

REASONS, PRACTICAL REASON, AND PRACTICAL REASONING

Robert Audi

Abstract

The concepts of reasons as supporting elements, of practical reason as a capacity, and of practical reasoning as a process are central in the theory of action. This paper provides a brief account of each. Several kinds of reason for action are distinguished. Practical *reason* is characterized both as a capacity whose exercise is largely constituted by a kind of responsiveness to reasons and as governed by certain normative principles; and practical *reasoning* is described as a kind of mental process in which reasons figure as premises and, from those premises, a practical conclusion is drawn. Much of the paper undertakes two related tasks: to distinguish the main kinds of practical reasoning and the associated criteria of assessment and to formulate some important substantive principles of practical reason. These principles yield criteria of several sorts: logical, inferential, epistemic, and material. On the theory presented, although any (non-basic) intentional act can be grounded in practical reasoning, the same acts can be performed for the relevant reason(s) *without* being so grounded, and in either case their rationality depends on adequate support by the reason(s). One kind of reason is commonly thought to be captured by Kantian hypothetical imperatives. The final sections explore what constitutes a hypothetical imperative and what other principles are needed to account for practical rationality. A major conclusion is that in the domain of practical reason, if there are no categorical imperatives, there are no hypothetical imperatives either.

Reasons are central in understanding both practical reason and theoretical reason. But there are many kinds of reasons. Partly because of this, philosophical writing is often unclear about what counts as a reason and about how reasons are connected with various closely related elements. One of these is *reason*, as a general capacity. A second is reasoning, as an exercise of that capacity. My concern is the practical domain, but much of what I say also applies to the theoretical realm. I begin with the nature of reasons for action, proceed to connect these with practical

reason, and then connect both with practical reasoning. I am particularly interested in the assessment of practical reasoning and in how such reasoning bears on practical rationality.

I. Three overlapping categories of reasons for action

There are at least three main kinds of reason for action (similar conceptual categories may be identified for belief, which in this respect is analogous to action). The first kind is normative, the second motivational, the third explanatory.

Normative reasons are reasons (in the sense of objective grounds) there *are* to do something. They are reasons for anyone, or at least anyone of a certain general description, to do certain things, for instance (for normal persons) to wear coats in a cold wind and to make amends for wrong-doing. Some normative reasons are *person-relative*: reasons there are for a specific person, say me. The fact that it will help *my* friend can be a reason for me to do an errand.

A second broad category is that of motivational reasons. There are two main kinds. The first is *possessed reasons*: reasons someone *has*, such as my reason to wear a coat, which I have in virtue of needing warmth. The second kind is both possessed and motivating (and is described below). A possessed reason is subjective if based on a desire (which is not to imply that just any desire provides a reason). Reasons based on desires may be called *internal* to contrast them with normative reasons viewed as *independent* of what the agent wants and in that sense external and objective; but 'internal' can mislead because some possessed reasons are also normative, and because normative reasons must be *capable* of being possessed and hence of being in that sense internal. Clearly, a reason we possess may or may not be an actual basis of action. When we act for such a reason, it is not only motivational in kind, but *motivating*, and it plays an explanatory role. The paradigm of a reason's motivating action is one's doing something *in order to* realize the desired state of affairs.

The category of possessed reasons overlaps the third main category of reasons pertinent to understanding practical reason, that of *explanatory reasons*: reasons *why* an action occurs, say why one dons a hat. These are typically also possessed, hence reasons one *has*; but explanatory reasons need not be possessed. Something very different, say certain brain manipulations, might explain why

one does something, without constituting or providing a reason one has to do it and without constituting a practical reason even in the widest sense. But typically, explanatory reasons for action are motivating. These are the richest kind of reasons for action. They are called *reasons for which* we do something.¹ They are not only reasons we have; they actually motivate our doing something on the basis of them and thereby ground a motivational explanation of our doing it. They are explanatory, possessed, and commonly also normative. They are also the kinds of reasons for which we act when we act on the basis of practical reasoning.

One thing common to the three kinds of reasons is that, strictly speaking, they are abstract elements – in the case of contents of beliefs and other cognitive attitudes, propositions, and in the case of the contents of desires and of other conative attitudes, states of affairs. Normative reasons are objective in at least this sense: when a normative reason is propositional, the proposition constituting it is true; when it is not propositional, it in some way corresponds to a truth. For instance, suppose there is an objective reason for me to help a friend and that it is expressed, as it might be, using an infinitive: to fulfill my promise. This reason corresponds to the truth that it is a promissory duty to give this help. As to the third case, that of reasons one has, these are expressed by one's intentional states, such as desire, hope, and intention, and they are possessed in virtue of being the contents of the appropriate intentional states. These states may or may not exercise causal power on conduct.

In the theory of practical reason, the main focus of analysis is normative reasons for action. These are *practical reasons*, by which I mean the kind that determines what we have (some) normative reason to do. Correspondingly, they determine what it is *rational* for us to do when, in virtue of one or more of them, we have *adequate* (normative) reason for an action, as where we have a practical reason to do something and no such reason not to. (The notion of its being *irrational not* to do a particular thing

¹ These five kinds of reasons are introduced and discussed in my "Acting for Reasons," *Philosophical Review* XCV (1986), reprinted in my *Action, Intention, and Reason* (Ithaca and London: Cornell University Press, 1993), and, in a different way, in *The Architecture of Reason* (Oxford and New York: Oxford University Press, 2001), esp. ch. 5. I might add that since motivating reasons as here described operate in producing or sustaining action, one might also call them *activating reasons*; and since subjective reasons may or may not activate behavior, but are the appropriate kind to motivate it, we can call them *motivational* as opposed to *motivating*.

yields a concept of a *compelling* reason, but I leave this notion aside here.²⁾

If reasons are *contents* of such propositional attitudes as desires and beliefs, why is it often natural to say, in answer to ‘What was your reason for doing that?’ things like ‘I wanted to show appreciation’, where we cite a desire (though its content coincides with the reason)? For one thing, this reply *both* gives a reason and indicates that it was mine (*I* wanted to show appreciation). For another, ‘I wanted’ contrasts with different attitudes I might have had that express the same practical reason, all of them different from wanting even if compatible with it; for instance, hoping that the action would show appreciation and feeling obligated to show it. When ‘reason’ designates desires (and other attitudes) that express the sorts of abstract elements which constitute reasons of the abstract kinds just described, I propose to speak of *reason states*.

It is often thought that desires can provide all five kinds of reason. Some desires do not. Irrational desires, even if they can motivate, provide no normative reasons. Suppose an agent (*S*) can readily see that the desired object is impossible to realize, say because it is internally inconsistent.³ An irrational desire of this sort surely does not provide any kind of normative reason for action aimed at satisfying it. A normal desire, by contrast, for instance to read a good novel, can provide (which is not to say it can ground) a normative reason, at least *for* certain people. (It can also provide a subjective, explanatory, and motivating reason, but that point is not my concern here.⁴⁾

² What about the possibility of a reason so minor that, by itself, it cannot count toward the rationality of an action? I am inclined to think that if a candidate for such a minor reason deserves the name ‘reason’, then there is a possible action minor enough in some circumstance to enable that reason to count toward its rationality. On a mountain hike, the slightest curiosity about the shape of a leaf can make it rational to examine it.

³ I do not say that a *desire* whose object the person knows or believes to be impossible is irrational, because I think that, for an object we believe is impossible, we can have only a *wish* rather than a *desire* (apart from such lapses as temporary forgetting of the belief). We could consider *any* such conative attitude irrational, but I leave open that one might believe, for instance, that one cannot square the circle yet still rationally wish one could. In any case – and this is the main point – such a wish would not render action to achieve its object rational.

⁴ Beliefs can also provide all five kinds of reason, either for further belief or for action, though they may do so for action only because of what one does or should want, where the ‘should’ is that of rationality. To be sure, beliefs, by contrast with wants, are not quite as naturally described as reasons for action. I think this is because, apart from what we

II. Practical reason

On any plausible view of practical reason, it is a rational capacity, specifically the kind in virtue of which agents respond to (normative) reasons for action. It is also widely agreed to be in some way parallel to theoretical reason, conceived as a rational capacity to respond to (normative) reasons for belief. Theoretical reason is not only analogous to practical reason but also essential to its operation. Beliefs are needed to guide action: desires represent a destination to be reached, but by themselves indicate no routing. An agent will be practically rational, then, at least in good part on the basis of having certain desires and beliefs.

Among the beliefs crucial for practical rationality, instrumental beliefs of a kind that tell us how to satisfy our motivational attitudes are especially important; without such beliefs we could realize our aims only by good luck. Given rationality in these beliefs and in other psychological elements (motivational as well as cognitive), one would expect rationality in practical reasoning and indeed in action. I shall explore rationality in all three cases, but let me first lay out some basic assumptions.

There are several kinds of theory of practical reason. They can be, for instance, subjective or objective, internalist or externalist (about motivation or justification or both), and cognitivist or noncognitivist. The most natural kind, in my view (set out in detail elsewhere⁵) is pluralistic, objectivist, and, in a sense, internalist. It recognizes, for instance, pleasure, pain, and other sorts of rewarding or "punishing" elements as grounds of normative reasons and takes such grounds to be internally accessible to agents in a way that helps to explain their role in motivating action.

A contrasting theory is Humean instrumentalism: it takes non-instrumental desires as (with special exceptions) grounds of reasons for action almost regardless of their content. To be sure, since desires are internally accessible, the theory is internalist. But

want or should, in some presumably objective sense, want, there cannot be reasons for action. Even if, e.g., a belief that listening to an aria will be enjoyable provides, by itself, a reason to listen to it, it also provides a reason to want to listen it, and it could not yield the former reason apart from producing the latter. By contrast, a belief can express a reason for a further belief quite apart from what one wants or even should want, or from any non-cognitive attitudes; and this evidential role of beliefs is usually taken to be their primary reason-giving function.

⁵ A comprehensive theory of rationality is set out in my *Architecture of Reason*.

it is subjectivist: it takes no particular kinds of contents as providing reasons for action independently of the subjective (conative) disposition of the agent.⁶

One might, however, treat rationality as a virtue concept and then try to understand reasons for actions as the kind that would actuate a person having the virtue of being rational in the relevant sense. One might also frame a procedural notion of practical reason, such as a kind of Kantian constructivism, on one side,⁷ or, on another side, a constrained instrumentalism.⁸ The former would take the framework of the categorical imperative as yielding practical commitments for agents who properly employ it; the latter would have us begin with non-instrumental desires and expose them to reflection guided by factual information and broadly logical standards.

All of the plausible theories of practical reason give theoretical reason a role in practical rationality. For all of them, what actions are rational for us depends on our beliefs – especially our rational beliefs – about the consequences of our alternatives. But on my view instrumental beliefs are not the only kind that bear on the rationality of action. Certain normative beliefs also have practical authority. If we rationally believe that something is good (desirable, valuable, worthwhile), we thereby have a reason for action to realize it, even apart from any independent desire. In virtue of these normative and instrumental roles of belief in determining rational action, one might think that practical rationality is reducible to theoretical rationality. I doubt this (and have argued against it and defended a view on which, although belief has a special role as the connective tissue of rationality, practical and theoretical reason are strongly parallel⁹).

My purposes here do not require denying the *possibility* of accounting for practical rationality in terms of a suitably broad

⁶ I have explicated and critically assessed Humean instrumentalism in “Prospects for a Naturalization of Practical Reason: Humean Instrumentalism and the Normative Authority of Desire,” *International Journal of Philosophical Studies* 10 (2002).

⁷ For accounts of constructivism, see John Rawls, *Lectures on the History of Moral Philosophy* (Cambridge: Harvard University Press, 2000) and “Kantian Constructivism in Moral Theory,” *Journal of Philosophy* 77 (1980); Christine Korsgaard, *Sources of Normativity* (Cambridge: Cambridge University Press, 1996); and Onora O’Neill, “Kantian Constructivism,” in Samuel Freeman, ed., *The Cambridge Companion to Kant* (Cambridge: Cambridge University Press, 2003).

⁸ For an instrumentalistic view of the qualified kind I have in mind, see Richard B. Brandt, *A Theory of the Good and the Right* (Oxford: Oxford University Press, 1979).

⁹ In my *Architecture of Reason*, esp. chs 3–5.

theory of rational belief. One might argue, for instance, that, using theological or Platonic or Kantian criteria,¹⁰ we can ascertain what is intrinsically good for us, and that desires and actions are rational on the basis of their role in realizing the good. This intellectualist view allows that our rational beliefs ground our rational desires and rational actions, but it is consistent with the conceptual autonomy of practical reason and does not undermine the distinctions and principles I shall defend. We might, for instance, be able to see that there is *prima facie* reason to avoid pain even if we lack an intellectualist account of *why* pain is a bad thing in human life. Moreover, the principle that there is better (*prima facie*) reason to prefer a more efficient means to an end over a less efficient one can be seen to be sound even apart from the kind of grand theoretical claims about action that might justify it from a point of view on which the rationality of “practical” elements like desires is not basic. One intuitive explanation would be that there is reason to avoid wasting effort; but what I say about practical reason will not depend on this explanation.

III. Practical reasoning

I have spoken of practical reason as a capacity, and I take it as clear that it is commonly and distinctively manifested in practical reasoning. But – although the literature has long contained instances in which the two are not distinguished – practical reason and practical reasoning are quite different things.¹¹ Even if all

¹⁰ A divine command theory would be an example of a kind that would permit deriving reasons for action from convictions meant to be accessible to theoretical reason. For an account of how such a theory might work, see Robert Merrihew Adams, *Finite and Infinite Goods* (Oxford: Oxford University Press, 2000). An indication of how a Platonic reduction of practical to theoretical reason might work is provided in my “Moral Judgment and Reasons for Action,” in my *Moral Knowledge and Ethical Character* (Oxford: Oxford University Press, 1997). For accounts of Kantian constructivism, see the works by Rawls cited above. Cf. Christine Korsgaard’s view that “The . . . truly Kantian strategy is to first give an account of rationality – as we will see, as the autonomy of the human mind – and then to define reasons in terms of rationality, say those which can be autonomously willed, or as those considerations which accord with autonomous willing.” See “The Normativity of Practical Reason,” in Garrett Cullity and Berys Gaut, *Ethics and Practical Reason* (Oxford: Oxford University Press, 1997), p. 243. This does not entail, though it suggests, the primacy of theoretical reason; but it is clearly a sketch of a top-down strategy for understanding practical reason.

¹¹ Even as careful a writer as R. Jay Wallace sometimes slides from points about practical reason to points about practical reasoning without noting the difference (though there is no doubt that he does not consider them equivalent). See “How to Argue about Practical Reason,” *Mind* 99 (1990), e.g. pp. 356 and 380–81.

practical reasoning manifests practical reason, the former does not exhaust the latter. To *A* (where *A* is an action) or even form the intention to *A*, when one knows (but has temporarily forgotten) that *B* is a better way to achieve one's aim is to make a practical mistake, but need involve no reasoning; to judge that one is obligated (overall) to do something and have no accompanying motivation is another kind of (prima facie) failure of practical reason, though it need not involve reasoning;¹² and surely a (non-instrumental) desire for pain – say, one induced by posthypnotic suggestion – is (prima facie) irrational on practical grounds.

The question that naturally arises here is what constitutes practical reasoning and how it manifests practical reason. Such reasoning is widely conceived as reasoning aimed at deciding what to do, by contrast with theoretical reasoning, which is widely conceived as reasoning aimed at determining what is the case. This is broadly correct but must be qualified: first, to account for cases in which reasoning is *exploratory* or has the character of an *exercise*, as where we simply consider a projected plan of action as we might a proposed argument whose premises we do not accept; and second, because 'aim' is metaphorical and chiefly functions to call to mind criteria on which practical reasoning is properly judged.

Capturing what is practical about practical reasoning is difficult.¹³ It is also difficult to capture what makes it *reasoning*. A good starting point is to conceive reasoning as equivalent to inference. In broad terms, an inference is roughly a certain kind of passage of thought from one or more propositions to another, guided by a sense of some relation of support between the former – call it the premise set – and the latter, the conclusion. It is in part because of this guidance element that we may also take inference to be a kind of transition that – by contrast with, say, free association – is appropriately assessed by deductive or inductive standards. Inferences are typically *belief-forming*, but one can infer something one already believes, or draw an inference from something one disbelieves to something else one disbelieves. Rather than explicate these difficult notions here, I take the idea to be

¹² I have defended these views in "Weakness of Will and Rational Action," *Australasian Journal of Philosophy* 68 (1990) and "Moral Judgment and Reasons for Action," in *Moral Knowledge*. In the former, however, I indicate why the failure – as in certain cases of weakness of will – is only prima facie.

¹³ My *Practical Reasoning* (London and New York: Routledge, 1989) contains much discussion of what makes practical reasoning practical. See esp. ch 4.

clear enough to proceed to cases. Two preliminary points, however, are needed.

First, both 'reasoning' and 'inference' have uses in which they designate *processes* and other uses in which they designate the abstract *contents* thereof. If you and I each reason from the same premises to the same conclusion, we make the same inference, and do the same reasoning, in the abstract senses of those terms, the senses in which we can speak of valid inference (and reasoning). But the real mental processes in question are at most qualitatively identical. There are, in my terminology, two inferential *tokenings* of the same inference in the abstract sense. (No sentential tokening is presupposed, though commonly one will occur; the tokens are typed semantically and might be conceived as simply *representations* of an argument in the abstract.) We might also say there are two tokenings of the same argument, bearing in mind that 'argument' also has the same duality of use as 'reasoning'.¹⁴

Second, if practical reasoning is like theoretical reasoning in being broadly propositional, then its concluding element is not an action or intention, though it may be a practical judgment, say that I must now speak up for a friend, which immediately yields action. Anscombe, following her reading of Aristotle, has taken the concluding element to be the action¹⁵; but although this view can be accommodated to ascriptions of validity and cogency, accounting for these favors a propositional interpretation of practical reasoning, and it is preferable to work within the latter view. Much (though not all) of what I say below will in any case be re-expressible in some version of the action-as-conclusion view.

IV. The dimensions of assessment of practical reasoning

It is important to distinguish the assessment of practical reasoning from that of the action (if any) taken on the basis of the reasoning. We can reason quite cogently in favor of an action relative to one goal but lose sight of another goal and hence do something that, though based on a (limitedly) good piece of practical

¹⁴ A further problem is how to characterize the content of a *de se* first-person belief, say that I* need an airline ticket, where I conceive myself in an irreducibly personal way. This paper provides room for various ways to deal with this problem, but cannot propose one.

¹⁵ See G. E. M. Anscombe, *Intention* (2nd edn, Ithaca: Cornell University Press, 1963).

reasoning, is rationally defective. We can also make mistakes in reasoning which we cannot reasonably have been expected to avoid, as where the reasoning is highly complicated; this may result in an action that is rational by plausible standards though based on defective reasoning.

Some patterns of practical reasoning

It will help to note some common patterns – schemata, in my terms – of practical reasoning. It is useful to have a simple kind of basic schema for reference.¹⁶ It has three elements:

1. A purposive (roughly, end-expressive) premise (e.g., I want [need, must achieve, have a duty to realize] *G*, where *G* is a goal);
2. An instrumental premise (e.g., *A*-ing will achieve *G*);
3. A practical conclusion (e.g., I should *A*).

Intention will serve as well as desire in the “major” premise; and the “minor” may indicate not only instrumental means but *constitutive* means: the kind essential in the end itself, as singing is for the pleasure of singing. Let us consider some of the important varieties of practical reasoning.

In one common kind of practical reasoning, the major premise, say that on balance I must accept the invitation, expresses an overriding need, i.e., one taking priority over all competing ends relevant at the time, and the minor premise says that *A*-ing, for example making another trip, is necessary to satisfying the need. This instantiates a *necessary condition schema*. It is plausibly considered valid (though not formally so) because its conclusion simply says that one should, on balance, do something necessary to realizing what one needs on balance – hence something one should on balance attempt to realize. (Rule schemata, such as those representing an action as required by a rule, may also be valid provided the rule figuring in the major premise expresses a similarly overriding demand.)

It is more difficult to identify valid schemata where no such necessary condition is represented. Suppose the major premise sets out a (normatively) overriding end and the minor says that *A*-ing is *sufficient* for it. This would yield a *sufficient condition schema*. It does not follow from these premises that *S* should *A*. An easier

¹⁶ In *Practical Reasoning*, on which I draw substantially in this section, I called this the simplest “basic schema” (ch 4).

alternative might be preferable. Some kind of prima facie judgment does follow, since *S* has *some reason* to *A* (at least where *S*'s belief of the minor is not irrational). I call such judgments practical, but they may often fail to lead to action. Indeed, if there is an obvious alternative means that is far better than *A*-ing as a way to achieve the end, then normally, one would be unreasonable to judge unconditionally that one should *A*.

Suppose, on the other hand, that the minor premise says that *A*-ing is the best way to achieve the end. If 'best' has a suitably broad sense, wider than, say, 'most efficient', then it apparently does follow that *S* should, on balance, *A*. For that is the overall best way to realize the overall best end. The reasoning would exhibit a second kind of practical reasoning, an *optimality pattern*. Depending on whether the end is objectively or subjectively optimal, for instance is "really" best or merely best in *S*'s opinion, the practical judgment will express objective or subjective reason for action.

In the more usual instances of practical reasoning, where the major premise does not represent an end as overriding in the strong sense sketched, even an optimality claim in the minor premise would not suffice for validity. For there might be some competing end in the situation in the light of which, all things considered, *S*'s doing something other than *A* is more reasonable. Granting that we often do posit ends as overriding, we are frequently too cautious to do this and hence can validly infer at best a strong prima facie judgment favoring the action that our minor premise represents as best for achieving our end.

If practical reasoning had only prima facie conclusions, its assessment would be in one way simpler. For these conclusions are often sufficiently weak to follow from the sorts of premises we actually employ, such as that we want to help a student with a paper and, to do so, must work late. However, in the actual context of practical problems, we are trying to determine what to do, and here it is often natural to draw unconditional conclusions yielding a kind of definite directive on which we find it natural to act straightaway.

Often, then, we naturally conclude practical reasoning with an unqualified judgment even if the judgment is not entailed by our premises. Thus, inductive – in the broad sense of 'non-deductive' – standards are more appropriate than deductive standards for appraising the reasoning. In these instances, the basic criterion for good practical reasoning is rather loose. It is the *reasonableness*

of the conclusion relative to the premises, by which I mean that given the premises, the conclusion is quite likely to be true, in a sense implying that it is what is commonly called a “reasonable inference” from them.

The relevant notion of reasonableness is important for the appraisal of practical reasoning. The notion is related to justification as an epistemic concept. If *S* is rational, relevantly informed, and has nothing to go on but the premises, then *S* has at least minimal justification for the conclusion. Reasonableness is normally the appropriate standard for good practical reasoning, and it goes with justification. This point may be taken to imply that the conclusion is probable relative to the premises; but ‘probable’ is misleading in suggesting that we can commonly assign probabilities here. At best, we are likely to be warranted in saying that relative to the premises, the conclusion is more likely than not; but not even this qualified warrant is clearly entailed by the notion of a reasonable inference.

More must be said about reasonableness. It may imply that the premises make it at least as reasonable to believe the conclusion as to believe its negation (and not unreasonable to believe the former). We might call any pattern that meets this standard a *minimal adequacy pattern*. This is a quite permissive standard. Practical reasoning whose underlying argument only meets, and does not exceed, this standard, is not unqualifiedly adequate. For one thing, such patterns allow that it might be *more* reasonable to suspend judgment on the conclusion. Where the premises support the conclusion to the extent that it would be unreasonable *not* to draw it, we might speak of a *standard adequacy pattern*. Here, relative to the premises, it would be a mistake to suspend judgment on the conclusion; and though it might be clear that the premises do not entail the conclusion, they would surely provide adequate reason to draw it.

There is still another gradation. Suppose an argument barely meets the demands of standard adequacy, in the sense that, given its premises, it is only just barely unreasonable not to draw the conclusion. One might say that such an argument is short of being *cogent*. When, on the other hand, the premises give more support than standard adequacy requires, we might speak of a *cogency pattern*. Many such patterns will also be necessity patterns. But there will be disagreement, as with theoretical reasoning, over whether every cogency pattern must be valid. I am inclined to believe that the premises of a cogent argument need not entail,

but only give strong support to, the conclusion. I thus characterize a cogency pattern so as to include inductively strong arguments. For most non-skeptics, at least, there are certain good arguments whose premises do not entail their conclusions; cogent arguments may be of this kind.

Consider some examples of adequacy and cogency patterns. Granting that one would not want to base any important judgment on premises that are only minimally supportive, suppose one is choosing between two almost equally attractive small gifts for a friend. Here, believing that one of them is fairly likely to please, and a bit more likely to do so than the other, minimally warrants the judgment that one should give it. It might be slightly more reasonable to suspend judgment, but one would be reasoning in a minimally adequate way if one judged in favor of the more promising gift. If we vary the case so that one believes that the more promising gift is very likely to please, we would have a standard adequacy pattern, though not necessarily a cogency one. Now consider a cogency pattern. Suppose that my end is (normatively) overriding, say to protect my children; I would then like to have a minor premise that decisively favors one alternative over another, say by indicating a means that is necessary and sufficient for my end. If I find such a premise, my practical argument would be cogent. In a situation of forced choice, however, for instance between paying ransom and sending the police, one might have to act on a slim difference in value, say between a certainty of avoiding financial ruin and a low probability of better protecting the children. A cautious reasoner might then infer a weak *prima facie* conclusion, say that *prima facie* one should call the police. This would preserve validity and would yield a cogent argument for a weak conclusion.

Criteria for assessing practical reasoning

In the light of the kinds of practical reasoning noted, the broadly logical assessment of practical reasoning should address at least five patterns it may have – and many distinct subcases. There are necessity, optimality, and adequacy patterns, and two kinds of cogency patterns (valid and inductively strong). Some generalizations may be drawn immediately. Where the underlying argument is valid, the broadly logical assessment may be fairly straightforward. It may be easy to tell that a practical argument is valid, as with a necessity pattern. But formal criteria alone do not

suffice for the logical assessment of practical reasoning: it may also involve difficult questions about what kinds of ends and means imply various sorts of practical judgments. In the case of an underlying argument appropriately assessed inductively, there is no question of the conclusion's following from the premises; the logical question is how much support the premises give to the conclusion. The answer will rarely if ever be quantitative, and it may be difficult to determine.

The criteria for a broadly logical appraisal of practical reasoning concern the relations between the (propositional) premise and conclusion elements and thus apply to the practical *arguments* expressed in the reasoning. But there are also non-logical evaluative criteria concerning practical reasoning *processes*. Here the problem is roughly how much support *S*'s believing the premises gives to *S*'s believing the conclusion, where the minimal requirement is that the premise beliefs render *S* at least as reasonable in believing the conclusion as *S* would be in believing its negation. We can speak of an *inferential criterion*, since the concern is transmission of support from attitudes toward the premises (typically beliefs of them) to an attitude toward the conclusion (again, typically belief). This is an epistemic matter. Appraising the overall reasoning process requires using inferential as well as logical criteria.

My main point here has already been suggested: it is that however good the argument underlying one's reasoning, the reasoning process is not successful overall if it does not meet an appropriate inferential standard. For instance, if it is merely a rationalization, and one holds the conclusion on some basis other than the premises, then the reasoning fails to produce knowledge, or justified belief, of that conclusion. One could still know or justifiably believe it, but not through the reasoning. In short, one's conclusion, even if validly *inferable* from the premises, is not *inferential* on the basis of them. Thus, whatever support the premises might give it, it derives none from them. A cognition not *based* on premises is not justified *by* them. (I omit discussion of partial basing, in which case the justificatory power of the premises relative to the conclusion is "proportional" to the degree of basing, other things being equal.) This brings us to the issue of epistemic criteria.

The inferential assessment of reasoning, like its purely logical assessment, is indifferent to the actual truth or falsity of its premises, though not to the agent's justification for them. The overall

appraisal of reasoning is not indifferent to truth and falsity; and the rest of this section addresses mainly the epistemic assessment of practical reasoning, including the relations among the truth or falsity of its constituents and *S*'s justification for believing them. Epistemic assessment of reasoning overlaps inferential assessment, since one factor in *S*'s justification for believing the conclusion is the inferential relation this belief bears to the premises. But there are many other aspects of epistemic assessment.

Consider first the premises. There are complicated factors that affect assessment of them. First, these premises may be mistaken without the argument's ceasing to be practical. Moreover, since one may rationally believe certain false propositions, the falsity of a premise does not preclude *S*'s justifiedly drawing the conclusion. If we call a practical argument that is valid and has true premises *sound*, we may say that unsound practical reasoning may nevertheless confer justification on its conclusion, where this implies, minimally, yielding greater warrant for believing it than for withholding it. For *S* may have excellent grounds for believing the premises (and conclusion) even if they are false. Indeed, even an argument that is not valid may instantiate some adequacy pattern and be inductively strong. Moreover, our having sufficient warrant to *take* an argument to be valid may (given justified beliefs of the premises) justify our believing its conclusion. There may be only a very limited range of cases in which such a logical error concerning a practical argument can have the required degree of warrant (or at least of excusability). But there apparently are some cases. If my premises are true and I justifiably believe them, then if I am mistaken in thinking the argument valid only because I miss a very abstruse source of invalidity, perhaps I can still justifiably believe – though I could not know – the conclusion on the basis of these premises.

We have, then, three dimensions of assessment for practical reasoning. The first, the abstract (argumental) dimension, concerns the embodied argument, viewed logically in terms of its validity or inductive strength, and viewed materially in terms of the truth and falsity of its propositional constituents. The second, the inferential dimension – which is governed by both psychological and epistemic criteria – concerns *S*'s justification for inferring the conclusion from, and for believing it in virtue of, the premises. Roughly, the question is how much justification the reasoning process gives to *S*'s belief of the conclusion – typically by transmitting justification from beliefs of the premises to a belief

of the conclusion. The third, the purely epistemic dimension, concerns the overall justification of *S*'s beliefs of each of these propositions (or, if the reasoning is suppositional, the justification *for* believing them, roughly in the sense that if *S* believed them for the reasons constituting this justification, the resulting beliefs would be justified). In a given piece of reasoning, these dimensions may vary independently. For instance, certain reasonings offered in rationalizing an error may be logically and materially adequate, yet inferentially and epistemically defective.

One would hope that one's practical reasoning is adequate in all three dimensions. When it is sound (or at least has true premises and is inductively strong) and, on balance, *S* justifiably believes its premises and, on that basis, also justifiably believes its conclusion, I shall call it *cogent*. This overall notion of cogent reasoning is quite rich, but the intuitive idea is that in this case our premises provide a cogent reason for our conclusion and we hold it *for* that reason. Such reasoning instantiates some cogency pattern, has true premises *S* justifiably believes, and yields *S*'s believing the conclusion on the basis of those premises. The reasoning is logically, materially, epistemically, and inferentially adequate. It satisfies all four kinds criteria and so is adequate in all three dimensions of assessment (three rather than four because truth and falsity are placed in the logical category as crucial for assessing the content of the reasoning). Earlier I indicated how practical reasoning of various sorts may satisfy the logical requirements, and some of the inferential requirements, for cogency. The satisfaction of the (material) truth requirements needs no special comment, but the epistemic requirements for justifiably believing the individual propositions do need it.

Consider the major premise first. We have seen in discussing logical criteria that there is a tradeoff: the stronger our premises (in content), especially in representing our goal as overriding, the better the prospect of validity, particularly if our conclusion is *prima facie*; yet the stronger the premises are, the less likely it is that we justifiably believe them. If my major premise says only that I want to accomplish something, or simply that I have a *prima facie* obligation to do something, then (if I am in normal circumstances) it is not likely that I am unjustified in believing the premise. Nevertheless, clearly we are often justified in believing that we want something, or that we have a *prima facie* obligation to do a certain deed. Often we are also justified in believing that something is currently our overriding end. If I see a child about

to ignite a curtain, I would normally be amply justified in believing that on balance I must intervene.

In many cases, however, I would not have a practical problem if I could easily discern my overriding end, or what, on balance, is my overriding obligation. I may be comparing recreational prospects, so the question may be mainly what I most want (or should want in the light of what I enjoy). A week at the beach would be relaxing, but might be too slow; a week in a favorite city would be too expensive; and there may be several other options. If I decide that all things considered, the country would be best, I might be neglecting other prospects, or wrongly appraising my own reactions to the envisaged situation itself. The mere possibility of a mistake does not undermine my justification for settling on the country, but the point is that an avoidable and unjustified mistake is possible here.

This point can also be illustrated with respect to obligations, as where one must devote time and resources to one child as opposed to another, or to a parent rather than a spouse. It is easy to go astray in such cases and unwarrantedly conclude one's reasoning. And just as we sometimes correct an earlier stance, or retrospectively admit an unwarranted view, regarding what we wanted on balance, we may revise our views on what we are obligated to do.

It is one thing to point out basic kinds of mistakes that can be made in holding the major premise; it is another to give criteria for the degree (if any) of the unjustifiability of holding them. No simple formula suffices, and a case can be made for any of a number of standards ranging from demanding to permissive. Plainly, there is a tendency to insist that, as a reasonable person would, one meet a higher standard where more is at stake. Making a change of career is far more important than choosing a birthday cake. Moreover, if the major premise does not represent the end as overriding, the risk of mistake is reduced. However, if the reasoning is undertaken in the normal way in the course of answering a practical question, then even if the words I use, or would use if I expressed the reasoning, do not indicate an overriding end, the belief I actually express by the words indicates one. I may say simply 'I *believe* I'd better concentrate on the older child's problems today' to express the difficulty of being certain what I should, on balance, do, but my belief may well be to the effect that so doing is my overall obligation. The proposition that it is my overall obligation is thus a good candidate for my major

premise. Similar points apply to the use of expressions like 'I want', 'my aim', and 'I really should', in expressing practical reasoning. Their common modesty of tone may obscure their frequent unconditionality of intended content.

In assessing *S*'s justification for the minor premise, the task is simpler where the premise represents *A*-ing as necessary for the end. In that case the question is *S*'s justification for taking a certain instrumental or constitutive relation to hold. However, we commonly conclude practical reasoning in favor of actions that we do not consider necessary conditions for realizing our end, but regard only as something like our best bet, or good, or adequate, for achieving this end. In the latter case, in which we do not take the means to be necessary for the end, there are at least three criteria. They parallel those cited for the major premise: we may overlook a relevant feature of the action, say its unpleasantness; we may fail to see one or more relevant consequences of it, such as its eliminating the chance of realizing some other end of ours; and we may neglect a consequence for something that, on reflection, we *would* want. The general point is that where *S* does not consider the action necessary for the end, and particularly where *S* believes that it is not necessary, the question of its suitability is comparative. This holds whether or not *S* actually makes a comparison. Thus, from errors of either commission or omission, *S* may unjustifiably believe that the action is, say, a good way to realize the end, hence fail to be justified in believing the minor premise.

V. A range of substantive principles of practical reason and practical reasoning

Many philosophers have proposed standards governing reasons and reasoning. I want to consider some of these that bear on all the dimensions of practical reasoning just considered, but particularly on the inferential and epistemic dimensions. I will formulate the standards mainly as principles that do not essentially refer to reasoning. This is because the normative (e.g., evidential) relation between one set of beliefs (or judgments) and another is not affected by the difference between *reasoning* from the content of the first set to that of the second and, on the other hand, simply *holding* the second set of beliefs (or judgments) on the basis of the first set. Reasoning from one belief to a second,

for instance, yields justification for holding the second if and only if it satisfies the conditions for being justified in holding the second on the basis of the first.

My main point here concerning the practical domain is that if a practical reason justifies believing one should *A*, it does so whether one simply holds this on the basis of the reason – thereby having a *belief for a reason* – or whether one engages in practical reasoning from a proposition expressing the reason to that belief – thereby having a *reasoned belief*, one based on an inferential process (an inferential tokening) of an argument. A belief (or practical judgment) can have the same foundation whether one has climbed to it by the ladder of inference or reached it more directly, simply propelled, as it were, by the force of the evidence.

It may seem that one is better justified in the former case because, having earned one's conviction through reasoning, one is better positioned to justify it. Perhaps we are often better able, or more readily disposed, to justify a belief when we have reasoned to it than when we have simply formed it on the basis of a reason; but this is a contingent matter. We can forget a premise from which we reasoned, for instance, and we can often readily see what belief of ours grounds one whose justification is queried even if we did not arrive at the latter by reasoning.

Hypothetical imperatives

A good place to start in identifying some basic kinds of practical principles is with Kant's famous hypothetical imperative. In one version, it might be called a *principle of the scope of the will in rational persons*: He who wills the end wills the (necessary) means, at least so far as reason has decisive influence on him. This may be plausibly called an *imperative* because it implies that if willing an end is not accompanied by the corresponding instrumental willing, one is in some way deficient in rationality. Kant also says:

HI₁ [W]hoever wills the end wills also (necessarily according to reason) the only means to it which are in his power (see the *Groundwork*, esp. sections 417–418).

Both formulations lack temporal variables. Suppose, however, that (as it appears) Kant means to include cases in which we reflect on what to do or in which we will an end *before* being aware of a means. These cases are common, and we need a standard for them. We should add temporal variables that allow the principle to apply across time. A plausible candidate would be

HI₂ If, at t , S wills an end, E , then, for any necessary means to E which S (a) considers then or over an interval beginning at t and (b) *takes* to be a necessary means in S 's power, S wills that means at t or by the end of that interval.

We may assume that the reference is to means that are not *merely* necessary, as with flipping one of two switches jointly required to turn on a light; here presumably S would will to flip both. The principle is plausible only on the assumption that S takes the means in question to have a significant chance of realizing the end. Kant is apparently also thinking of cases in which we are following through on something like an act of will or an occurrent intention. We might call HI₂ a *Kantian principle of volition transfer* (it seems broadly Kantian even if it is not exactly coincident with one of the principles Kant had in mind).

There are other versions of the hypothetical imperative. In a recent account of it, Christine Korsgaard has said:

[W]illing an end just is *committing* yourself to realizing the end . . . to give oneself a law, hence to govern oneself . . . What about Kant's own formula? If it is to be like my first formula, the one that works [i.e., "if you *have a reason to pursue* an end, then you have a reason to take the means to that end"], then . . . you must think that the fact that you will an end *is a reason* for that end.¹⁷

This passage suggests (though it does not entail) the principle that

HI₃ If, at t , S wills an end, E , and believes that the fact that S wills this is a reason to pursue the end, then, for any necessary means to E which S (a) considers at t or over an interval beginning at t and (b) *takes* to be a necessary means in S 's power, S wills that means at t or by the end of that interval.

It is not implied (by HI₃ or by Korsgaard) that believing there is a reason entails that there *is* one, only that this belief is required for willing an end to generate (rationally) willing a means and to ground the imperatival character of such principles, in virtue of which those who will the end but not the instrumental means are in some way deficient in rationality.

¹⁷ Korsgaard, "The Normativity of Practical Reason," p. 245.

A counterpart Humean formulation represents reasons for action as arising more directly from a motivational state. Consider this “means/ends rule” (M/E): If you desire to *A* and believe that by *B*-ing you will *A*, then you ought to *B*.¹⁸ Viewed as a kind of imperative, this is most plausible if the desire is taken to be predominant. For if one had merely a weak desire to *A*, massively outweighed by desires for objects one knows cannot be realized if *A* is, ‘ought’ would be unwarranted (perhaps even understood as *prima facie*). A principle suggested by this qualification would be

HI₄ If, at *t*, *S* has a predominant desire for a state of affairs, *E*, and believes that *A*-ing will realize *E*, then, at *t*, *S* has reason to *A*.

One might claim that M/E or something similar is a basic principle of practical inference and indeed that “Someone who does not accept the M/E principle cannot be given reasons of any sort.”¹⁹

A principle related to HI₄ which uses intention rather than desire and is closer to HI₁ (on one reading) than is HI₂ or HI₃, is that

[I]f you intend to do something and you do not repudiate this intention, your intention normatively requires you to do what you intend. Unrepudiated intentions normatively require to be acted on.²⁰

¹⁸ James Dreier, “Humean Doubts about the Practical Justification of Morality”, in Cullity and Gaut p. 98 (his variables have been altered to correspond with mine).

¹⁹ Dreier, *op. cit.*, p. 98. Cf. Peter Railton’s comparison (in “On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action,” in Cullity and Gaut) of Moore’s “his true but I don’t believe it” with “*E* is an end of mine, but that’s nothing to me in my deliberation” (p. 68), where I would put ‘good’ or something like it in parallel with ‘true’. More tellingly, he says that

E is an end of mine;

Means *M* would secure *E*;

So: there is that much to be said for my doing *M*, or against my having *E* (p. 77)

is a valid schema *and* defends it by Lewis Carroll’s point. He does not raise the question whether having an end provides any reason at all and so licenses inferences. It would seem to provide some reason – something “to be said for” doing *M*. If it does not, we at best have the prohibition against simultaneously having the end, the belief, and no desire (this would give ends deliberative weight only in a psychological sense).

²⁰ John Broome, “Are Intentions Reasons? And How Should We Cope with Incommensurable Values?”, in Christopher Morris and Arthur Ripstein, eds., *Practical Rationality and Preference: Essays for David Gauthier* (Cambridge: Cambridge University Press, 2001), p. 112.

As explained by Broome, however, the conditional here is material. The requirement is thus that *either* one cease holding an unrepudiated intention or do the intended thing. It is not implied that even unrepudiated intentions generate reasons for action.²¹

Neither M/E nor the unrepudiated intentions principle is temporally qualified. This is important, particularly if reasoning cannot be instantaneous. For then there will be a major difference between normative principles governing practical reasoning and those practical principles simply applicable to agents at a given time, which may govern practical *reasons* but not practical reasoning. I will return to temporal considerations. We should first explore a different dimension of normative assessment.

Three kinds of normative principle

Normative principles differ in a way that is not yet clearly in view. To bring this out, let me contrast two kinds of case, first in the theoretical domain and then in the practical realm. If belief is taken as a counterpart of intention, then a theoretical analogue of HI₁ is

T₁ If, at *t*, *S* believes both that *p* and that *p* entails *q*, then, at *t* (so far as reason has “decisive influence” on *S*), *S* also believes *q*.

Call this *the principle of closure of belief (in rational persons) under believed entailment*. Compare it with a related principle superficially like the unrepudiated intentions principle:

T₂ If, at *t*, *S* believes both that *p* and that *p* entails *q*, then, at *t*, *S* has reason to believe *q*.

This is not a closure principle but a *generation principle*. It says in effect that beliefs generate (normative) reasons via entailment. One plausible cross-temporal counterpart is

T₃ If, at *t*, *S* believes both that *p* and that *p* entails *q*, then if, at or immediately after *t*, and with an awareness of holding these beliefs, *S* considers whether *q*, *S* has prima facie reason to believe *q*.

²¹ Broome specifies (op. cit., 106) that the conditional is material, an important point in distinguishing the principle from the superficially similar one that takes an intention to generate a reason for the action constituting its object. Further pertinent discussion of the kind of principle at issue here is found in his “Normative Requirements,” *Ratio* 12 (1999).

In appraising T_1 – T_3 it is also helpful to distinguish them in relation to time. Call T_1 , T_2 , and other principles applying at a *given* time *synchronic*. Call T_3 , which applies *across* time, *diachronic*. I believe all three are false. Let us start with T_1 .

Consider a moment at which S believes that p and forms the belief that p entails q . T_1 allows that (a) both beliefs are irrational, and (b) S might, on considering q (which we may assume S can do at the same time) justifiedly find q implausible and thereby acquire a reason not to believe it stronger than any reason S has to believe p . Why, then, must S 's belief that p , which is irrational, give S any reason to believe q ? One answer would be that since T_1 holds, S rationally must (is rationally required to) believe q , given S 's believing that p and that p entails q , and that 'must' is the basis of the reason-generating power of the "premise" beliefs in T_2 and T_3 . Is that so?

A principle in this vicinity that clearly *is* true is

T_4 At any given time, there is reason *not* to believe: that p , that p entails q , and that not- q .

But T_4 implies nothing about whether the beliefs that p and that p entails q normatively support believing q . T_4 is what might be called a *coherence principle* (or an *incoherence principle*): it prohibits a kind of incoherence. It says nothing about what one has positive reason to believe. Indeed, p might be obviously false. Moreover, this might be discovered by considering its entailment of q , which one might already take to be false or might readily see to be false upon considering the entailment. That point, in turn, helps to show why T_2 and T_3 are false. Once we allow for S 's *considering* q , the possibility arises that through doing it, S will have reasons for disbelieving q that outweigh whatever reason S has for believing p .

Is practical reason different from theoretical reason on the points that have now emerged? I do not think so. Consider a counterpart of T_2 suggested by some of the practical principles cited above:

P_1 If S intends to A and believes that B -ing is necessary and sufficient for A -ing, then S has reason to B .

Sufficiency is, to be sure, a closer counterpart of entailment than necessity and sufficiency combined; but since, where B -ing is not necessary for A -ing, some other sufficient means to A -ing could

be vastly preferable, we have in P_1 a more plausible principle than we would without the double-barreled condition. In any case, are there not the same kinds of considerations here that apply to T_1 ? It is true that there is a kind of practical incoherence in simultaneously intending to A , believing that B -ing is necessary and sufficient for A -ing (where one takes it that one can B)²², and, aware of the intention and belief, intending not to B (or perhaps even failing to intend to B). But this practical coherence principle implies nothing about what one has reason to intend to do.

Now consider the diachronic counterpart of T_3 :

P_2 If, at t , S intends to A and believes that B -ing is necessary and sufficient for A -ing, then if, at or immediately following t , and with an awareness of having this intention and belief, S considers whether to B , S has reason to B .

P_2 does not rule out the possibility that the intention to A , or the related instrumental belief, is irrational. In this case, the consequent may be false. S also might, on considering whether to B , find the act highly objectionable and thereby acquire an overriding reason *not* to intend to A .

We can now see the importance of the distinction between synchronic and diachronic principles. The former apply to a time slice of the agent; they take no account of change. Incoherence is possible at a given time, and there are sound principles that prohibit it; but considering a proposition or prospect *and* forming a belief on the basis of it or an intention to bring it about is at least normally not possible at a single time. Perhaps it is never possible at a single time if we take it to include *beginning* to consider a proposition or prospect and inferentially forming the relevant belief or intention, as opposed to forming it *while* still considering a belief or prospect. The former – call it *episodic inference* – is probably more common than the latter – call it *emergent inference*. If episodic inference is by its nature barely possible at a single time (a matter that can be left open here), the common kind of inference that qualifies as at least a minimal case of *reflection*, is not.

An important general point that emerges is that consideration may lead to change, and change, in turn, may alter the rational-

²² It would be possible to fail to take it that one can B ; but if one intends to A and believes B -ing necessary for this, yet fails to take oneself to be capable of B -ing, this would in itself imply some *prima facie* deficiency in rationality.

ity status of an intention. This bears on standards governing practical reasoning. It is doubtful that reasoning is commonly instantaneous. This certainly holds if we think of a piece of (episodic) reasoning, as opposed to just the constituent inference conceived as a passage of thought from the premise(s) to the conclusion. But suppose a piece of reasoning can be instantaneous. This would still allow that at the very moment one considers *B*-ing in the light of intending to *A* and of believing that *B*-ing is necessary and sufficient for *A*-ing, one could have the thought that *B*-ing is morally repugnant just as quickly as one could form the intention to *B*. How much the mind can do, or respond to, at a given time, is largely a contingent matter.

More generally, if it is theoretically possible for reasoning to be instantaneous, it is also theoretically possible that, at the relevant time, a thought or realization can occur that provides reason not to form the intention, or do the deed, in question. This bears on how one can avoid the kinds of incoherent triads of propositional attitudes we have been considering. If concluding practical reasoning in the usual way that favors the act figuring in the minor premise can be instantaneous, acquiring a reason to reject one of the premises that favor so concluding can be also. We can instantiate modus tollens just as quickly as modus ponens. The need to avoid inconsistent triads does not by itself favor one pattern of reasoning over the other.

There is another way to put one of my conclusions. Even if reasoning need not be diachronic, it is *dynamic*. It entails a developmental change, at least when it is belief-forming. That change can bring with it new reasons, as where *S* arrives at a new justified practical judgment; these in turn can alter what the agent ought to intend (or believe). There is a sense, then, in which the assessment of reasoning is holistic. This is why so many criteria figure in its proper assessment and why it is defeasible in the light of new considerations.

Toward sound principles of practical reason and practical reasoning
 What principles, then, might we rely on in appraising practical reasoning? There are many (including some for each pattern of practical reasoning considered above), but close analogues of simple deductive closure principles do not seem adequate. What the sound principles are can be seen only in the light of the three interrelated dimensions of assessment for practical reasoning

and the four kinds of evaluative criteria sketched in Section III. The first dimension concerns practical arguments conceived as abstract structures. We have noted several patterns of argument, representing necessity, optimality, adequacy, and cogency. The arguments corresponding to episodes of practical reasoning may be appraised by logical criteria, whether deductive or, in a broad sense, inductive. Secondly, assessment occurs in the inferential dimension: the reasoning process itself may be appraised by inferential criteria. These concern both the conditions for justifiably drawing the inference and the requirements for holding the conclusion in virtue of believing the premises. Thirdly, there is the epistemic dimension: the agent's beliefs of the premises and conclusions may be appraised by epistemic criteria that may or may not concern the reasoning process or corresponding argument. These criteria concern both what (if anything) justifies the agent's beliefs and what might defeat that justification. These epistemic criteria leave open whether any of the beliefs is inferentially justified, through either practical or theoretical reasoning. And fourth, there are material criteria, those concerning the truth or falsity of the propositions in question. Ideally, good practical reasoning expresses a valid underlying argument with premises that are true and justifiably believed and with a conclusion that is true and both justifiably inferred from them and justifiably held on the basis of them.

No one manageably simple principle for assessing practical reasoning takes account of all the variables just noted. But we can discern two that are implicit in what has been said and cover a considerable portion of the common kinds of practical reasonings. The first is suggested by a synchronic theoretical counterpart:

P₃ If, at *t*, (1) there *is* a (normative) reason for *S* to realize *G*, and (2) *S* has a justified belief, or justification for holding a belief, that *A*-ing will realize *G*, then, at *t*, there *is* a reason for *S* to *A*.

This is a principle expressing *closure of practical reasons under instrumentally justified belief* (actual or hypothetical), and it is particularly relevant to appraising practical reasoning having a sufficiency pattern. One counterpart synchronic principle for reasoning would presuppose that inferential and epistemic criteria are met; and a stronger counterpart (of greater interest here) would pre-

suppose the logical and material soundness of the practical reasoning. The latter would yield:

P₄ If, at *t*, *S* does valid practical reasoning, from true premises which *S* justifiedly believes, to a practical conclusion (say, a practical judgment) which, at *t*, is held on the basis of them, then, at *t*, *S* justifiedly holds that conclusion.²³

Where truth and validity are not presupposed – as they often cannot be in appraising practical reasoning – we are forced to be less nearly precise. We might perhaps say that

P₅ If, at *t*, *S* does practical reasoning, from premises *S* justifiedly believes, to a practical conclusion which those premises *adequately* support, then, at *t*, if *S* holds this conclusion on the basis of those premises, *S* justifiedly holds it.

Both of these are synchronic; but they differ in that the second does not require either the truth of the premises or their entailing the conclusion. Diachronic versions, allowing temporal passage, require further qualifications. Since reasoning normally occurs over time, we must take account of what may happen if *S* considers the propositions and prospects in question *after* the time at which *S* forms (or begins moving toward an inference based on) *S*'s beliefs of the premises. A candidate for a diachronic principle for overall practical reasoning might be this:

P₆ If, at *t*, *S* does practical reasoning, from premises *S* justifiedly believes, to a practical conclusion which those premises adequately support, then if, at or after *t*, *S* holds this conclusion on the basis of those premises and does not acquire grounds which defeat that support, *S* (on balance) justifiedly holds it.²⁴

P₆ will be useful only insofar as we understand what kinds of elements defeat the support supplied by the premises. But for a full understanding of practical reasoning, we need a theory that clarifies that in any case. Whether we have such a theory or not,

²³ I take *S*'s justification in P₄ and the principles to follow to be defeasible, but it may still be justification on balance and may be quite strong.

²⁴ Suppose there is a long time between the reasoning and the time in question. How can we tell whether the conclusion is held on the basis of its *premises*? It is not sufficient that the conclusion be held on the basis of (belief of) the relevant propositions; the reasoning process itself must in some way figure in the basis or we will have only a judgment based on the same reasons, but not the same *reasoning*. How to tell is a challenge both to philosophical theory and empirical inquiry, but the task seems possible.

even synchronic principles like P_6 provide critical standards that are of considerable help in appraising practical reasoning.

VI. Practical reasoning and rational action

Implicit in my treatment of practical reason and practical reasoning so far has been the assumption that an action based on good practical reasoning is rational. This section sketches the kind of connection we might expect between a rational action and practical reasoning on which it is based. This is not to imply that rational action *requires* a basis in practical reasoning. Far from it. In broad terms, rational action may be conceived as action that is well-grounded in reasons (including the case of things done for their own sake, say for pleasure).²⁵ This, in turn, is a matter of its being explainable on the basis of rational propositional attitudes, above all desires and beliefs. The contents of these attitudes represent reasons and presumably do not have causal power; the attitudes, as reason *states*, have explanatory power (and, I assume, causal power).

The rationality of desires and other motivational attitudes, intentions in particular, is practical; that of beliefs is theoretical. This last point implies that the notion of rational action cannot be fully explicated without appeal to epistemological considerations. I cannot here discuss the conditions for rational belief, but there is no shortage of informative theories, and understanding rational belief is a central philosophical problem on whose resolution many special theories in philosophy depend to some degree. If, however, we can presuppose the notion of rational belief, we may plausibly conceive rational action as grounded, by rational beliefs, in rational propositional attitudes.

There are, to be sure, questions of degree. Suppose A is well-grounded in this sense, but S should have seen that B would be far better given everything relevant in the context. This does not imply that S 's A -ing is *irrational*, but there is a kind of mistake: choosing B would have been better. Must we, then, maximize some value, to act in a fully rational way? This does not follow. It may seem to follow because, if we are choosing between two otherwise equally acceptable options and it is plain that one

²⁵ I have developed this view in detail in *Architecture* and draw on that book in the next few paragraphs.

conduces more to (say) human flourishing, as where a charity is more efficient than its competitor with the same concerns, then we should choose the better one. That we need not always be maximizing does not permit us to ignore opportunities to advance the good that present themselves in the course of our everyday activity or our discharge of ordinary duties.

It would be a serious mistake to infer, from the defeating role of the perceptible inferiority of an option in conducing to flourishing (or any other value), that we are obligated always to maximize some value or even that we must always positively aim at maximizing some value.²⁶ Commitment to a *preferential standard* in making concrete choices does not entail commitment to adopt a maximizing standard as either a criterion of rightness or a general policy of deliberation. It is one thing to avoid choosing a lesser alternative when we consider options; it is quite another to take maximization as governing our choices and deliberations at the outset.

It will be apparent that on many points my view contrasts with Humean instrumentalism, which conceives non-instrumental desires as grounding reasons for action and hence endorses practical generation principles like P_1 and M/E. This is not the place to argue for the superiority of an objectivist view. Moreover, Humeans can grant many of my points: that there is a plurality of basic reasons (even if they consider them all desire-based); that practical reasoning is sound only when the agent rationally believes its premises and is justified in taking them to justify the conclusion; that action in accord with the conclusion of practical reasoning is rational in virtue of that reasoning only if it is based on the reasoning; and that in general action is rational only when grounded, by rational belief, in certain propositional attitudes. The chief difference is that (apart from such defects as internal inconsistency and the obvious empirical impossibility of realizing the desire) for Humeans any non-instrumental desire provides a basis of rational action.

The contrast between the kind of objectivist theory I have sketched and a subjectivistic instrumentalism bears directly on the

²⁶ This is reminiscent of the error in epistemology of inferring, from the capacity of incoherence to *defeat* justification, that coherence is the *ground* of justification. Rejecting the fallacious inferences here does not, of course, commit one to denying that we should even try to maximize or that coherence never plays a positive role in justification. Analysis of the epistemological case is provided in my *Structure of Justification* (Cambridge: Cambridge University Press, 1993), esp. chs 3–4.

treatment of hypothetical imperatives. On my view, there are no hypothetical imperatives – as opposed to conditional coherence principles – unless there are categorical ones: nothing is required *given* a condition unless something is categorically required. Similarly, there is nothing we ought to do given what we want, unless there is something we ought to want in the first place. Even conditional coherence principles owe their normative status to the categorical requirement that we avoid incoherence. Kant might well have thought this, particularly if he conceived of what one *rationally* wills as rational in the light of an application of the categorical imperative framework.²⁷

It might seem that even Humeans may grant the dependence of hypothetical imperatives on categorical ones, since for them the mere presence of a suitable non-instrumental desire “categorically” *constitutes* a ground for a reason for action. But this would be a mistake: Humeans do not conceive reason as categorically calling for our having any desires at all, much less any particular ones. There is nothing intrinsically worth wanting; we are not even rationally required to want desire satisfaction itself. We may seem to be, because it is a truism that given a desire, we want its satisfaction *in the sense that* we want realization of its object. But nothing in the Humean theory requires that we have the relevant *higher-order* desire for the satisfaction of one or more of our desires.

I do not claim to have demonstrated here that hypothetical imperatives depend for any normative power they have on categorical ones or, more accurately, that desires alone do not ground reasons for action and, by themselves, constitute only psychological rather than normative bases for action. But I have shown that there is a subtle difference, which has apparently not been generally noted, between, on the one hand, practical and theoretical coherence principles and, on the other, normative generation principles, such as hypothetical imperatives as usually understood. Coherence principles are not imperatives, except insofar

²⁷ Some of the central passages in the *Groundwork* do not make it clear that this is how Kant sees it, and one might wonder whether here he is influenced by Hume – or at least by instrumentalist ideas – more than is usually recognized. Korsgaard, as quoted above and elsewhere in the same article, seems to take Kant to be presupposing that what we will is in some sense rational; but this is not entirely clear, since autonomous willing is taken to be a source of reasons, rather than reflecting them. This gives autonomous willing a role much like that of desires on an instrumentalist view, though to be sure one that is more like Brandt’s constrained instrumentalism than like Hume’s instrumentalism in the *Treatise*.

as we may *presuppose* reason to avoid incoherence and take this to be in some sense imperative.

Once the difference between coherence principles and genuinely normative generation principles is noted, hypothetical imperatives like HI₄ and similar principles are less likely to seem normatively sound. It appears that neither desires nor even intentions have independent normative authority. The counterparts of such hypothetical imperatives in the theoretical domain are not plausible; and in that domain, experience – especially perceptual, memorial, and reflective experience – plays a major role in grounding rational belief.

There is a similar range of reasons to think that experiences – particularly rewarding experiences of enjoyment and aversive experiences of pain and suffering – play a major role in normatively grounding rational desires and rational intentions. If this is so, then practical reason is substantive. It supplies criteria for rationality that go beyond those calling for avoidance of incoherence; and these criteria yield a variety of constraints on practical reasoning and, more generally, on rational action. Rational action is action well-grounded in reasons; it may or may not also be based on practical reasoning; but when it is, the standards governing its rationality must take account of the same kinds of grounds that are central in the general theory of practical reason.²⁸

*Department of Philosophy
University of Notre Dame
Notre Dame, Illinois, Indiana 46556
raudi@nd.edu*

²⁸ Earlier versions of this paper were presented at Texas A and M University and at the University of Reading. Comments from the audiences in both universities were of much help in revising. I particularly want to thank John Cottingham, Jonathan Dancy, Brad Hooker, Michael Meyer, Christian Miller, Elizabeth Radcliffe, and Philip Stratton-Lake.