# Practical Reasoning

Antti Kauppinen (a.kauppinen@gmail.com)

Properly final, re-revised version, March 13, 2016

Here is a train of thought that someone might have:

**Paris**

I'm going to go to Paris in two weeks.

Flying with Ryanair is the cheapest way of getting to Paris.

But Ryanair are greedy bastards who will rip you off if you close your eyes for a second.

Flying with Aer Lingus is the cheapest and most convenient of the remaining options.

To fly with Aer Lingus in two weeks, I must book now.

So, I'll book now. (Or action: straightaway, I book the flight)

And here is another:

**Advance**

The advance for the book has come in.

It would be great to have a new guitar.

On the other hand, the family needs a new car.

I can't get both a car and a guitar.

Family comes first.

So, I'll buy a new car.

I believe both of these trains of thought could be manifestations of good (if possibly enthymematic) *practical reasoning*, the exercise of our capacity of reason pertaining to action. Not everyone would agree. Virtually everything about practical reasoning – among other things, its scope, premises, conclusions, and norms – is subject to widespread disagreement. In the face of such diversity, it seems to me that the best way to proceed is to stick close to everyday examples like Paris and Advance, and try to come to understand what if anything makes them instances of (good) practical reasoning. As the cases suggest, practical reasoning comes in many varieties. Paris begins with a proposition that is presumably the content of an *intention* to go to Paris, and concludes with an intention or action in service of that original intention. It thus is or contains an instance of *instrumental reasoning*. Advance, in contrast, begins with propositions that are contents of *evaluations*, broadly speaking, and concludes with a decision. I will stipulatively label non-instrumental practical reasoning *deliberation*. Much of what I say will concern instrumental reasoning, as it turns out to be difficult enough to make sense of its standards and psychology.

My aim is not to survey the vast field as a whole. I will have to leave aside important questions, such as the role of decision theory in practical rationality. Instead, I will focus on two competing conceptions of practical reasoning that have emerged in recent philosophical literature. According to what I'll call the Rule-Guidance Conception, roughly, practical reasoning aims at acquiring (or retaining or giving up) intentions so as to meet the synchronic requirements of rationality we ought to meet. John Broome, in particular, has recently developed an impressive account of reasoning as a rule-guided operation on contents of psychological states. According to the Reasons-Responsiveness Conception, in contrast, practical reasoning aims to get us to respond correctly to the reasons there are. For some in this tradition, rationality is an evaluative rather than normative notion, and reasoning is guided by our take on reasons rather than rules of rationality. Some Reasons-

Responsiveness theorists believe standards of correctness for rational response derive, roughly, from what one would have objective reason to do, were things as one supposes.

I'll begin by examining the Rule-Guidance Conception, in particular the role, nature, and justification of the requirements of practical rationality that are central to it. As we'll see, under pressure of criticism, emphasis has shifted from synchronic and symmetrical requirements to diachronic basing permissions that govern the process of reasoning when it goes well. This has brought the view closer to the Reasons-Responsiveness Conception. But important differences remain. Reasons-Responsiveness theorists do not form a unified group, but I propose that a good way to make a number of arguments cohere is to say that requirements of rationality all derive from one's subjective take on reasons, while the standard of good reasoning is given by how one ought to respond to things on the assumption that belief-inputs to reasoning are true and there is sufficient reason for non-doxastic inputs such as intentions. On this view, we need to distinguish between three kinds of normative failure: irrationality, bad reasoning, and objectively mistaken conclusions. If this is right, merely distinguishing between rationality and reasons does not suffice to capture all the normative relations in play.

In the final section, I argue that while Broome's work offers important insight into the nature of reasoning as a conscious activity, the Reasons-Responsiveness Conception has an advantage when it comes to capturing the non-monotonic nature of practical reasoning, among other things. I also sketch an account of the process of instrumental reasoning, according to which there is no premise corresponding to the intention for the end. Instead, intention works in the background, motivating the reasoning and guaranteeing that the agent takes the fact that the means-action facilitates the end as a reason to take the means.

## 1 Good Practical Reasoning: The Rule-Guidance Conception

To understand when a train of thought amounts to reasoning, we have to understand what makes reasoning *correct* or good, since even bad reasoning aspires to be good reasoning. The Rule-Guidance and Reasons-Responsiveness Conceptions offer radically different takes on what makes practical reasoning correct, and thus on the nature of practical reasoning in general. I will start with the former. Here are some theses that have been endorsed by its proponents:

> *Requirement Conformity*
> The point of reasoning is to come to conform to requirements of rationality.
>
> *Wide-Scopism*
> Requirements of rationality govern combinations of psychological states.
>
> *Inference Basis*
> Requirements of rationality derive from inferential relations among contents.
>
> *Objective Normativity*
> Rational requirements are objectively normative: one ought to (or has reason to) conform to them.
>
> *Rule Guidance*
> Reasoning is a process of revising thoughts that is guided by the rules of rationality.

These theses are logically independent, so one might endorse Wide-Scopism without endorsing Inference Basis, for example. (In fact, the leading proponent of the view, John Broome, now rejects Inference Basis.) But they form a coherent picture.

To see the appeal of this view, it is useful to consider the parallels between practical and theoretical reasoning. The two are closer than many used to think. To begin with, in much of older literature on the topic, it is claimed that the premises of practical reasoning are of the form "I want to…" or "I intend to…" and "I believe…". For example, Georg Henrik von Wright says that in practical inference to necessary means "the premises are a person's *want* and his *state of knowing or believing* a certain condition to be necessary for the fulfillment of that want." (von Wright 1963, 168). His example of a first-person inference is the following:

I want to make the hut habitable.
Unless I heat the hut, it will not become habitable.
Therefore I must heat the hut.

Engaging in the inference is meant to reveal a 'practical necessity' for the agent, and consequently motivate action.

Alas, there are a number of problems with this pattern of inference. Most obviously, we can rationally want incompatible things, and surely are under no rational demand to take the means to each of them. (Quite the opposite.) Later writers in this Psychologistic tradition, such as Robert Audi, take the major premise to be of the form "I intend…" (Audi 2006). Intentions are a much more plausible candidate than desires for giving rise to rational demands, since their function is to enable intrapersonal and interpersonal coordination (Bratman 1987). But it is nevertheless problematic to psychologize the premises by including the agent's intentions and beliefs in the *content* of the reasoning.

To see this, consider the parallel with theoretical reasoning. Suppose you know Shoshanna's sense of humour, and are confident that she'll laugh out loud if she watches *Veep*. She's just about to watch it. So you reason that Shoshanna will laugh. How should we describe your reasoning? On the Psychologistic model, your thoughts would be the following:

I believe Shoshanna will watch *Veep*.
(Optional: I believe that) if Shoshanna watches *Veep*, she'll laugh.
So, Shoshanna will laugh.

In terms of content, this is clearly an invalid argument (the premises could be true without the conclusion being true). Nor does the truth of the premises raise the probability of the conclusion without further assumptions. Consequently, it is plausibly a case of bad reasoning. Something has gone wrong. And it is evident what: this reconstruction fails to distinguish between your psychological state and its content. It is the latter that you reason with. The picture looks like this:

| Psychological States | Propositions |
|---|---|
| Belief that Shoshanna will watch *Veep* | Shoshanna will watch *Veep* |
| Belief that if Shoshanna watches *Veep*, she will laugh | If Shoshanna watches *Veep*, she will laugh |
| Belief that Shoshanna will laugh | Shoshanna will laugh |

Call the psychological states the *premise-attitudes* and the *conclusion-attitude*. In theoretical reasoning, they are beliefs with premises as their content. It is natural to think that what makes the transition from premise-beliefs to the conclusion-belief an instance of good deductive reasoning is that the contents of the premise-beliefs *entail* the content of the conclusion-belief. In inductive reasoning, the contents of premise-beliefs make the content of the conclusion-belief more *probable*. What theoretical rationality requires thus seems to derive from (likely) truth-preserving transitions from a set of believed propositions to another (Inference Basis). Further, what rationality requires seems to be that we *either* form the conclusion-belief *or* give up one of the premise-beliefs: it is the combination of having one but not the other that is guaranteed to be irrational (Wide-Scopism). Rule-Guidance theorists typically believe this rational requirement is normative in the sense that everyone objectively ought or at least has reason to conform to it (Objective Normativity). When we reason in the Shoshanna case, we lose or acquire beliefs in response to some *modus ponens* based rule (Rule Guidance) so as to come to conform to the wide-scope requirement (Requirement Conformity).

On the Rule-Guidance Conception, *practical* reasoning consists in a psychological transition from premise-attitudes to the conclusion-attitude, where the latter is an intention. Take the last steps of the Paris case. They might be reconstructed as follows, assuming your name is Daniel:

| Psychological States | Propositions |
|---|---|
| Intention to fly in two weeks | Daniel will fly in two weeks |
| Belief that if I don't buy a ticket now, I will not fly in two weeks | If Daniel doesn't buy a ticket now, Daniel will not fly in two weeks |
| Intention to buy a ticket now | Daniel will buy a ticket now |

There are two obvious but possibly superficial differences from the theoretical case. First, propositions feature in practical reasoning in a first-personal mode of presentation (as on the left). This is plausibly a matter of character rather than content in Kaplan's (1989) terms, so I will represent the contents in non-indexical terms (as on the right). Second, linguistic expressions of intention typically do not have propositional form – we do not ordinarily say "I intend that I will fly in two weeks". But I'm going to grant here that intentions can be said to have propositional content.[1] Broome (2004) takes it to be the same as the content of a belief about one's future action. If *p* is the proposition that one ϕ-s, to intend to ϕ is (*inter alia*) to have a making-true attitude towards *p*, while to believe that *p* is to have a taking-true attitude towards *p*.

In the instrumental case, then, it seems the *premise-attitudes* of practical reasoning are an intention and a belief about a possible action of one's own that facilitates its realization, and the premises are the contents of those attitudes. What about the *conclusion*?[2] One popular view is that the conclusion is the proposition (in the first-personal mode of presentation) that is the content of the derived intention. The instrumental intention can then be called the *conclusion-attitude* of the reasoning. The basic argument for thinking this way is that forming an intention is as far as reasoning can get us (Broome 2000, 87). If reasoning tells me to press this key right now and I form the intention to do so, but I don't press the key because of sudden paralysis, no amount of further reasoning will help. There is no mistake of reasoning involved. As Broome puts it, "[w]hat rationality requires of you is proper order in your mind" (2013, 152).

---

[1] However, see Dancy (2009) for several reasons to doubt this.

[2] A conclusion in the sense we're asking about isn't just anything that concludes the process of reasoning – a heart attack might do that. As Dancy (2014) emphasizes, we're asking about the *proper* conclusion of practical reasoning, the thing that stands in the relevant normative relation to the premises.

There is, however, a venerable tradition going back to Aristotle that takes the conclusion of practical reasoning to be an *action*. After all, it is natural to think that practical reasoning concerns what to do, not what to intend (except in special cases). But even if we leave aside the issue about reach of reasoning, this view faces the challenge that actions do not, on the face of it, have propositional content. They are not propositional attitudes but events. As such, they cannot follow from the contents of any premise-attitudes.[3] If Inference Basis or something analogous holds, this suffices to rule action out as conclusion.[4] The third popular view, according to which practical reasoning concludes with a *normative judgment*, is compatible with what is said so far. The basic challenge to it is simple: why wouldn't the process of forming an evaluative or normative belief be an instance of theoretical reasoning? For cognitivists about normative judgment, such beliefs concern putative facts, for which we have more or less evidence. We engage in standard belief reasoning and conclude that we ought to or have conclusive reason to do such and such, but there is nothing distinctive about the requirements of rationality governing this process. Once we have such a judgment, however, many think we can reason from it to intention or action. In such practical reasoning, however, the judgment functions as a *premise*-attitude rather than as the conclusion. And there is little reason to think that practical reasoning *necessarily* involves such premises.[5]

Let us assume, for now, that practical reasoning proceeds from the contents of premise-attitudes and concludes with an intention. The Rule-Guidance Conception says it is a conscious activity that aims at conforming to requirements of practical rationality

---

[3] This argument goes back to Hume, who argued that since reason is concerned with truth and falsehood and actions cannot be either true or false, they cannot be "either contrary to or conformable to reason" (*Treatise* 3.1.1, 458). Maria Alvarez (2010) defends the view that actions do have content, which is the same as the content of the intention with which one acts. For problems with such proposals, see Dancy 2009.
[4] For other arguments against the Aristotelian view and responses to them, see Tenenbaum 2007.
[5] Dancy (2014) also observes that it is often the very same things that favour thinking that one ought to do something that favour doing it. He is surely right that favouring action is more basic – the facts favour *believing* that you ought to do something because they favour your *doing* it.

(Requirement Conformity). I will discuss the Rule Guidance conception of the nature of this activity in Section 3, and will briefly examine Objective Normativity in the next section. But in order to understand and assess these claims, we'll first have to have an idea of what the requirements of practical rationality might be and why we should think they are such. Here are versions of two commonly suggested schemas:

> *Wide Means-Ends Concordance (WME)*
> Rationality requires of you that [*if* you intend to E at t *and* believe that you will not E at t unless you M at t, *and* that you will not M at t unless you intend to M at t, you intend to M at t].
>
> *The Wide Enkratic Principle (WEP)*
> Rationality requires of you that [if you believe that you ought to $\phi$, and that you will $\phi$ if you intend to $\phi$ and will not $\phi$ unless you intend to $\phi$, you intend to $\phi$].

Let us begin with the content of these requirements. For them to have *wide scope* means that rationality requires one not to have a certain *combination* of states. Suppose you intend to get your uncle's money and believe that you will not get it unless you kill him, and believe (most likely implicitly) that you won't kill your uncle unless you intend to kill your uncle, yet lack the intention to kill your uncle. You are in violation of WME. Since (by Objective Normativity) rationality is normative, you have reason to *either* intend to kill your uncle *or* give up your intention to get his money *or* give up one of your beliefs. Note in particular that it doesn't follow from your having the end-intention and the means-beliefs that you have reason to (intend to) kill your uncle. That would only follow if the rational requirement had a *narrow scope* – that is, if you would be required to have the means-intention whenever you have the end-intention and the means-beliefs. The core argument for Wide-Scopism is precisely that narrow-scope requirements would result in illegitimate *bootstrapping*: we could create reasons or oughts to do silly or evil things simply by having goals for which we

took them to be necessary.[6] Similarly, a narrow scope version of the Enkratic Principle together with Objective Normativity entails that whenever I believe I ought to do something, I am rationally required and thus ought to or have reason to do it. But I am not infallible about what I ought to do. Wide-Scopism (seemingly) avoids bootstrapping and infallibility.

Why would rationality require us to have or avoid certain combinations of intentions and beliefs? There are many ways to argue for this. Inference Basis says that the requirements derive from inferential relations among the contents of intentions and beliefs. Here, I will only briefly discuss a tempting thought along these lines. The thought is this: Suppose you believe that if you won't kill your uncle, you won't get his money. Then it seems that simply in virtue of inferential relations among propositions, you can't rationally believe that you'll get your uncle's money without believing that you'll kill him. As you see it, getting his money entails killing him. Consequently, it seems you can't rationally set out to *make true* that you'll get your uncle's money without setting out to make true that you'll kill him. So the requirement to [intend the necessary means if you intend the end] seems to derive from inferential relations among propositions. (I'm leaving out a lot of potentially quite problematic assumptions here.[7]) Alas, even if there were no problems with auxiliary assumptions in this argument, there is a gap between logical validity and what appear to be correct instances of practical reasoning. Most obviously, it is common and apparently correct to reason to best means, even if the best means isn't necessary. Maybe I think I could get my uncle's money by slaving for him until he dies, but just think killing him would be a better way to do it. However, if killing my uncle is just *a* way (even if it is the best one) of getting my uncle's money, then even if I will get his money, it doesn't follow that I will kill him. So

[6] Broome offers another reason why bootstrapping is problematic by drawing a parallel with the theoretical case. Suppose it was a narrow scope requirement of theoretical rationality that if you believe that *p* and that *p*-> *q*, you ought to believe that *q*. Now, every proposition trivially follows from itself (*p* -> *p*). So if you believe that *p*, then by the narrow scope requirement, you *ought* to believe that *p*! In other words, whatever you believe, you ought or have reason to believe. Such bootstrapping is evidently objectionable.

[7] For developments of this type of argument, see Harman (1976), Wallace (2001), Broome (2004), Audi (2006), and Setiya (2007). For problems, see Brunero (2009) and (2014) and Broome (2013).

without making a logical mistake, I could set out to make true that I'll get his money without setting out to make true that I kill him. Thus, inferential relations among contents can't explain why reasoning to best or sufficient means isn't always bad reasoning. Going beyond instrumental reasoning presents even more daunting challenges. So Inferential Basis seems false. (Broome, for example, now rejects it.) Insofar as the correctness of *theoretical* reasoning drives from inferential relations among contents, there will thus be a major difference between theoretical and practical reasoning.

How about Wide-Scopism? It doesn't fall with Inference Basis, since wide-scope requirements could be defended on other grounds, such as appeal to intuitions about rationality. But it has been directly challenged on the basis of the *synchronic* and *symmetrical* nature of wide-scope requirements. Simplifying a little, you can come to satisfy WME either by intending the means or giving up the end, and WEP either by forming the intention or giving up the belief. Suppose you find yourself not intending to do what you believe you ought to. Which way should you go? WEP gives you no guidance. This is problematic in two different ways. First, Niko Kolodny argues that principles that don't guide us by telling us what to do to comply with them cannot be *normative* (2005, 517; 2007). At best, they provide a standard for *evaluating* the state we're in, like the 'requirements' of beauty. Normative principles of rationality have to be *diachronic*: to tell us what to do going forward from where we are now.

The second, related issue with Wide-Scopism is that it seems that in many cases, the different ways of coming to satisfy a wide-scope requirement are not in fact equally rational (Schroeder 2004). Suppose you believe you have conclusive reason to stop eating chocolate but don't intend to stop eating chocolate. There is a conflict among your attitudes that you can rationally resolve by forming the intention to stop eating chocolate. But you cannot rationally resolve it by giving up your belief that you have conclusive reason to stop eating

chocolate on the basis of lacking the intention. After all, a different principle of rationality tells you to base your beliefs on your evidence, and that hasn't changed. Here a Wide-Scoper might say, to be sure, that changing the belief is only *epistemically* irrational, but as far as *practical* rationality goes, either way of coming to meet WEP is on a par. The rejoinder is that if you change your ought belief to match your intention, your problem is not merely epistemic (Schroeder 2004, 346).

Broome has responded to these criticisms by introducing his own version of diachronic and asymmetric requirements, or rather permissions. On his current view, a move in reasoning is correct if it is allowed by a *rational basing permission* interacting with a requirement of rationality. The form of such permissions is the following:

> Rationality permits N that
> N has attitude A at some time and
> N has attitude B at some time and
> N has attitude C at some time and
> …
> N has attitude K at some time and
> N's attitude K is based on N's attitudes A, B, and C. (Broome 2013, 190)

For an attitude to be based on another attitude is, roughly, for it to be caused by the other attitude in a way that is sensitive to the content of the attitude. Leaving aside questions about the precise nature of this relation, appealing to it allows Broome to say that basing your belief that you ought not $\phi$ on your not intending to $\phi$ is rationally *prohibited*, while, roughly, basing your intention on your ought-belief is rationally *permitted*. While he admits there is no general method for deriving basing permissions from synchronic requirements, he believes the two are closely related (2013, 248). After all, on the Rule-Guidance Conception, the point of reasoning, which is directly governed by basing permissions and prohibitions, is to come to meet synchronic requirements. I will return to how this might happen in Section 3, having examined the alternative view of correct reasoning in the next section.

## 2 Good Practical Reasoning: The Reasons-Responsiveness Conception

Reasons-responsiveness theorists of rationality and reasoning endorse the contemporary consensus view that *reasons* are facts or states of affairs or possibly true propositions that count in favour of action and attitudes like belief, desire or intention. This conception of reasons is compatible with many different explanations of what it is for R to favour ϕ-ing and what makes it the case that R favours ϕ-ing. Most agree that it is possible for R to be a reason for some S to ϕ even if S doesn't know R or see that R favours ϕ-ing – for example, I may have a reason to give to charity even if I don't think so. It is plausible that what we (objectively) ought to do is what we have conclusive reason to do. If we recognize and respond to all our reasons correctly, we will thus act (and believe) as we should.

Yet almost everyone agrees that there are dimensions of normativity beyond doing the right thing. We can be rational or engage in good reasoning while nevertheless ending up doing what we shouldn't or not doing what we should. Suppose there is an elaborate trail of planted evidence that leads me to blamelessly but falsely believe that my neighbour has killed his wife (see *Gone Girl*), and will escape without punishment unless I shoot him in the hand. Because of this, I shoot an innocent man in the hand. I ought not do so, in the sense that there is conclusive objective reason for me not to harm an innocent person.[8] Yet it could be rational for me to think otherwise, and I might be able to form the intention to shoot by correct reasoning from justified assumptions. Indeed, as Conor McHugh and Jonathan Way (forthcoming) emphasize, the starting points of good reasoning can be even worse: if you irrationally believe the your plane will crash, you can reason well to a decision not to board it.

[8] The issue is complicated, however, since at least in some cases, it appears that what I ought to do hangs on my epistemic situation. Some argue that this means we should distinguish between subjective and objective oughts, while others hold that either evidence-relative or evidence-independent oughts are merely apparent. For discussion, see the Sepielli chapter and Kiesewetter 2013.

How, then, could we hope to make sense of rationality or good reasoning in terms of responsiveness to reasons? I will here focus on two kinds of answer, which I'll stipulatively label Subjectivism and Subjunctivism.[9] Roughly, Subjectivism says that we are rational when we have the attitudes we take ourselves to have sufficient reason (of the right kind) to have, and reason well when transitions from one set of attitudes to another are governed by our take on reasons. The paradigm of irrationality is thinking that I have most reason to believe or intend X and failing to do so. In the shooting case, my evidence might support the belief I ought to shoot, and it might make it rational for me to form the intention to shoot. Subjunctivism, in turn, says that what rationality requires is having attitudes one has conclusive reason to have if one's beliefs are true (or: having attitudes one would have conclusive reason to have, were one's beliefs true[10]). In the shooting case, perhaps it is true that were there no other way to prevent my neighbour from escaping punishment for murder, I ought to shoot him in the hand because of this. In that case, it seems that I've responded to the misleading evidence the way I should when I form the intention to shoot based on it. If I've reasoned from the supposed facts to acting the way I did, I've reasoned well.

Reasons-responsiveness theorists, among whom I include Maria Alvarez, Jonathan Dancy, Benjamin Kiesewetter, Niko Kolodny, Errol Lord, Joseph Raz, Thomas Scanlon, and Mark Schroeder, form a fairly heterogeneous group. Here are some theses that can be gleaned from their work, though none would endorse all of them:

> *Reason Conformity*
> The point of reasoning is to come to conform to objective reasons in one's thought and action as well as one can.
>
> *Subjectivism*

[9] For other views I lack the space to discuss here, see Kiesewetter 2013, Lord 2014, and Sylvan 2015.
[10] Strictly speaking, the counterfactual construal invites the conditional fallacy – the closest possibility in which I have true beliefs might be one in which I ought not do what I actually ought to do. That's why it's better to say, as Schroeder (2009) does, that one has what I'm calling a subjunctive reason if one has beliefs such that if they are true, one has objective reason. I'll nevertheless use subjunctive formulations for convenience.

Subjective reasons are considerations one takes to be or treats as objective reasons. Rationality requires believing or intending what one takes oneself to have conclusive reason for.

*Subjunctivism*
Subjunctive or would-be reasons to F are propositions one believes which, if true, are objective reasons to F. Rationality requires believing or intending what one has conclusive subjunctive reason for.

*Error Theory About Structural Rationality*
There are no structural requirements of rationality; the appearance is merely the consequence of how subjective or subjunctive reasons for attitudes predictably stack up.

*Subjective Normativity*
Doing or believing as rationality requires is normative only from the agent's first-personal perspective.

*Value Thesis*
It is not necessarily the case that we *ought* to do what we have conclusive subjective or subjunctive reason to do, but acting on subjective or subjunctive reasons is nevertheless a *good* disposition or virtue.

In the rest of this section, I will examine the competing theses of Subjectivism and Subjunctivism in more detail, discussing the other theses as they come up. The next section will consider implications for the nature of practical reasoning.

Subjectivism, as I have defined it, claims rationality is to be understood in terms of subjective reasons, which represent our take on objective reasons. Here are its core claims, a little more precisely stated (based on Kolodny 2005, 524 and Kolodny 2008, 457):

*Core+*
If one believes at *t* that one has conclusive reason to have attitude A, then one is rationally required to form or sustain, going forward from *t*, on the basis of this belief, A.

*Core-*
If one believes at *t* that one has conclusive reason not to have A, then one is rationally required to revise or refrain from forming, on the basis of this belief, going forward from *t*, A.

(These principles should be read with the proviso that things don't change in relevant respects in the course of reasoning.) The rationale for these principles is simply that, as

Thomas Scanlon puts it, "it seems clearly irrational to fail to have an attitude one explicitly judges oneself to have conclusive reasons for" (2007, 90) Note that I have formulated the principles so that not only do rational requirements have narrow scope, but they are also diachronic: they tell you to form an attitude on the basis of your beliefs. Such requirements do not set a state as a target for the process of reasoning, but rather directly govern the process of reasoning. They thus play a role similar to Broome's basing permissions. If all rational requirements have this form, it follows that it is not necessarily irrational to *have* inconsistent intentions or beliefs, for example. If you are hit in the head and acquire the belief that the Earth is flat without giving up your existing belief it's not, you do not thereby become irrational. Instead, as Kolodny puts it, "Being rational just is responding in the ways that process-requirements call for." (2005, 517) Rational requirements govern the way we conduct our thinking in the light of our evidence and commitments (assuming that having evidence or a commitment entails judging oneself to have reason to believe or intend).

The Subjectivist idea is that other rational requirements can be derived from the Core ones. For example, it is a short step from Core+ to the following:

> *Narrow Enkratic Principle (NEP)*
> If you believe that you have conclusive reason to $\phi$, and that you will $\phi$ if you intend to $\phi$, and will not $\phi$ unless you intend to $\phi$, rationality requires of you that you intend to $\phi$.

The most obvious objection is that this allows bootstrapping – by NEP, if I falsely believe I have conclusive reason to poison the water supply, rationality requires me to intend to do so on the basis of the belief.[11] But surely I have no reason to do so! Avoiding this was the key motivation for Wide-Scopism. Subjectivism avoids bootstrapping reasons or oughts into existence in another way: by denying that rationality provides reasons, and endorsing Subjective Normativity instead. On this view, rationality is only *seemingly* normative, since

[11] More precisely: intending is rational, as long as I have the belief. If my evidence doesn't support the belief, theoretical rationality requires changing it.

whenever a rational requirement applies to someone, it will seem to *her* that she ought to do what it tells her to do. (Kolodny (2005) labels this the Transparency Account.) For example, NEP only applies to you if you believe you have conclusive reason to poison the water supply and that poisoning it depends on your intending to do so, and it tells you to intend to poison. In any such situation, it will seem to you that you ought to intend to poison the water supply. From your perspective, it is as if the rationally required thing is the thing to do.

Note that the fact that rationality requires you to intend something doesn't itself provide either a subjective or an objective reason for anything. The subjective reason to intend is whatever you take to be the conclusive reason to poison the water supply, combined with the action's being up to you. So rationality itself isn't normative in any way. It is, after all, surprisingly difficult to say *what* is the reason to be rational such that we ought to satisfy rational requirements in every case, even when doing so results in our doing or believing something we should not, as it certainly might.[12] This is why subjectivist Reasons-Responsiveness theorists reject Objective Normativity. But rationality may still be an evaluative notion, as the Value Thesis says. There's something to be said for you as an agent in virtue of satisfying rational requirements. It follows, importantly, that rational process requirements themselves cannot *guide* reasoning, since they are not normative. Subjective reasons can, since they appear to you to be normative, and when they do guide your reasoning, you are rational.

There are two other main challenges to Subjectivism. The first is *attitude scarcity*. Broome notes that we rarely have explicit thoughts about what we have sufficient reason (of

[12] See Kolodny 2005, Raz 2005, Broome 2007, and Bratman 2009 for discussion. Jonathan Way (2010) provides an ingenious answer to the problem in the case of instrumental rationality. Briefly, he argues that the fact that M is necessary for E is a reason to either give up the intention to E or to intend to M. Assuming that rationality requires us to do what we would have reason to do, were things as we suppose, if you believe that M is necessary for E, rationality requires you either to give up the intention to E or to intend to M. This is a Subjunctivist move, and Way's view does, in my view, ultimately belong to the Reasons-Responsiveness camp. If he's right, there are wide-scope requirements that derive from subjunctive reasons and have jurisdiction over only some individuals, which makes them very different from the kind of wide-scope requirements central to the Rule-Guidance Conception.

the right kind[13]) to intend or believe, so even if such thoughts suffice to rationalize intending, appeal to subjective reasons won't account for why rationality always requires intending as we believe we ought, for example (2013, 97–98). The obvious move in response to this is to appeal to implicit reasons thoughts, which might not be equally scarce. I will briefly return to this below when discussing reasoning.

The second objection is *getting the extension of rational requirements wrong*: it seems we are sometimes rationally required to believe or intend in certain ways even if we *don't* think we have reason to do so, and are sometimes not rationally required to believe or intend in certain ways even if we *do* think we have reason to do so. Take means-end coherence. I intend to get my uncle's money and believe I have to kill him to get it – no beliefs about reasons for attitudes are evident, so why would rationality require intending to kill, according to Subjectivism? Scanlon provides material for an interesting response based on a thesis about the nature of intention: "[A] person who believes that doing *A* would advance some end of hers will … see the fact that *A* would advance this end as a reason for doing *A* simply because she has the end in question" (2007, 86).[14] So: one of the things that makes my attitude concerning getting my uncle's money an *intention* is that I do take the fact (as I believe) that killing my uncle would facilitate getting his money as a reason to kill him, and the fact (as I believe) that intending to kill my uncle is a condition of my killing him as a reason to intend to kill him. So whenever I have the end-intention and the means-beliefs, I will have (conclusive) subjective reason to intend to take the believed means. That is part of what it is to be *committed* to the end.[15] It doesn't yet follow that I am irrational *all*

---

[13] For reasons of space, I leave aside the issue of reasons for intending that are not reasons for acting. For my take, see Kauppinen (2014).

[14] I have elided, among other things, Scanlon's qualification "insofar as she is rational". I don't think it can be made sense of in this context, if all rational requirements are to derive from Core+ and Core-. (See next footnote.)

[15] On this point, Scanlon seems to be treating commitment as a distinct normative primitive. On his view, it's irrational to fail to treat M's facilitating E as a reason to M, *because* that would "involve acting contrary to a commitment that one has made (and not revised)" (2007, 93). On my variant, again, being committed to E *is*,

*things considered* if I don't so intend, since I might take myself to have more reason not to kill. In that case, rationality would require giving up the end-intention instead. But that is just how it should be – it isn't all things considered irrational to fail to take the means to one's end, if one takes the cost of the means to exceed the benefit of the end. (In that case, it is all things considered irrational not to give up the end.)

This quick Subjectivist account of means-end coherence will certainly not quell all the doubts. Part of the case will be providing an error theory about Wide-Scopism. Such a theory claims that whenever there seems to be a wide-scope requirement prohibiting a combination of attitudes, there is in fact subjective (or subjunctive) reason to drop one of the conflicting attitudes. (See the Way chapter for discussion.) Even if that is accepted, however, Subjectivism has some counterintuitive consequences. As Parfit (2011, 120–121) famously notes, according to Subjectivism, if you believe you have no reason to avoid agony on Tuesdays (or don't believe you have such a reason), it is not rationally required of you that you choose a painless operation on Wednesday rather than an excruciatingly painful one on Tuesday. Yet surely someone who doesn't choose the painless operation is irrational, or so the thought goes. Rationality appears to have some sort of substantive aspect that transcends our opinions concerning reasons.

My suggestion is that at least some of this pressure goes away if we distinguish between *rationality* and *reasoning well* in addition to distinguishing between rationality and what we objectively ought to do. The rationale for this distinction is that there are (at least) three different kinds of normative failure in this vicinity. First, I will be *objectively mistaken* if I fail to act or think the way I have conclusive objective reason to act or think, regardless of my beliefs and commitments. Second, there is something wrong with me if I don't act or

---

among other things, treating M's facilitating E as a reason to M, so that what is irrational is failing to M (unless the subjective reason is outweighed). On my version, the only irrationality in acting contrary to the commitment is acting against one's take on reasons. It is thus compatible with Core+ and Core- being the source of all rational requirements.

believe as I take myself to have conclusive reason to act or believe, regardless of what reasons there objectively are for me. This is what I propose to call *irrationality* in the narrow sense.[16] But there is a further way in which I can go wrong that is not captured by either of these categories. Suppose again that I believe an operation on Tuesday would cause agony, while a Wednesday one would be only inconvenient. It is very plausible that if this belief is true, I have conclusive reason not to have the operation on Tuesday. But it's possible that I somehow fail to recognize the badness of pain on Tuesday, and instead reason from *An operation on Tuesday would be painful but convenient* to *I will have the operation on Tuesday*. Further, if it happens to be the case that my belief is false, and that the Tuesday operation is both painless and convenient, I'm not objectively mistaken in forming the intention either – the reasons there are for me to favor choosing the Tuesday operation. Nevertheless, *had* things been as I took them to be, I *would* have made a terrible mistake (call this the Tuesday Case). So it seems clear that there is something wrong with my thinking on this matter. My suggestion is that it is natural to describe this as an instance of *bad reasoning*. My misguided normative belief led me to reason to the wrong conclusion from a true non-normative belief. Good reasoning requires some kind of sensitivity to the reasons there are if things are as we suppose them to be, not just forming attitudes conforming to our take on reasons. After all, when we reason badly, it's typically precisely because we mistakenly treat something as a reason when it's not, or treat something as a weightier reason than it actually is.

The last kind of failure is the starting point of what I'm calling Subjunctivist theories. However, they are standardly formulated as theories of rationality (or subjective ought), and thus presented as competition to Subjectivism. Here is how Derek Parfit formulates the view:

---

[16] Although in work in progress, I develop Kurt Sylvan's (2015) idea that rationality's demands are determined by a *competent* take on reasons, where competence is fallible sensitivity to objective reasons.

> We are rational insofar as we respond well to reasons or apparent reasons. We have some *apparent* reason when we have beliefs about the relevant facts whose truth would give us that reason. (Parfit 2011, 5)

I think it is misleading to talk about *apparent* reasons in this way. Suppose I believe there is gin and tonic in the glass. If these beliefs were true, I might have reason to drink from the glass. This is the case even if the glass actually contains petrol (call this the Gin Case). But I don't necessarily *take myself to have reason* to drink gin and tonic – the supposed fact need not *appear* as a reason to me even if it is one, relative to my beliefs. So it is misleading to say I have an apparent reason to drink in the circumstances. For the same reason, Schroeder's term *subjective* reason (2009, 233) is also somewhat misleading. So I prefer my own *subjunctive* reason, or Kurt Sylvan's (2015) *would-be* reason. I also disagree with Parfit about *rationality* being connected to subjunctive reasons, as long as the beliefs in question concern non-normative facts. There may be nothing irrational about believing that the glass contains gin and tonic and not drinking from it, even if that fact is in context a conclusive reason to drink, as long as one does not *take* the fact that there is gin and tonic to provide a reason. In addition, we need *some* label for the failure involved in forming attitudes or acting against our own take on reasons. It might be worth reserving the term "irrationality" for this.

However, even if Subjunctivism is problematic as an account of rationality, it might offer a standard for *good reasoning*. Reasoning involves drawing a conclusion on the basis of (the contents of) premise-attitudes. The contents of at least some of the premise-attitudes may be such that were they to obtain, they would sufficiently favour the conclusion. (As Dancy (2014) observes, if we think in these terms, the conclusion favoured by the reasons may after all be an action rather than just intention.) Here is a suggestion in the Subjunctivist spirit:

*Tracking Sufficient Would-Be Reasons*

> Reasoning from the contents of premise-attitudes to conclusion C is good if and only if, and because, if non-normative facts are as the agent assumes, the contents of belief-premises are sufficient reasons for C.[17]

This accounts straightforwardly for deductive and (good) inductive reasoning, since if the premises of such reasoning are true, they necessitate the truth of the conclusion-belief or make it probable, and thus are sufficient reasons for it, assuming that the right kind of reasons for believing have to do with truth. It also certifies the agent's reasoning in the Gin Case, since had the agent's belief been true, she would (we're assuming) have had sufficient reason to drink. Also, reasoning in Tuesday Case turns out to be bad by this standard, since had the non-normative facts been as the agent took them to be, she wouldn't have had sufficient reason to choose the painful operation. How about defeasible reasoning? McHugh and Way (forthcoming) point out that in the absence of defeaters, it can be good reasoning to infer from *Somebody said that p* to *p*. But surely, one might protest, the mere fact that someone said that *p* isn't sufficient reason to believe that *p*! True – it is easily defeasible. That's why the principle says "if non-normative facts are as the agent assumes", not just if the agent's premise-beliefs are true. If the agent additionally believes that the speaker is a liar, then if non-normative facts are as the agent assumes, the fact that the speaker said that *p* is not a sufficient reason to believe that *p*, and the reasoning is not good. But if the agent has no such additional beliefs about defeaters, their reasoning is okay. This is arguably the case even if the agent should, in the light of their evidence, realize that the speaker is lying. When we assess the quality of a piece of reasoning, we focus on local rather than global features of an agent's thinking. Even if you've reasoned badly to *p*, or adopted *p* on non-rational

[17] Here is how Dancy puts it in the context of discussing someone who begins from false premises: "What they deliberate from is not the case, but they may still be correctly tracking what that situation would have favoured had it been the case, and deciding to act accordingly. In such a case, they will have reasoned well, but since they started in the wrong place they may well end up doing the wrong thing, or a wrong thing anyway." (Dancy 2014, 6).

grounds, you can still reason well from *p* to *q*, if the truth of *p* would give sufficient reason to believe that *p* in the context of your other assumptions.

Yet things are not so rosy for Tracking Sufficient Would-Be Reasons when it comes to practical reasoning. Suppose you intend to get your uncle's money. In the best kind of case for this view, you also have (false) non-normative beliefs about your family affairs such that were they true, you ought to get your uncle's money. Next assume that you believe that killing your uncle is necessary and sufficient for getting his money. If both these beliefs are true, you plausibly have sufficient reason to kill your uncle. Why? Following Raz's (2005) insightful discussion, many have endorsed the following kind of principle:

> *Transmission*
> If you have a reason to E and doing M would non-superfluously facilitate E, you have a reason to M, of a strength proportionate to the strength of the reason to E and the probability of E-ing given M.

(For more precise formulation and argument, see Kolodny's chapter.) Roughly, transmission explains how your reason to get your uncle's money generates a reason to kill your uncle, if your beliefs are true. Hence, by Tracking Sufficient Would-Be Reasons, reasoning from your end-beliefs and means-beliefs to intending to kill your uncle is good reasoning – it mirrors the structure of objective reasons if things are the way you think (call this the Uncle Case).

But hold on! The content of the intention wasn't mentioned in the story, and means-ends coherence surely has to do with means to an end one has! It is indeed hard to find a place for intention in Tracking Sufficient Would-Be Reasons, since intending doesn't seem to amount to supposing that things are in a certain way. The basic move is to say that intending *ensures* the presence of certain beliefs. For example, Mark Schroeder suggests the following hypothesis about the nature of intention: "If you intend to do A, then you have some beliefs which are such that, if they are true, then you objectively ought to do A." (2009, 236; see the consequent discussion for modifications). But this is a highly dubious

assumption. We can, after all, have intentions for which we don't even *think* there is sufficient reason, as when we akratically intend to take another piece of chocolate.

What if the restriction to non-normative facts were removed from Tracking Sufficient Would-Be Reasons? The thesis would then be that what makes reasoning good is that were both the agent's normative and non-normative beliefs true, she would have sufficient reason for the conclusion. This would help to explain the goodness of instrumental reasoning if intending involves believing that one has sufficient reason for the end, or at least the means. (These assumptions would require a lot of defence, of course.) It would also explain why it's good reasoning to move from *I ought to jump off the train* to *I shall jump off the train* even if one doesn't actually have sufficient reason to jump off the train (Train Jump Case). Alas, it would also certify the bad reasoning in the Tuesday Case – after all, were the agent's false normative belief that pain on future Tuesdays doesn't matter true, she would have sufficient reason to choose the Tuesday operation. While it makes sense to bracket the truth of the agent's non-normative beliefs in assessing the quality of the transition from premises to conclusion, it doesn't make sense to bracket the truth of the normative beliefs that guide the transition itself. Reasoning well involves precisely taking some fact to be a reason of a certain strength when it actually is a reason of that strength, or at least something close enough.

Note, however, that not all normative beliefs that are involved in reasoning govern transition from (what one takes to be) non-normative facts to a conclusion-state. Belief that one ought to do something, for example, doesn't play a guiding role in this way. Instead, it functions relevantly like non-normative beliefs – as *input* to reasoning rather than as a part of our perspective on what favours what. Thus, for a non-vacuous account, it suffices that we don't bracket the truth of those normative beliefs that guide the transition from premises to conclusion. This results in a view according to what makes reasoning good is that there is

sufficient reason for the conclusion, if the agent's non-normative beliefs and normative beliefs that don't guide the transition from non-normative beliefs to the conclusion are true. For short, we can say that in good reasoning there is sufficient reason for the conclusion, if the *doxastic inputs* of reasoning are true. Insofar as reasoning involves beliefs or belief-like attitudes like *Pain on future Tuesdays doesn't matter* or *That a car goes fast is a decisive reason for buying it*, they are not inputs to reasoning, but attitudes that shape our response to how we take things to be.

With this addition, Tracking Sufficient Would-Be Reasons explains why enkratic reasoning in the Train Jump Case is good – were the doxastic input *I ought to jump off the train* true, the agent would have sufficient reason to intend to jump off the train. But unless controversial views about beliefs involved in intention are true, practical reasoning remains a mystery. To avoid these assumptions about the nature of intention, a reasons-responsiveness theorist might try a different tack. Maybe what makes reasoning good is that it's *sufficient-reason preserving*, in something like the way good deductive reasoning is truth-preserving. On this variant, the standard of good reasoning is that if one has sufficient reason for the premises, one will have sufficient reason for the conclusion, other things being equal. However, as McHugh and Way (forthcoming) point out, this won't work, because one may have sufficient reason to believe something false. That means that when one falsely but with sufficient reason believes that there is gin in the glass, one won't have sufficient reason to intend to drink from it, because the glass, after all, contains petrol. The move from the belief (and prior intention to drink gin) to the intention to drink from the glass is thus not sufficient-reason preserving. Nevertheless, it is intuitively good reasoning.

So what should a reasons-responsiveness view say about the role of intention in practical reasoning? Here's something that seems true about the Uncle Case, for example: if you did have sufficient reason to intend the end, and if your non-normative beliefs were true,

you would have sufficient reason to take the means. More precisely, by Transmission, the fact that killing your uncle is necessary and sufficient to get his money would be a sufficient reason to kill him. Similarly in the Gin Case, if you did have sufficient reason to drink gin, and were your non-normative beliefs true, you would have sufficient reason to drink from the glass, derived from your reason to drink gin. What this suggests is that beliefs and intentions play very different roles in making reasoning good. For intentions, the model of preserving sufficient reasons works, because there is no analogue of a further dimension of assessment like truth in the case of belief.

So perhaps the best form of a reasons-responsiveness account will be disjunctive, reflecting the different roles of doxastic and non-doxastic inputs:

*Tracking Sufficient Would-Be Reasons**

Reasoning from the contents of premise-states to conclusion C is good if and only if, and because, if the doxastic inputs are true *and* the agent has sufficient reason for non-doxastic inputs, the contents of belief-premises are sufficient reasons for C.

On this picture, then, instrumental reasoning to believed necessary means, for example, is good roughly because were one to have a sufficient reason for the end-intention, it would guarantee that the contents of the means-beliefs (means M is necessary for end E, and up to me) would together sufficiently favour taking (and intending) the means. Conversely, reasoning is *bad* if there wouldn't be a sufficient reason to have the conclusion-attitude even if the agent's non-normative beliefs and non-guiding normative beliefs were true, and even if the agent had sufficient reason for her intentions. So someone who reasons from *I will get rich* and *I will only get rich if I buy a lottery ticket now* to *I will buy a lottery ticket tomorrow* reasons badly, because even if the inputs are correct, the content of the belief won't provide a reason for the conclusion.

**3 The Process of Reasoning**

Suppose I intend to steal my neighbour's lawnmower, and believe that scaling the fence is necessary for me to do it, and that I won't scale the fence without intending to do so (for short, that it's up to me). Everyone agrees on two points: I may come to have the intention to scale the fence through some effortless, automatic process, or by way of reasoning. But what exactly do I do when I engage in reasoning? Rational requirements (or permissions) or perceptions of reasons evidently play some kind of role. But when we try to spell it out, we soon encounter a dilemma. The first horn I'll call *mere conformity*. It could happen that I have, in succession, the following mental states:

> Intention to steal my neighbour's lawnmower
> Belief that scaling the fence is necessary to steal my neighbour's lawnmower and up to me
> Intention to scale the fence

My psychological states, and their evolution over time, conform to diachronic rationality. But this isn't yet a description of a process of reasoning: I have a permissible conclusion-attitude, but for all that's been said so far, I didn't reason to it. I could have acquired the second intention as a result of a hit on the head. Further yet, bad reasoning is also possible. So mere conformity to a rational norm is neither sufficient nor necessary for reasoning. Nor will it suffice if the second intention is *caused* by the end-intention and the means-belief. It wouldn't be a case of reasoning if the intention and belief caused me to intend to scale the fence by virtue of some bizarre association that might as well have resulted in an intention to eat a bowl of pasta. For the process to be one of reasoning, the premise-attitudes have to cause the conclusion-attitude *in virtue of* my somehow taking them to rationally require conclusion-attitude or provide a reason for the conclusion. To avoid this horn, there is a temptation to conceive of the process as involving something like the following mental states:

Intention to steal my neighbour's lawnmower
Belief that scaling the fence is necessary to steal my neighbour's lawnmower and up to me
Belief that rationality requires (or permits) me to (intend to) scale the fence, given my intention and belief
(Optional: Intention to conform to requirements of rationality)
Intention to scale the fence

The additional belief here is what Broome (2013) labels a normative higher-order belief, since it is a belief about normative relations among attitudes. He rightly raises two key objections to it. First, we can reason without such beliefs, which require a level of sophistication few of us have. Second, there's good reason to think we are not even able to form beliefs, in particular, in response to such higher-order beliefs about rational requirements or permissions, even if combined with the optional intention to believe or intend in the required way. Beliefs aren't susceptible to that type of voluntary control. Nor will adding a first-order belief about what follows from the *contents* of our attitudes do the job. That way lies Carroll's famous regress – a rule that we use to derive a conclusion from premises cannot itself be among the premises, or we would need yet another rule to derive the conclusion from the new set of premises, and so on to infinity.

What the dilemma shows is that in some way, the psychological transitions involved in reasoning must be guided by rational requirements or subjective reasons without their featuring in the contents of premises or as higher-order beliefs. Broome offers perhaps the most detailed account aiming to meet this constraint. For him, again, practical reasoning is an activity of bringing ourselves to conform to the requirements of rationality. To characterize this activity, it is insufficient to say that insofar as we are rational, we are *disposed* to move from premise-attitudes to the conclusion-attitude when rationality so requires (Smith 1994). This does not yet distinguish reasoning from automatic, non-actional conformity with rational requirements. Such a disposition might be regularly activated in a creature incapable of practical reasoning. The same goes for us: we often come to believe the

obvious consequences of our beliefs or (perhaps) intentions without engaging in reasoning, as a result of subpersonal processes. Psychological research suggests that this is true particularly of some simple forms of inference, such as modus ponens – even modus tollens requires more attention, and results in more errors (Evans 2006).

So there is a further question about how conscious reasoning works. Following Broome, let us begin with a simple case of theoretical reasoning. You are at an aquarium, where there are fish and sea creatures of all kinds. To your surprise, you see some dolphins repeatedly surfacing to breathe. So you form the belief (A) dolphins breathe air. It follows from your background zoology, based on vague memories of what you learned in school, that (B) if dolphins breathe, they are mammals. So you engage in a bit of quick reasoning that concludes with the belief that (C) dolphins are mammals (perhaps you had forgotten about this). What is that you do to get to C from A and B? Plausibly, one element in this is that you *attend* to the propositions, perhaps by engaging in a kind of inner speech, as Broome suggests. You say to yourself, whether silently or aloud, "Dolphins breathe" and "If dolphins breathe, they are mammals", and then, as a causal consequence, "So, dolphins are mammals". But what role does the relevant rule of inference play in this move? According to Broome, what makes the causal process reasoning is that the conclusion-belief is caused by consciously attending to the contents of the premise-attitudes and performing *a rule-guided mental operation* on them. It is performing this operation that licenses saying to yourself "So, …" on reaching the conclusion.

But what is it to perform an operation on contents? For Broome, operating on contents is simply an activity of deriving a new content from them by way of applying a rule to them. When you are guided by a rule, he says, the derivation *seems right* to you, though at the same time you regard it as open to correction if checking results in its no longer seeming right (2013, 237–239). You're following a rule correctly if the derivation has a steady

disposition of seeming right, Broome suggests. These remarks point to the deep waters of the rule-following debate. I will dip into them below, but this point, I will only register my disagreement with Broome's claim that inferring in a way that conforms to *modus ponens* as a result of its robustly seeming right to do so entails *believing* a proposition of the form "If (p, and if p then q), then q" (229). Why should we think so? Many marks of belief that *p* have been proposed, such as willingness to avow *p* in suitable conditions and reliance on *p* as a premise in reasoning. Neither is true of Broomean linking propositions. So it is better not to assume that belief in them is necessary.

One problem that Broome himself observes with this picture is that insofar as the contents of thoughts are propositions, there is nothing to distinguish between reasoning from beliefs to beliefs and from intentions to intentions. Suppose you say to yourself "I will go to Paris in two weeks". This could express either a belief about your future action or an intention to act. But the outcome of the mental operation of combining the content with the content of the belief that buying a ticket now is necessary for going to Paris is a different state depending on which kind of state the sentence expresses. To overcome this challenge, Broome introduces the idea that there is a *marker* in the content of the attitude that indicates whether it is the content of a belief or intention, for example (2013, 252). The content of a belief is a proposition together with a belief marker, while the content of an intention is a proposition together with an intention marker. So the content of a belief might be [I will go to Paris in two weeks; belief] and of the corresponding intention [I will go to Paris in two weeks; intention], where the markers are in fact inaudible. If we accept this, then it makes sense that contents with an intention marker are subject to different rule-governed operations than contents with a belief marker, and can yield different outputs.

Broome's picture of reasoning is very liberal in one respect: for him, *any* rule-governed mental operation on contents counts as reasoning. *Correct* reasoning, however,

operates with rational basing permissions. Take instrumental reasoning. According to Broome, the following is an instance of rationally permissible basing:

> *Instrumental Permission*
> N intends at some time to E, and
> N believes at some time that M is a means implied by E (one won't E unless one does M)
> N believes at some time that doing M is up to her herself, and
> N intends at some time to M
> N's intention to M is based on N's intention to E, and belief that M is a means implied by E, and belief that M is up to herself. (Cf. Broome 2013, 257)

Basing one's intention to M this way is one way to come into conformity with WME. Such basing need not involve reasoning, as noted above. On Broome's take, when it does involve reasoning, the operation is governed by the following rule (using square brackets to indicate psychological states):

> From:
> [E; intention] and
> [M is a means implied by E; belief]
> [M is up to me; belief]
> to derive
> [M; intention] (Cf. Broome 2013, 259)

Now we have a full picture of Broome's view. There are synchronic requirements of rationality that we ought to conform to, and corresponding (though not strictly speaking derived) diachronic basing permissions that tell us how we may come to do so. From basing permissions we derive rules that define which operations on contents amount to correct practical reasoning.

Broome's work on reasoning is pioneering and in many ways enlightening. Yet there are a number of challenges to it that suffice to motivate developing a corresponding account based on the Reasons-Responsiveness Conception. I will here leave aside potential issues with introducing marked contents, which threaten to collapse the attitude/content distinction highlighted in Section 1. The first problem I'll discuss has to do with the form of the rational basing permissions. On Broome's conception, again, the rule you follow when reasoning

says that from the *entire set of premises* you may derive the conclusion – say, A, B, and C permit deriving D. Although this is appealing in the case of deductive reasoning, it is at least unlikely to suffice for practical reasoning (or indeed other types of theoretical reasoning). This is because practical reasoning, at least, is *non-monotonic*, as is widely acknowledged (Brandom 2000, Dancy 2004). What this means is that adding premises can turn a correct piece of reasoning into a bad one. It might be that A, B, C, and X do *not* permit deriving D, while A, B, C, X, and Y again do. Take the instrumental reasoning rule above. Once again, I intend to get my uncle's money, believe that killing him is a means implied to get his money, and that killing him is up to me. Following Broome's rule, I can reason well to the intention to kill him. But now add the belief that killing someone for money is never permissible. Now it seems I can only reason well to giving up the intention to get my uncle's money. But Broome's rule is either silent on what to do or still permits me to intend to kill (since I still have all the attitudes it features). To make sense of non-monotonic reasoning, we need to show how *each individual premise* contributes to establishing the conclusion. And there is no way basing permissions could capture this, since they apply to the entire set. We could, to be sure, imagine a potentially infinite list of basing permissions and corresponding rules governing every possible combination of contents, but then we wouldn't be talking about a possible human activity any more.

The second problem concerns the picture of reasoning as a rule-governed operation on contents. It seems to me that what really does the work of distinguishing reasoning from mere causation of the conclusion-attitude by the premise-attitudes is the idea that we form the conclusion-attitude because it *seems right* to do so on the basis of the considerations we have consciously rehearsed, perhaps with the aim of finding out what they support. As Paul Boghossian puts it, using 'inferring' for the activity of belief reasoning, for S to infer from p to q "is for S to judge q because S takes the (presumed) truth of p to provide support for q."

(Boghossian 2014) Broome himself says very little about what seeming right amounts to. It is not a phenomenal state, but just an "attitude of yours towards the mental process you go through when you reason", which amounts to "a sort of personal endorsement from you" (2013, 238). I agree that an element of endorsement is necessary to distinguish reasoning from non-reasoning. But there is a gap between endorsing a transition from one content to another and performing a rule-guided mental operation, or so it seems. It is not clear to me why the latter would be necessary. It may of course be, and most likely is the case in something like *modus ponens* reasoning, that you are disposed to endorse any transition with the same formal features. But I see no reason to believe that endorsing a transition in a particular case requires this sort of generalizing commitment. If so, what distinguishes reasoning from non-reasoning may just be that you take the contents of the premises, taken together, to sufficiently favour the conclusion-belief or intention (if not action). And that brings us to the Reasons-Responsiveness Conception.

So how might reasoning work on the Reasons-Responsiveness Conception? According to it, the point is not to come to conform to synchronic requirements of rationality, but to come to act and think in ways one has most objective reason to act and think, starting from one's assumptions about how things are. In this process we work with subjective reasons, since from a first-person perspective they are indistinguishable from objective reasons. Bearing in mind that we can and typically do respond to reasons without engaging in the activity of reasoning (as Kate Manne emphasizes in her chapter), the process of conscious reasoning will begin with some *triggering event*: you're puzzled or unsure about how to go on in thinking or acting. Consider the everyday cases I started this chapter with. In Advance, the triggering event is the realization that there's money available that you haven't budgeted for, so you need to decide what to do with it. In Paris, I believe, it is becoming aware that you have a goal that you don't yet know how to reach – though this is

likely to be a controversial take on the situation. The awareness could coincide with forming the goal.

The triggering event is followed by a *search for reasons*. You rely on your automatic processes to bring to mind considerations that might bear on the problem. One way in which they can do so is by favouring a course of action. But as Dancy (2004) has rightly emphasized, favourers are not the only considerations bearing on what to do. There are also modifiers (*attenuators*, which weaken favourers, and *intensifiers*, which strengthen them) and conditions (*enablers*, which 'turn on' other facts as favourers, and *disablers*, which turn them off). You call these things to mind, among other things, by imagining what would happen if you did A or B, get a clearer idea of their force by calculating the odds of outcomes conditional on actions, and test your convictions by considering potential objections. The activity of reasoning thus involves many sub-actions in its service.

I believe that in phenomenological terms, the most realistic view on the next operation, *weighing reasons*, is that it is performed alongside the search – each new consideration, as it were, moves the deliberative gauge in one direction or another, or perhaps silences other considerations. (For one detailed picture, see Cullity's chapter in this volume.) In theory, it could be a separate stage. But the concurrent weighing model has a number of attractions. It fits nicely with the non-monotonicity of reasoning, as each individual consideration in turn may change the direction of the gauge. Now you lean towards A-ing, now towards B-ing. Where you end up when the music stops is your take on the balance of reasons. It is important that the process of seeking and weighing reasons has a diminishing marginal utility – you'll soon be in a situation in which you'll better conform to your objective reasons by engaging in some other activity than reasoning. (Consider how much time gets wasted when two or more philosophers try to decide which restaurant to go to.) To be reasons-responsive, you have to somehow register that you've run through enough

considerations. You might form the explicit belief that what you've rehearsed are all relevant facts, but that isn't necessary. Let's say there is a *subjective termination point* for reasoning.

I agreed with Broome and Boghossian that for the process to be reasoning, your conclusion-action or conclusion-attitude must be caused by your premise-attitudes *in virtue of* its seeming right to you to base your conclusion on them. You somehow treat the contents of the inputs to reasoning, taken together, as a reason for the act or the intention. As Dancy puts it, "when someone deliberates well and then acts accordingly, the action done is the one favoured by the considerations rehearsed in the deliberation, taken as a whole. It is a response to those considerations as together calling for or favouring it." (2014, 4) But you don't need to have the *belief* that they favour or call for the conclusion. So what is it to treat the premises as a reason? The Reasons-Responsiveness theorist will think it is the same sort of thing as treating anything as a reason. You have a positive attitude towards performing the act or forming the intention or belief, given the premises, and negative attitude towards not doing so. These attitudes are manifest in many different ways in how you play the game of giving and asking for reasons. When the fact that you move from accepting the premises to the practical conclusion is in part explained by your having such attitudes, you treat the former as a reason for the latter.

I will next work through the concrete cases I started with to make the case for the plausibility of this sketch. First, in Advance reasoning is triggered by the thought of as-yet unbudgeted windfall: what to do with it? The first option that comes to your mind is buying a new guitar. You may have seven already, but perhaps not an authentic Sienna Sunburst American Standard Stratocaster. That it would be nice to have one favours buying one. But then another thing occurs to you. Your family is growing and the old car is small, so you need a new car. This favours buying a car. Could you do both? You ruefully say to yourself "I can't have both". So you're pulled in two different directions. Next, you make the

comparative judgment that the good at stake in buying the car is more important – as a decision theorist would put it, the expected utility of buying the car is greater than the expected utility of buying a guitar. The deliberative gauge now points towards buying a car, as you take it to be more strongly favoured than buying a guitar. You don't see the need for further consideration. In response to taking the balance of reasons to favour buying a car, you form the intention to do so, expressing this to yourself by saying "So, I'll buy a car." By Tracking Sufficient Would-Be Reasons*, this is good reasoning, because if things are as you assume, you do indeed have sufficient reason to (intend to) buy a car – you're correct in taking the balance of reasons to favour buying it, and so respond to inputs to your reasoning as you should.

Instrumental reasoning, such as Paris, is actually harder to make sense of on the Reason-Responsiveness Conception. The big question is: what is the major premise of such reasoning? In Section 1, we saw that it's not that you intend to E. One possibility would be to follow Broome and make use of the notion of marked content. But it is not such contents that favour doing something, but rather potential states of affairs, as Dancy emphasizes. So I think the best way is to make what may seem like a radical move and drop the major premise altogether.[18] My suggestion is that the intention plays a different, non-premise role. When you intend to E, you take the fact that M-ing facilitates E-ing as a reason to M, of a strength proportionate to the degree of facilitation. The intention will also motivate you to engage in reasoning in the first place.[19]

---

[18] Interestingly, Aristotle's examples of the practical syllogism also lack any premise corresponding to intention or desire. Anscombe, too, says that "the end ought to be specified, but the specification of the end is not in the same position as the premise" (1989, 378). Inspired by her, Alvarez says that "the premises of practical reasoning jointly show what good (or apparent good) there is in my doing something. But among those premises we won't normally find the mere fact that I want to do something because that fact does not normally contribute to the good-making characteristics of the action" (Alvarez 2010, 362).

[19] Cf. Alvarez, who says that when we state the agent's desire or intention in describing practical reasoning, "we state the *point* of the piece of practical reasoning – which is the goal for the sake of which the agent engages in practical reasoning and for the sake of which he acts, if he acts on that reasoning" (2010, 365).

Take the Paris case. It may be that the content of your intention is "I will be in Paris in two weeks", as Broome claims. But on the view I'm sketching, it is not a premise in your practical reasoning. It is rather a part of the set-up. If you say it to yourself, you're just getting yourself to focus on the topic you'll be reasoning about. Then you say to yourself "Flying Ryanair is the cheapest way to get to Paris." Given your intention, you take this fact (as you believe) to favour flying Ryanair, perhaps sufficiently. Alas, you next remind yourself that "Ryanair are greedy bastards who will rip you off if you close your eyes for a second." This may either be a subjective reason not to fly Ryanair or an attenuator for the subjective reason you do have. Then you say to yourself "Flying Aer Lingus is the cheapest and most convenient of the remaining options". This state of affairs, as you see it, favours flying Aer Lingus in the same way as the content of the first instrumental premise favoured Ryanair. Next, you weigh the reasons against each other, and say "So, I'll fly Aer Lingus" to indicate your decision – formation of a more specific intention on the basis of the weighing operation that concluded with taking there to be sufficient reason for your conclusion-response. Assuming that you don't have evidence that considering more options would significantly improve the chances of doing what you have most objective reason to do, so that it's rational to close off deliberation, you are now rationally required to intend to fly Aer Lingus, and a bit of additional instrumental reasoning takes you to the intention to book a flight now. (As is typical, you may need to engage in yet further instrumental reasoning to figure out the best way to book next.)

Again, by Tracking Sufficient Would-Be Reasons*, you have reasoned well, since assuming you have sufficient reason to intend to travel to Paris, and that your non-normative beliefs about the various ways of getting there are correct, you do indeed have sufficient reason to (intend to) book a flight with Aer Lingus now.

## 4 Conclusion

In this chapter, I have presented two competing contemporary views about the nature and correctness of practical reasoning. My discussion has taken place at a relatively high level of abstraction – I have said very little about what practical rationality or good reasoning requires. Indeed, on the Reasons-Responsiveness view there isn't much that could be said about good reasoning in advance of substantive inquiry into objective reasons. According to the version I've tentatively sketched, what makes a thought process practical reasoning is that it is (fallibly) guided by what one takes to be reasons for action or intention, and what makes it good reasoning is that these reason-takings take one to an action or intention that one really does have sufficient reason for, if the doxastic inputs to reasoning are true and non-doxastic inputs permissible (in the sense of there being sufficient reason for them). Although I have pointed to some advantages of the Reasons-Responsiveness Conception over the Rule-Guidance Conception, the debate is very much open.

I will end with a brief observation about the broader significance of this debate. Only the Rule-Guidance Conception grants normative significance to requirements of rationality as such. This means that any attempt to derive authoritative moral demands from practical reason must presuppose that some form of Rule-Guidance is correct. If, in contrast, the Reasons-Responsiveness Conception is correct, the authority and content of morality cannot derive from rationality. What makes reasoning correct derives from truths about reasons that are independent of rationality. The Reasons-Responsiveness Conception is thus bad news for Kantians and good news for Aristotelians and, arguably, Humeans.

## References

Alvarez, M. (2010), 'Reasons for Action and Practical Reasoning', *Ratio* 23/4: 355–73.

Anscombe, E. (1989), 'Von Wright on Practical Inference', in P.A. Schilpp (ed.), *The Philosophy of Georg Henrik von Wright* (Chicago: Open Court): 377–404.

Audi, R. (2006), *Practical Reasoning and Ethical Decision* (London: Routledge).

Boghossian, P. (2014), 'What Is Inference?' *Philosophical Studies* 169: 1–18.

Brandom, R. (1994), *Making It Explicit*, (Cambridge, MA: Harvard University Press).

Brandom, R. (2000), *Articulating Reasons*, (Cambridge, MA: Harvard University Press).

Bratman, M. (2009), 'Intention, Belief, Practical, Theoretical', in S. Robertson (ed.), *Spheres of Reason: New Essays in the Philosophy of Normativity* (Oxford: Oxford University Press): 29–61.

Broome, J. (2000), 'Normative Requirements', in J. Dancy (ed.), *Normativity* (Oxford: Blackwell): 78–99.

Broome, J. (2004), 'Reasons', In J. Wallace, M. Smith, S. Scheffler, and P. Pettit (eds.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz* (Oxford: Oxford University Press): 28–55.

Broome, J. (2007), 'Does Rationality Consist in Responding Correctly to Reasons?' *Journal of Moral Philosophy* 4/3: 349–74.

Broome, J. (2009). 'The Unity of Reasoning?'. In S. Robertson (ed.), *Spheres of Reason: New Essays in the Philosophy of Normativity* (Oxford: Oxford University Press): 62–92.

Broome, J. (2013), *Rationality Through Reasoning* (Blackwell: Oxford).

Brunero, J. (2009), 'Against Cognitivism about Practical Rationality', *Philosophical Studies* 146: 311–25.

Brunero, J. (2014), 'Cognitivism About Practical Rationality', *Oxford Studies in Metaethics* 9: 18–44.

Dancy, J. (2009), 'Action, Content, and Inference', in H.J. Glock and J. Hyman (eds.),

*Wittgenstein and Analytic Philosophy* (Oxford: Oxford University Press): 278–98.

Dancy, J. (2014), 'From Thought to Action', *Oxford Studies in Metaethics* 9: 1–17.

Harman, G. (1976), 'Practical Reasoning', *Review of Metaphysics* 29/3: 431–63.

Hume, D. (1739–40/1978), *A Treatise of Human Nature*, ed. L. A. Selby-Bigge, 2nd rev. edn., P. H. Nidditch (Oxford: Clarendon Press).

Kaplan, D. (1989), 'Demonstratives', in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan* (Oxford: Oxford University Press): 481–563.

Kauppinen, A. (2014), 'Fittingness and Idealization', *Ethics* 124/3: 572–88.

Kauppinen, A. (2015), 'Favoring', *Philosophical Studies* 172/7: 1953–71.

Kiesewetter, B. (2013), *The Normativity of Rationality* (PhD Dissertation, Humboldt University).

Kolodny, N. (2005), 'Why Be Rational?' *Mind* 114: 509–63.

Kolodny, N. (2007), 'State or Process Requirements?' *Mind* 116: 371–84.

Kolodny, N. (2008), 'Why Be Disposed to Be Coherent?' *Ethics* 118: 437–63.

Kolodny, N. (this volume), 'Instrumental Reasons'.

Korsgaard, C. (1997), 'The Normativity of Instrumental Reason', in G. Cullity and B. Gaut (eds.), *Ethics and Practical Reason* (Oxford: Clarendon Press): 215–54.

Lord, E. (2014), 'The Coherent and the Rational', *Analytic Philosophy* 55/2, 151–75.

McHugh, C. and Way, J. (forthcoming), 'What Is Good Reasoning?' *Philosophy and Phenomenological Research*.

Parfit, D. 2011, *On What Matters, Vol. 1* (Oxford: Oxford University Press).

Raz, J. (2005), 'The Myth of Instrumental Rationality', *Journal of Ethics and Social Philosophy* 1: 1–28.

Scanlon, T. (2007), 'Structural Irrationality', in G. Brennan, R. Goodin, F. Jackson, and M. Smith (eds.), *Common Minds: Themes from the Philosophy of Philip Pettit* (Oxford: Oxford University Press): 84–103.

Schroeder, M. (2004), 'The Scope of Instrumental Rationality', *Philosophical Perspectives* 18: 337–64.

Schroeder, M. (2009), 'Means-End Coherence, Stringency, and Subjective Reasons', *Philosophical Studies* 143/2: 223–48.

Setiya, K. (2007), 'Cognitivism About Instrumental Reason', *Ethics* 117/4: 649–73.

Smith, M. (1994), *The Moral Problem* (Oxford: Blackwell).

Streumer, B. (2007), 'Inferential and Non-Inferential Reasoning', *Philosophy and Phenomenological Research* 74/1: 1–29.

Sylvan, K. (2015). 'What Apparent Reasons Appear to Be'. *Philosophical Studies* 172: 587–606.

Tenenbaum, S. (2007), 'The Conclusion of Practical Reason', in S. Tenenbaum (ed.), *Moral Psychology* (Rodopi, New York/Amsterdam): 323–43.

Wallace, R. J. (2001), 'Normativity, Commitment, and Instrumental Reason', *Philosophers' Imprint* 1/4: 1–26.

Way, J. (2010), 'Defending the Wide-Scope Approach to Instrumental Reason', *Philosophical Studies* 147/2: 213–33.

von Wright, G. H. (1963), 'Practical Inference', *Philosophical Review* 72/2: 159–79.