. The reason for selecting the less costly of otherwise equally good ways of pursuing a goal only depends on the assumption that there is *some* other goal that could be pursued through monetary means, not that all goals can be evaluated in

3.22 *Before Economizing*

Though it is perhaps the most familiar, economizing or attention to costs is not the simplest kind of case in which more than one objective may bear on action. This is evident most plainly in the stipulation above that the agent has determined upon the pursuit of some objective. So far as the discussion has gone, it is only after some objective has been selected that cost enters the picture,²⁷ namely, the relative costs of different means for pursuing that objective in terms of other objectives. But questions can be raised on two fronts, first, as to whether there are ways that more than one objective may be relevant to a decision *prior* to the selection of some one of them for (immediate) pursuit, and second, as to whether and how, beyond that, the various objectives an agent has are relevant to the selection itself. Each deserves at least brief treatment.

3.221 *The Relevance of Multiple Ends: Prior to Selection*

Suppose an agent has some array of ends, E_1 , E_2 , and E_3 , and that each of these represents one disjunct of a binary alternative: each will be achieved or else not achieved, and there are no intermediate degrees to which it may be promoted or advanced. Suppose further that this is a complete list: there are no other ends (nor any other considerations)

terms of money.

²⁷ Cost may enter the picture in another way when it serves as a limiting factor on whether an objective is selected. For example, when I inquire as to the ticket price in order to determine whether to go to the concert, that presupposes that there is at least some other objective that bears on the decision. A case where cost is relevant to the initial selection of an objective is best understood in terms of the relevance of some combinatorial principle, as discussed in §3.222.

that bear on how she should act. Can we say anything about how it is rational for her to act in such a case without assuming that she has already selected one of them for immediate pursuit? Or, equivalently, since ‘rationality’ takes its meaning in part from the contrasting cases of that which is less or not rational, can we say anything about how it would be *unreasonable* for her to act? It depends on what options for action she has. Plainly, there are eight possible outcomes that could be correlated with options she faces. She can achieve

- A. none of the ends, E_1 , E_2 , or E_3 ;
- B. E_1 and neither E_2 nor E_3 ;
- C. E_2 and neither E_1 nor E_3 ;
- D. E_3 and neither E_1 nor E_2 ;
- E. E_1 and E_2 , but not E_3 ;
- F. E_1 and E_3 , but not E_2 ;
- G. E_2 and E_3 , but not E_1 ;
- H. E_1 , E_2 , and E_3 .

If we let the letters, A through H, stand for options she could have that would bring about the corresponding outcomes, then it would be unreasonable for her to select A if she has any other option and unreasonable for her to select anything but H if H is one of her options. In between, E and F are both better than B; E and G are both better than C; and F and G are both better than D. With no more information than has been given, there is no way to tell which, if any, is better or best of B, C, and D or of E, F, and G. Nor is there any way to determine how B compares with G, how C compares with F, or how D compares with E.²⁸

²⁸ There is an interesting parallel here to the economists’ notion of Pareto-improvements. A change is Pareto-improving when it is advantageous to some – at least one – and disadvantageous to none. The same patterns of relative superiority between some options and indeterminacy between others appear.

Despite the indeterminacies just noted, the general point is that, given some set of objectives and some set of options that differently affect the realization of those objectives, there are or may be rational requirements pertaining to which of the options should be selected that do not presuppose that the agent has *already* selected some objective for immediate pursuit. There are ways in which an agent can be reasonable and also ways in which she can fail to be reasonable.

3.222 The Relevance of Multiple Ends: Combinatorial Principles

What, though, about the indeterminacies? This is a way of raising the question as to how to select one objective for immediate pursuit when the answer is not given simply by specifying the objectives and the agent's options. It seems that we often do make judgments of this sort, considering not just the fact that we have multiple ends which can be differentially advanced by the options available to us, but also considering their relative importance.

Suppose for the sake of illustration that the agent with the set of ends discussed above has only *B* and *C* as options. Assume also that the case differs in that the ends can be promoted to varying degrees. If she selects *B*, she will achieve one of her ends (to some degree); if *C*, then a different one (to some degree). Putting matters the other way around, selecting *B* is a decision against promoting E_2 , and selecting *C* is a decision against promoting E_1 .

If one of the options is to be rationally better than the other, then, in addition to the

relevant ends, E_1 and E_2 , there must be some combinatorial principle on the basis of which a choice of which to promote can be made.²⁹ This can take either of two principal forms. It can be a priority rule according to which one of the objectives, such as E_1 , is more important or valuable than the other and thus is to be promoted in the event of a conflict. Or it can be some weighting function that assigns greater importance or value to promoting, say, E_2 to *this* degree than promoting E_1 to *that* degree.³⁰ If it is the latter, there will be no implication that promotion of E_2 is always to be ranked ahead of the promotion of E_1 ; it may be that there is some degree to which E_1 can be promoted that would, if it were an available option, be more important or valuable than the promotion, to the degree possible in the circumstances, of E_2 .³¹ Whatever the form it takes, I think all that is required – that is, all that is required for it to *be* a combinatorial principle, though there are surely additional requirements upon its being reasonable or acceptable – is that it reduce indeterminacy. It need not decide all issues in order to genuinely decide some.

To this point, I have pointed to the need for one or more combinatorial principles if certain indeterminacies are to be overcome and have implied that not all possible combinatorial principles are equally acceptable, but have said nothing about why some

²⁹ I take no position here on whether the combinatorial principle is itself best conceived as being or being based upon some end.

³⁰ There are other possibilities. One would be a hybrid principle that applies, say, a priority rule below some threshold and a weighting function above it. Still others can be imagined.

³¹ There may be ways in which a weighting function can be represented as a (very complicated) priority rule or in which a priority rule can be represented as a weighting function. So far as I know, nobody is interested in the reduction of weighting functions to priority rules. It looks like a lot of work for no theoretical payoff. The possibility of reducing priority rules to weighting functions is more interesting but is crucial only if it is supposed, as many decision theorists might suppose, that rational choice in cases for which a combinatorial principle would have to be invoked must somehow be based on a weighting function. I think that is unlikely, but am content for the present to leave it an open question, while relying simply on the intuitive difference between weighting functions and priority rules.

particular combinatorial principle should be judged to be correct or better than some alternative. And I do not think we are in a position to do this here. At least some imaginable, if not very plausible, combinatorial principles can be ruled out as clearly unacceptable, such as the principle that tells an agent to pick one of his objectives at random and rank achieving it below all others or one that instructs him to list options in alphabetical order and select the first. Thus, we can say that the agent would be mistaken to select one of these principles, but that hardly narrows the field. In part, this is because they are so clearly unacceptable that no one is tempted by them in the first place, but more importantly, there are many, perhaps infinitely many, possible combinatorial principles that are not obviously unacceptable. All that can be said at this point is that *if*, in addition to some set of objectives, an agent accepts some applicable combinatorial principle, we will be able to say more about what it is reasonable for her to do, in terms of those objectives together with that combinatorial principle, than could be said in terms of the set of objectives alone.

In this and in the preceding sections, I have discussed the simple instrumental paradigm, where only a single objective bears upon action, and extensions of it to accommodate the relevance of multiple objectives. In the next section, I consider the relation of a more controversial kind of practical reasoning to an instrumental framework, *constitutive reasoning*.

3.3 Constitutive Reasoning

Consider constitutive reasoning.³² A constitutive means to some objective is one that at least partially constitutes the objective to which it contributes. The objective cannot be adequately specified entirely independently of the constitutive means. A constitutive means can be contrasted with an external means in that, for an external means, its contribution to promoting or achieving some objective is causal and the objective can be specified independently of the means. The distinction will receive further discussion in the next chapter, and, for the most part, the points I wish to make here are obvious and, I hope, uncontentious. Still, in considering the simple instrumental paradigm and its extensions, only external means to objectives were in view. When constitutive reasoning is brought into the picture, some issues take on a new form or have to be reconsidered from a different perspective. There are two principal questions which tend to flow together. The first is whether constitutive reasoning is genuinely *reasoning*. The second is whether it is genuinely *instrumental*. Obviously, if it is not reasoning at all,³³ then it is also not instrumental reasoning. On the other hand, since it is not clear what it would be if it were not instrumental, doubts about its being instrumental are likely to spill over into questions as to whether it is reasoning at all. Of course, we can also run this argument in

³² Since part of what will be addressed in the present section is whether what I am calling constitutive reasoning is really practical reasoning at all, the label may appear question-begging, and it might be thought that quotes – “constitutive reasoning” – are more appropriate. I prefer to avoid cluttering the text with such devices in favor of noting the point here. My use of the phrase without the quotes may be taken as a promissory note to be redeemed by the subsequent argument.

³³ It is important for present purposes that it be *practical* reasoning, reasoning about what to do, and that is what I am assuming to be implied by the claim that constitutive reasoning is indeed reasoning. If constitutive reasoning were to turn out to be some kind of theoretical reasoning, that would be to no avail.

reverse to say that if it is instrumental reasoning, then it must be reasoning. Since I have already argued that the normative control of means by ends is definitive of instrumental reasoning, that is the path I will take: I intend to test the credentials of constitutive reasoning by examining whether it meets the normative control conditions. (Proceeding in this way has the advantage that I can set aside worries about how constitutive reasoning, if not instrumental, still qualifies as reasoning.)

3.31 How Constitutive Reasoning Satisfies the Normative Control Conditions

The conditions of normative control, remember, were (a) that the end functions as a principle of selection from among candidate means, grading them as better or worse in terms of the end, and (b) that reason-giving force flows uni-directionally from the end to the means. Why might it be doubted that constitutive reasoning qualifies as instrumental by meeting these conditions? There seem to be three sources of worry here. The first is that there might not be constitutive *means* to an objective at all. The second is that an account that assimilates constitutive reasoning to instrumental reasoning runs the danger of trivializing the conception by construing instrumental reasoning so broadly that it would apply to *anything*. The third is that, if constitutive means partially or wholly constitute the objective to which they are means, it becomes unclear what either of the normative control conditions actually requires.³⁴

³⁴ I have actually found only the second of the three raised in the literature. The remaining two are, so far as I know, my own (though Sarah Broadie [1987] touches upon the third, from a different angle and with different concerns than mine). I address all three in the attempt to pose the toughest challenge I can to the thesis – which I accept – that constitutive reasoning is best understood by assimilating it to the instrumental. (It is not difficult to find thinkers who *assume* either that there is no problem with the

The motivation for the first worry lies in the fact that, for anything to be an objective at all, something or other must be constitutive of it. There must at least be some state of affairs, believed or presupposed to be possible, which would be the realization of the objective. However, the fact that something or other would constitute the objective is not sufficient for part or all of what constitutes it to be a *means* to that objective. Means are adopted, selected or performed because of their (expected) contribution to achieving or promoting an objective. What reason do we have to suppose that something constitutive of an objective is really a means rather than just a constituent? Or, to put it differently, what is the import of the requirement that means be adopted *because of* their contribution to an objective? Why not say instead that there is just an objective, constituted by various actions, events or states of affairs, but that none of these are to be singled out and contrasted with others as means? On such a reading, the contribution of a means to an objective would then have to be understood as *causal* contribution, something that brings about or helps to bring about the objective, and so, all means would have to be external.

I think a full answer has to wait upon distinctions yet to be developed in response to the other two worries, but a beginning can be made here. The first point to note is that an objective may be some state of affairs to be brought about and with respect to which the only relevant actions are external means. There need be no *action* that is constitutive of the objective itself. (The objective is to have a flower garden; the planting, weeding and fertilizing may be external means.³⁵) A means also may be a state of affairs, an

assimilation or else that no such option is available; it is much more difficult to find *arguments* for one or the other position.)

³⁵ They could be constitutive means if the objective were ‘gardening’ or ‘being a gardener.’

objective, which is itself related to action only by way of further external means. (The lock on the jail cell is a means to prevent the escape of prisoners.) So long as this is the case, there is no place for introducing any notion of constitutive means.

Whatever a constitutive means is, it must be more intimately related to action than that. For a constitutive means to be distinct from a state of affairs to be promoted by external means (if at all), it must be an action, activity, practice or disposition with respect to action. Is this also enough? If, abbreviating, some action is part of the objective, is that sufficient for it to be a constitutive means to the objective? I think the answer is negative. Consider the objective of raising one's arm. There is an action without which one could not achieve that objective, namely, raising one's arm. This action is not some causal precondition for arm-raising; it just *is* the raising of one's arm. But it certainly sounds odd to say that one raises one's arm because of or for the sake of its contribution to arm-raising. Something more is needed if we are to make a respectable place for considerations of constitutive means to objectives, but to pursue it, we need to consider the other two worries.

Christine Korsgaard raises the second issue:

But the instrumental principle is nowadays widely taken to extend to ways of realizing ends that are not in the technical sense 'means', for instance to what is sometimes called 'constitutive' reasoning. Say that my end is outdoor exercise; here is an opportunity to go hiking, which is outdoor

exercise; therefore I have reason to take this opportunity, not strictly speaking as a means to my end, but as a way of realizing it. This is a helpful suggestion, but it should be handled with care. Taken to extremes, it makes it seem as if any case in which your action is guided by the application of a name or a concept to a particular is an instance of instrumental reasoning. Compare, for example: I need a hammer; *this* is a hammer; therefore I shall take *this*, not as a means to my end but as a way of realizing it. In this way the instrumental principle may be extended to cover *any* case of action that is self-conscious, in the sense that the agent is guided by a conception of what she is doing. (1997, 215-216)

An initial point to note is that what Korsgaard calls “in the technical sense, ‘means’” is equivalent to what I have been calling ‘external means.’ This is only a terminological difference between us, turning on the fact that I see no reason to deny that some things that are not external means are nevertheless genuinely means. It is, however, a terminological difference that makes it more difficult for her to state the next point, for what she means by ‘the instrumental principle’ is a principle that affirms, roughly, that we have reason to take the means to our ends. If only external means are genuinely (or “in the technical sense”) means, then it is awkward and possibly misleading to speak of extending the instrumental principle to cover constitutive reasoning. The extended principle would have to say that we have reason to take the means to our ends and also

have reason to adopt ways of realizing our ends that are not means. (Why is that not just two principles rather than an extension of one?) If, on the other hand, we take a means to be an action taken or state of affairs selected or brought about for the sake of some objective,³⁶ then there is no problem with extending the instrumental principle to cover constitutive reasoning. It is extension only in the sense of recognizing a relevant similarity between cases that may be taken to be central – the employment of external means – and others that may be taken to be more peripheral – the employment of constitutive means – rather like the extension of the term ‘mammal’ to include cetaceans.

Setting aside terminological issues, Korsgaard still has an important point pertaining to the danger of trivializing the conception of instrumental reasoning.³⁷ (It might be regarded as a version of the first worry viewed the other way around. The concern there was whether *anything* could count as constitutive reasoning. Here, the concern is, if we agree that anything counts as constitutive reasoning, whether we do not have to let in *everything*.) Her thought seems to be that what has to be avoided is the acceptance of some such principle as:

- (1) For any two things (actions, states of affairs, etc.), if one is an
objective and the other a way of realizing that objective through action,

³⁶ It is important that ‘means to’ and also ‘ways of realizing’ an objective be understood to be such from the agent’s point of view; a means or way of realizing is selected or brought about because of the relation in which it is understood to stand to the objective.

³⁷ Note that she does not put even this point as more than a reason for caution: describing ways of realizing ends as constitutive means to those ends can be taken to extremes; there is no suggestion that it has no place when the extremes are avoided.

then the second is a means to the first.

That principle, if accepted, would leave open the possibility that the two are identical.

Then, if the objective happens to be the performance of some action, since the performance of that action is (of course) identical to itself, the performance will be a way of realizing the objective, and so, a means to realizing it. Every action undertaken in virtue of its falling under some concept or description will then be instrumentally rational. (And only the inadvertent or unintended will remain to provide lodging for lapses of instrumental rationality!)

Now, it is fairly plain what needs to be done to avoid this. Some restriction upon (1) is needed to rule out the possibility that the means and the objective are identical. So, it might be suggested that what we need is:

- (2) For any two things (actions, states of affairs, etc.), if one is an objective and the other a way of realizing that objective through action, then the second is a means to the first, provided that the two are not identical.

I do not think that is quite adequate as it stands. At the least, it stands in need of clarification, but the best way to bring that out is to proceed to the third issue mentioned above. That was the problem of how to understand the conditions of normative control of

means by ends when it is supposed that the means partially or wholly constitutes the end. How, if we cannot say what the end *is* apart from the means, can it be that the end serves as a principle of selection from among options available to the agent – that is, to distinguish means from non-means and better from worse means – and that reason-giving force flows uni-directionally from end to means?

To sort these issues out, something further needs to be said about the relation of constitution. First, we need to avoid misunderstanding what is being claimed when some means is said to be constitutive of an objective. A means is adopted for the sake of the objective, and, when the means is constitutive, it at least partially constitutes that objective. But we should not think of the relation of the constitutive means to the objective on *this* model: The objective is *A* and *B* together; the means is *A*.³⁸ To hold that some means is constitutive of an objective is *not* just to hold that it is a member of a set of elements, perhaps including other means and perhaps including some things that are not means, with the compound of elements being the objective in question. The basic reason is that the introduction of the compound end or objective would not do any theoretical work unless it made a difference and therefore was more than just a compound – if, e.g., it introduced some combinatorial function for relating its constituents.

Rather, there must be some way of understanding the objective *as* an objective, as something sought, aimed at or to be performed and in terms of which some criteria (which need not be complete) can be specified for whether something further – something else

³⁸ And *B* is either some other means or else something constitutive of the objective though not a means to it.

sought, aimed at or to be performed – counts as contributing to it, advancing it or being necessary to it, and therefore as a means to it. To use an earlier example, I may play tennis for the sake of exercise, and the tennis I play is constitutive of the exercise I get – perhaps, if I get no other exercise, wholly constitutive of it. At first glance, this looks as if it will violate the non-identity condition between objectives and means proposed in (2) above. The tennis I play and the exercise I get are identical: the very same events are, on the one hand, my playing tennis, and on the other, my getting exercise.

A closer look at the example will, I think, suggest the way in which (2) needs to be revised. Though the tennis I play is wholly constitutive of the exercise I get, I have some understanding of what exercise is independently of understanding what tennis is. Other activities than the playing of tennis (e.g., volleyball, walking, swimming or thumb-twiddling) can be considered as forms of exercise and can be compared with tennis along various dimensions – how strenuous they are, how appropriate to someone in my physical condition and so on. In terms of what exercise is, in combination with other relevant parameters, tennis can be assessed as better or worse exercise for me. The tennis I play and the exercise I get are *extensionally* equivalent, but *intensionally* distinct. In aiming to get exercise, I aim to do something that satisfies, at least reasonably well, the criteria by which I distinguish exercise from non-exercise and better from worse exercise. So, the principle needed to replace (2) is:

(3) For any two things (actions, states of affairs, etc.), if one is an

objective and the other a way of realizing that objective through action, then the second is a means to the first, provided that the two are not intensionally identical.³⁹

Once it is recognized that the objective must be intensionally distinct from the means, even when the means partially or wholly constitutes the objective, there is no problem in seeing how the normative control conditions, that the objective serves as a principle of selection from among means and that reason-giving force flows uni-directionally from objective to means, can be satisfied by constitutive reasoning. Even when the objective is wholly constituted by the means, the two are only extensionally identical. The normative control conditions can be satisfied because the objective is not intensionally identical to the means.

3.4 In Search of Normative Underpinnings

We all find instrumental reasoning compelling. Whether it is a matter of causal contributions to or constitutive ways of realizing ends or objectives, we think that it is, at least *ceteris paribus*, rational to act on conclusions reached through instrumental reasoning and that we are subject to criticism on the count of rationality – that we are irrational or less rational – to the extent that we fail to guide ourselves by such conclusions. But why? What exactly is wrong with instrumentally irrational action? I

³⁹ The relevant descriptions in terms of which two items are judged to be intensionally distinct (or not) must be available to the agent in her acting and decision-making. See note 36.

shall try to provide a partial answer here, but, in approaching it, we need to see why there are problems (not, I think, insuperable problems) with the most obvious suggestion.

The natural response to questions about the normative force of instrumental reasoning is that instrumental reasoning has just the normative force of the ends from which it proceeds: we are justified in acting in instrumentally rational ways because otherwise, we will fail to realize (or are less likely to succeed in realizing) our objectives.

That natural response, however, may not be available to us. For consider the instructive analogy between correctness in instrumental reasoning and validity in logic. In a valid argument, truth is transmitted from premises to conclusion: if the premises are true, then the conclusion must also be true. Similarly, in correct instrumental reasoning, normative force is transmitted from ends or objectives to means. Instrumental correctness (call it *instrumental validity*) is a practical analogue of deductive validity.

Whether an argument is valid, however, is not the only, or sometimes the most important, question that can be raised about it. We can also ask whether it is sound, whether, in addition to being valid, its premises are true, so that we can be assured that its conclusion is true. One way to bring out this point with regard to logical validity is to add that falsehood is transmitted *backward* from the conclusion to the premises: if the conclusion is false, then so must be at least one of the premises from which it was validly inferred.⁴⁰ In a valid deductive inference, the truth of the conclusion follows from the

⁴⁰ This way of characterizing logic comes from Popper: “[D]eduction ... is valid because it adopts, and incorporates, the rules by which truth is transmitted from (logically stronger) premises to (logically weaker) conclusions, and by which falsity is re-transmitted from conclusions to premises.” (1965, 64) In this form, it is not sufficient to cover all systems of logic, since it does not address systems with truth-values additional to ‘true’ and ‘false’ or that employ, say, probability metrics. However, it is not necessary

assumed truth of the premises: the conclusion must be true, *given* the premises. But, though validity is a matter of the relation between premises and conclusions, the truth of the premises or of the conclusions themselves is not. At least sometimes, there are ways of checking whether the premises or conclusions are true that do not depend on their place in a particular argument.⁴¹

And here we may encounter a disanalogy between instrumental and deductive validity. What, if anything, is the instrumental analogue of falsehood (or of truth)? Can there be any defect – call it *normative deficiency*⁴² to put a label on it – of the conclusion of an otherwise correct course of instrumental reasoning that infects the end or objective with which the reasoning began? On one ground it might be thought that there cannot, for if there were any such defect, it would undercut one of the conditions of the normative control of means by ends, that normative force flows uni-directionally from ends to means. In terms of the means, there would be reason for rejecting, altering or qualifying the end.

This, however, would be a misunderstanding. First, when the normative control conditions were introduced, the claim was not that there could be no reason, based on for my purposes to enter into these complications.

⁴¹ Some other argument, perhaps unstated, may be involved in checking. (I do not think it must be. Checking does not always involve inference.) From the premises, “all cows are green” and “Bossie is a cow,” it follows that Bossie is green. But we can examine Bossie to determine that she is not green. It might be claimed that we are thereby relying implicitly on some such argument as “if Bossie were green, she would look green (given current lighting, etc.), but she does not look green; therefore, she is not green.” Even if this is so, it is sufficient to point out that this is a *different* argument and, therefore, that the truth of the first argument’s conclusion is not simply relative to the premises from which it was inferred.

⁴² To avoid confusion, this normative deficiency should be understood as relative to context. Whatever the feature in virtue of which some conclusion of a particular course of instrumental reasoning is judged defective, that feature may not affect all courses of instrumental reasoning in the same way. It may be relevant *as a defect* in some situations but not in others.

some feature of the means, that could be relevant to the acceptability of the end, but rather that it could have no such feature *just insofar as it is a means* – that is, just insofar as it is contributory to the end or objective. But that does not rule out the possibility that the means could be normatively deficient by virtue of some other feature. Second, to pursue the analogy with deductive validity, a defect in the conclusion of a valid argument – namely, that it is false – implies that there is a defect also in at least one of the premises. If no premise were false, then the conclusion would not be false, either. So, it runs the risk of misunderstanding to speak of the falsehood of the conclusion being transmitted backward to the premises, as if there were nothing wrong with the premises apart from the conclusion being or being found to be false. It is a misleading metaphor to say that the falsehood discovered in the conclusion infects the premises; rather, it shows that the premises were already defective. Similarly, if some defect, some normative deficiency, can be found in the conclusion of an otherwise correct tract of instrumental reasoning, that shows that there was already a defect in the objective or objectives⁴³ from which the reasoning proceeded.

So, in principle, we can admit the possibility of normative deficiency in the means to which we are directed by a course of instrumental reasoning, and, if we find it there, we will have to admit some normative deficiency in the ends as well. But this may seem not to be any progress at all, for it appears that all we have reached are *conditional* claims: *if*

⁴³ If some combinatorial principle is employed and if that principle is not itself to be conceived as an end (see note 29), then the defect might be there instead. I shall assume that if combinatorial principles are not to be conceived as ends, then their role is analogous to rules of inference, the non-observance of which undermines the claims of a course of reasoning to be considered valid.

we find a normative deficiency in the means, then there must be some normative deficiency in the ends. That tells us neither that there ever is any normative deficiency in means nor how we find that there is. Instrumental rationality appears to be, as Darwall says, a matter of principles of *relative* rationality,⁴⁴ rationality relative to ends, desires or preferences that are given prior to or apart from any particular course or instance of instrumental reasoning in which they figure. But then, as Darwall also says, principles of relative rationality are principles of the *transfer* of reasons. Whatever reason there is for the objective is transferred to the means. That seems to get us no closer to saying that there ever *is* any reason for an objective. And if there never is a reason for an objective, then, if instrumental reasoning has just the normative force of the ends from which it proceeds, it has none.

Matters, I think, are not so desperate. We can say something about the normative force of instrumental reasoning without prior assumptions about the reasons that there are or may be for ends.⁴⁵ We can see this by reconsidering the parallel with deductive validity. In introducing the parallel, I noted that the truth or falsity of some premise or conclusion of a valid argument is not just a matter of its relation to other elements of the argument. But we can still recognize that, if we reach a contradiction in the conclusion of a valid argument, *that* must be false,⁴⁶ and therefore so must be at least one of the premises

⁴⁴ Darwall 1983, 15-16.

⁴⁵ I do not mean that those questions are unimportant; I shall say more about them later. But the normative force of instrumental reasoning does not depend, at least not entirely, upon the answers to those questions.

⁴⁶ I omit from consideration paraconsistent systems of logic in which $\sim(P \ \& \ \sim P)$ is not a theorem.

from which it was inferred. For the special case in which a conclusion is contradictory, no further checking is needed to determine that some premise is false.

Is there any analogue to this in instrumental reasoning, any way to recognize normative deficiency in the conclusion or premises of an otherwise correct tract of reasoning? I think there is. Consider that in instrumental reasoning the premises will include both some objective or objectives to be achieved and claims about how the world is, about what is causally or probabilistically related to what. Then, there are at least two ways that we can recognize defects in the premises or conclusions of some course of instrumental reasoning.

First, the objectives from which the reasoning proceeds might be in conflict. This might be thought unlikely, but I do not actually think it is, especially when the reasoning is complex and takes into account the bearing of many objectives upon action. The reason is that testing for consistency is a problem subject to combinatorial explosion. If there are two propositions, *A* and *B*, then one test suffices – whether *A* is consistent with *B*. Add a third, *C*, and four tests are necessary – whether *A* is consistent with *B*, whether *A* is consistent with *C*, whether *B* is consistent with *C*, and whether *A*, *B* and *C* are consistent together. Add a fourth, and eleven tests are needed (which I won't detail). Matters only get worse from there. Since the consistency of a set of ends can be modeled in terms of the truth of a set of propositions (*E*₁ is achievable, *E*₂ is achievable, ..., *E*_{*n*} is achievable), the same problem applies. No one who has a multitude of ends can reasonably be certain that they are all consistent. For the simplest case, a person might have set himself both to

achieve and to prevent the achieving of some objective. Action taken for the promotion of *E* will defeat his attempts to prevent *E* and *vice versa*. If there is this kind of conflict between or among objectives, nothing he can do will serve to realize his objectives (though he may be able to realize one or some among them). That would surely be sufficient to show that he was mistaken to think or assume that *none* of the objectives in conflict (nor their combination) was normatively deficient.

Second, it may be that the pursuit of some objective or system of objectives cannot, though there is nothing inconsistent in the statement of the objectives themselves, be expected to be successful, due to the factual and causal relations that are also among the premises. For example, it may be that, given the way the world is, including in particular the way that people distribute trust, the single-minded pursuit of money is likely to have a lower monetary pay-off than some alternatives. Single-mindedly pursuing money would not then be a good way of getting money or not as good a way as, say, devotion to a career. If that is so, it would show that, *as the world is*, there is something normatively deficient in the single-minded pursuit of money.

We can identify a feature that is common to both types of case. They are, in different ways, cases in which action in the service of ends is *self-defeating*. A person who guides himself by certain objectives or sets of objectives, under certain conditions, either cannot succeed in realizing those objectives or is less likely to succeed in doing so than if his objectives were different. And this is, though minimal, a conclusion that can be reached without presupposing anything substantive about which objectives there is or is

not reason to pursue. Accordingly, I suggest that being self-defeating is the appropriate practical or instrumental analogue to the defect in theoretical reasoning that is displayed in reaching (and maintaining) contradictory conclusions.⁴⁷

In other words, instrumental reasoning has normative force *of its own* that does not depend upon prior assumptions about what it is reasonable to do or pursue. The claim that it has only the reason-giving force of the objectives from which it proceeds, and therefore none unless those objectives are presupposed to have reason-giving force, is mistaken.⁴⁸

3.5 Summary

Both because it is pervasive in and because it is perhaps the best understood department of our practical reasoning, instrumental reason deserves special attention from the constructivist. Centrally, what instrumental reasoning concerns is the relation between objectives, goals or ends and means – in the simplest and most paradigmatic cases, the relation between single objectives and means, actions or states of affairs selected for what is taken to be their causal contribution to bringing about the relevant single objective. More specifically, means or alleged means are graded as more or less (or not at all) appropriate to the relevant objective.

In these simple cases, we find two features which I label together as the normative

⁴⁷ In focusing upon the way that failures of instrumental rationality are self-defeating, I do not mean to be reducing something normative to something non-normative. If someone claims not to see what reason there is against a self-defeating course of action, I have no further argument to offer. (At least not of the present kind – I might be interested in arranging a series of bets as a contribution to his education!)

⁴⁸ I do not wish to suggest that there is nothing else to be said about the normative force of instrumental reasoning, just that at least this much can be said without appeal to any presupposed reason-giving force attached to the objectives that anchor tracts of instrumental reasoning.

control of means by ends, that an end serves as a principle of selection from among options available to the agent and that reason-giving force flows uni-directionally from end to means. I argue that these features, taken together, are both necessary and sufficient for a tract of practical reasoning to count as instrumental, and thus that we can make use of these conditions as markers to identify instrumental reasoning in cases that are less central or paradigmatic.

With the normative control conditions in hand, we can extend the account of instrumental reasoning to partial coverage of cases in which multiple objectives bear upon the selection of means and, by way of supplying a place-holder to stand for the results of further investigation, to the consideration of the relevance of some kind of combinatorial principles to more fully cover the bearing of multiple objectives. More importantly, I show that we can also comprehend within the account of instrumental reasoning what has been called constitutive reasoning, where the means deliberated about are taken to partially or wholly constitute the objective to which they are means, rather than simply to causally contribute to its production.⁴⁹

Finally, I consider more generally the normative force of instrumental reasoning with a view to addressing the concern that it has no normative force that is not derived from some normative weight or importance attached to the objectives from which it proceeds, and thus, if no such weight or importance is presupposed, that instrumental

⁴⁹ This is important for my purposes because the constitutive relation, as will become evident in Chapter Four, is needed especially for the eudaemonist account of the virtues. If the constitutive relation could not be understood as instrumental, that would at least be a source of further complication for the current project.

reasoning falls short of being, genuinely, *reasoning* about what to do. In response to this concern, I argue that, even without substantive evaluative presuppositions about the ends or objectives to which instrumental reasoning is anchored, we can see that there is normative force in such reasoning, because failures of instrumental rationality are or tend to be self-defeating. Instrumental reasoning has normative force that does not depend upon value attached to the ends it serves.