REASONING AND NORMATIVITY

Franziska Felicitas Poprawe

DPhil Thesis

Exeter College
University of Oxford

I, Franziska Felicitas Poprawe, hereby certify that this thesis of approximately 75000 words in length is based on my work and has been written by me. It has not been submitted for a higher degree before.



Oxford January 31 January 2020

Abstract

This thesis addresses three questions: (I) What is theoretical and practical reasoning? (II) What is it to respond to normative reasons? (III) What are standards of evaluation for reasoning? Regarding (I), I first critically assess two views which have recently dominated the debate on the nature of the mental activity of reasoning. I argue that different versions of the *Rule-Following View*, according to which reasoning is essentially a rule-following activity, are problematic. Further, I argue that the *Reasons View*, which holds that reasoning consists in responding to normative reasons, is assumed by many but defended by few and remains underdeveloped. In particular, it is explanatorily unsatisfactory because, while it explains reasoning in terms of responding to normative reasons, it lacks an account of what it is to respond to normative reasons. I then provide a new theory of the nature of reasoning, the *Approval Account*. Reasoning is a transition in thought from some attitudes to another which the agent approves of. Regarding question (II), I propose a novel account of what it is to respond to normative reasons: it is a mental process that leads an agent from some attitudes to a response, where this process manifests the agent's implicit conception of what is right to believe and do, and is under the agent's (virtual) control. This account of responding to reasons allows me to systematically establish the claim that reasoning is a way of responding to normative reasons. With regards to question (III), it follows that correct reasoning is correctly responding to (supposed) normative reasons. I further argue that we must distinguish between the correctness, the goodness and the aim of reasoning.

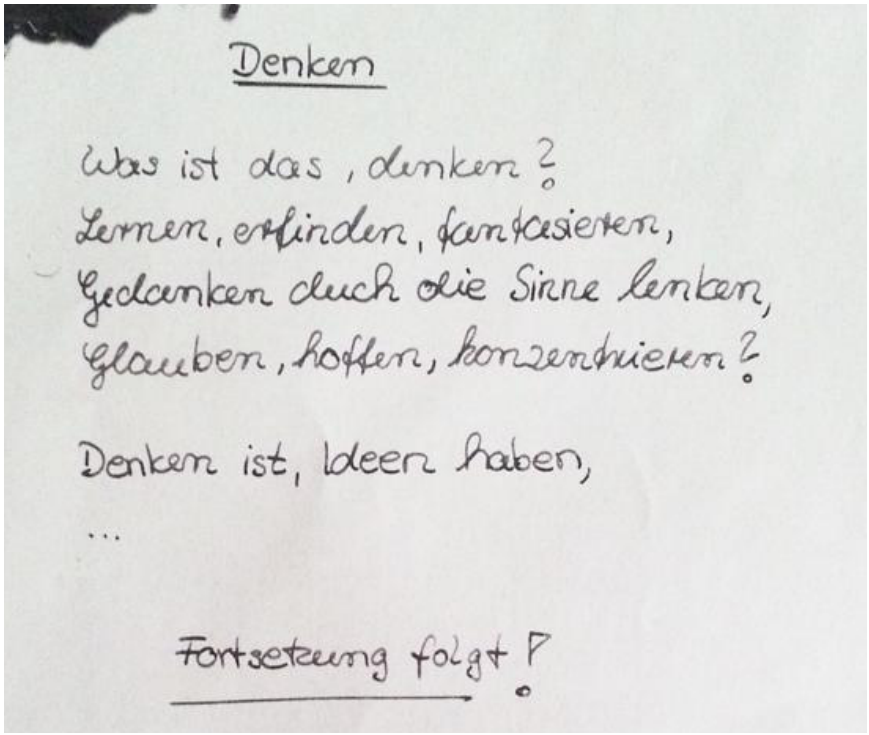
Acknowledgements

I am very grateful to many philosophers and friends who supported me during the writing of this thesis. While I was at Oxford, I've had inspiring discussions with many philosophers, but received especially helpful feedback on my writings from John Gibbons, John Hyman, Alison Hills, Natalia Hickman and Lucy Campbell, Annina Loets and especially Ben Brast-McKie.

I am very thankful to Nick Southwood for inviting me to visit the Australian National University in 2017. During my time there, I learned a lot from him, and also from Philip Pettit and Jesse Hambly. Last year, thanks to Jay Wallace, I had the opportunity to spent one semester at the University of California, Berkeley. I thank Jay for our discussions on responding to reasons, and also Niko Kolodny for commenting on a draft of a chapter. I thank Franz Altner, for reading so much of my work over the past three years. Most of all, however, I wish to express my sincere gratitude to my supervisors at Oxford, John Broome and Ralf Bader. They have been very generous with their time, and I feel extremely fortunate to have had the opportunity to learn from them. I will never forget the long, stirring and inspiring DPhil supervision meetings at Corpus Christi and Merton College. I am extremely grateful.

The writing of this thesis was financially supported by various institutions. Without the Hanfling Scholarship I could not have accepted the offer to the DPhil. I will always be thankful to Hanflings and the Department of Philosophy at Oxford for giving me this immense opportunity. I am also thankful to the Cusanuswerk and Studienstiftung des Deutschen Volkes for their support.

Finally, I thank my parents, my three brothers and Till for supporting me in writing this thesis. I dedicate this thesis to them.



Aachen, 23 January 2006

Contents

Introduction 1

1. The *Rule-Following View*

1.1. What is Reasoning? 9
1.2. The Taking Condition 11
1.3. Rational Requirements and Rules 18
1.4. Meta-Propositional Beliefs 28
1.5. Endorsing Transitions of Thought 31
1.6. The Fittingness View 34
1.7. Conclusion and Outlook 41

2. The *Reasons-View*

2.1. Introduction 44
2.2. The Popularity of the *Reasons View* 45
2.3. Paul Grice 48
2.4. Joseph Raz 50
2.5. Antti Kaupinnen 53
2.6. Jonathan Dancy 55
2.7. Rules or Reasons? 61
2.8. Reasoning and Normativity 62
2.9. Conclusion 64

3. The *Approval Account*

3.1. Introduction 67
3.2. The Approval Account 70
3.3. The Normativity of Reasoning 80
3.4. An Implicit Conception of Normative Reasons 92
3.5. Different Kinds of Premises 97
3.6. Meeting the Desiderata 100
3.7. Three Features of Reasoning 104
3.8. The Intentionality of Reasoning 109
3.9. Instrumental Reasoning 114
3.10. Conclusion 121

4. Responding to Normative Reasons

4.1. Introduction 123

4.2. Believing and Intending in Light of Considerations 129

4.3. Supposed Normative Reasons 138

4.4. A Conception of Rightness 149

4.5. A Conception of Epistemic Value 155

4.6. The Capacity to Respond to Normative Reasons 159

4.7. Reasoning and Responding to Reasons 162

4.8. Responding Correctly to Normative Reasons 166

4.9. Explaining the *Normative Leap* 168

4.10. Conclusion 171

5. Correctness, Goodness, and the Aim of Reasoning

5.1. Evaluations of Reasoning 181

5.2. Correctness 182

5.3. Correct Theoretical Reasoning 191

5.4. Correct Practical Reasoning 196

5.5. Goodness 205

5.6. The Aim of Reasoning 209

5.7. Summary 214

Conclusion 215

References

Introduction

1. Reasoning, Normative Reasons, and Evaluations

A great deal of what we think, want, and do is the result of mental processes that we have no immediate control over. But as rational agents, we are not entirely at the mercy of sub-personal, automatic mental mechanisms. We have some control over our cognitive life. We have the capacity to *reason*. This capacity allows us to figure things out, consciously form judgements about how the world is or ought to be, deliberately make decisions, and bring our minds into a coherent order. When we reason about what to do or what to believe, we can attend to considerations that are reasons for us to do this rather than that, and reasons to believe this rather than that. We can recognise and respond to the reasons we have and bring our minds and actions in accordance with what they demand. Both, the capacity to reason and the capacity to respond to reasons seem to be distinctive features of human nature. They make us subject to warranted criticism if we get things wrong. This picture of reasoning and reasons, as intuitive as it may be, raises some fundamental questions, in particular, the following three.

(I) What is reasoning?

Reasoning is not argumentation or proof, but a psychological process through which we, for example, acquire or change beliefs, and form intentions to do things (Harman 1984:115). But not any kind of thought process through which we form or change beliefs and intentions is reasoning. Reasoning is a mental activity, something we do, and not all thinking is of this kind. Following Kahneman (2011: 20-22), it has become common to distinguish between two types of thinking: system 1 and system 2.[1] They are described as follows:

[1] The distinction of system 1 and system 2 was first introduced by the psychologists Stanovich and West (2000).

> *System 1* operates automatically and quickly, with little or no effort and no sense of voluntary control.
>
> *System 2* allocates attention to the effortful mental activities that demand it, including complex computations. The operations of System 2 are often associated with the subjective experience of agency, choice, and concentration.

System 1- thinking allows us to do many things at cognitive ease, like understanding and speaking simple sentences, recognising faces, or even driving on an empty road. System 2- thinking, in contrast, is controlled processing; it allows us to "construct thoughts in an orderly series of steps." For example, System 2- thinking is involved in calculating your taxes, or weighing the reasons for and against a career in banking. The model of the two systems of thinking raises many questions, but it motivates the distinction between, on the one hand, unconscious, passive, automatic thinking (paradigm cases of System 1- thinking) and, on the other hand, thinking that is conscious, active and under the control of the agent (paradigm cases of System 2). In this thesis, I am concerned only with the latter. 'Reasoning' refers to the mental *activity* through which we consciously acquire, discard or revise attitudes, specifically beliefs and intentions. Psychologists like Kahneman say that when System 1 fails, we must invoke system 2 and listen to the 'voice of reason'. But they do not tell us what this means. What does the "effortful mental activity" of System 2- thinking consist in? What differentiates *acts* of reasoning from automatic, unconscious processes that also cause us to form, revise or discard attitudes? Suppose I have a train of thought which I express as follows:

> "If global warming causes harm, we must stop emitting CO_2.
>
> Global warming does cause harm.
>
> So, we must stop emitting CO_2."

where all three sentences express beliefs. I acquire the belief that we must stop emitting CO_2 on the basis of the belief that if global warming causes harm, we must stop emitting CO_2 and the belief that global warming causes harm. If I *derive* this conclusion through reasoning, the first two beliefs do not merely cause me to believe that we must stop emitting CO_2. There is something I *do*. But what? What makes this an instance reasoning, something I *do*?[2] This is the question of the nature of reasoning. It concerns reasoning that leads to a belief like here, which is called theoretical reasoning, as well as reasoning that leads to intention, which is called practical reasoning. For example, I could reason:

> "My boss gave me an unexpected bonus.
>
> I have no immediate need for the money.
>
> My friend needs surgery she cannot afford. She won't accept the money as a gift but will accept a loan.
>
> So, I shall lend her the money."

where the first four sentences express beliefs and the last an intention. When I ask what reasoning is, I refer to both, practical and theoretical reasoning.[3] Our capacity to reason has been associated with our free will, our responsibility for at least some of the beliefs we hold, our understanding of moral and other matters, the ability to motivate ourselves to do what we believe we ought to do, and our distinctive human nature in general. The question of what reasoning is has extensive relevance.

(II) What is it to respond to reasons?

There are facts in the world that are or at least appear to be reasons for us to do and believe certain things. For example, the fact that Sam is injured is a reason to call an ambulance. The

[2] As I use the term, reasoning is always active, so "active reasoning" is, strictly speaking, a pleonasm.
[3] If there is reasoning with to attitudes such as desires or hopes, for example, it is not discussed in this thesis.

fact that exercising keeps you healthy is a reason to exercise. And the fact that only Smith's fingerprints are at the crime scene is a reason to believe Smith is the guilty party. These reasons are so-called normative reasons.[4] They are facts that "count in favour" of certain responses (Scanlon 1998:17).[5] As humans, we have the capacity to recognise, understand and respond to these reasons. We can treat things as reasons to do or believe something, or to admire someone. We do not just *react* to the world (as all living creatures do), but we can do things for the reasons we have to do them.

Suppose, for example, Tom sees an elderly lady stumble and fall on the other side of the street, and crosses the street to help her. It seems plausible to say that Tom responds to a normative reason he has to cross the street – that the lady fell. He realises what the situation requires him to do. He sees and treats the fact that the lady fell *as a reason* to help her. If all of this is true of Tom, why is that so? What is it for Tom to 'see' something as a reason for him to do something, for example? Tom may reflect on the situation or act intuitively. But it is possible – indeed probable – that Tom does not think about what reasons he has to do what. No thoughts about 'reasons' or 'favouring' need to appear in his consciousness. All he may think about are the observable features of the situation: the lady stumbling, the spatial distance between him and her, etc. Why, then, would Tom's action count as a response to *normative reasons*, rather than as a reaction to these descriptive *facts*? How does normativity enter his mind so that his action counts as a response to the normative reason he has? The ability to do things for the reasons we have to do them is arguably necessary for our actions to be morally worthy (Markovits 2010) or virtuous (McDowell 1979). The ability may be the mark for

[4] Normative reasons are to be distinguished from motivating reasons and explanatory reasons, see e.g. Grice (2001: ch. 2), Alvarez (2010) Broome (2004) and Hyman (2015: ch. 6).

[5] I assume that reasons are facts. This is currently the mainstream view on the ontology of normative reasons, according to e.g. Kiesewetter (2017(; Lord (2018); Worsnip (2016). However, some epistemologists argue that reasons for belief are mental states, among them Davidson (2001b), Turri (2009) and Pollock (1974: pages). Arguments against these defences of so-called mentalism or psychologism about reasons are developed by Ginsborg (2006), Sylvan (2016), Littlejohn (2012: pages), respectively. I don't deny that some reasons are facts about our mental states, only the claim that all reasons are.

personhood, just as the ability to have second-order desires, i.e. to desire to desire things (Scanlon 1998:23; Frankfurt 1971). Further, we may say beliefs are fully justified only if we hold them for the reasons we have *to* hold them (e.g. Evans 2013; Silva 2015). An account of what it is to respond to normative reasons, that is, an account of what it is to act and believe for normative ('good') reasons, seems to get us a long way in understanding our place in the world.

(III) What are the standards of evaluation for reasoning?

It is common to evaluate reasoning. We say this person's reasoning is correct while that person's reasoning is not very good. But what is correct reasoning? On what grounds do we fail when we reason incorrectly? And what is good reasoning? It is not clear whether these questions refer to the same or different standards of evaluation. In addition, some philosophers say that reasoning has an aim. But what does it mean to say, for example, that reasoning aims at the truth and the good? And what is it for reasoning to reach its aim? Lastly, there are normative reasons that pertain to our reasoning, and we can fail to comply with those reasons. In some situations, we may have reasons to reason in certain ways (when you consider whether to save the many or the few, for example), while in others we may have reasons to not be reasoning at all (and only run from the fire, for example). One might think that different moral theories set a different standard for correct practical reasoning. Can we say what correct or good reasoning is, independent of substantive normative commitments?

2. This Thesis

This thesis provides answers to all three questions. Even though reasoning and responding to reasons has been examined within studies of structural rationality, epistemology and moral philosophy, we lack a comprehensive theory of (I) the nature of reasoning, (II) the nature of responding to normative reasons and (III) the standards against which we evaluate reasoning.

I develop a theory of reasoning that unifies theoretical reasoning (towards belief) and practical reasoning (towards intention), overcoming the dissociation of the two (due to Aristotle and later Hume), that continues to be enforced by the compartmentalisation of contemporary philosophy.[6] The theory establishes a close link between the mental activity of reasoning and our responding to normative reasons. More generally, the aim is to elucidate the relation between our reasoning and all that which is normative.

The first three chapters concern question (I) *What is reasoning*? In the recent philosophical debate on this question, two major views have emerged that are mostly assumed to be mutually exclusive. I discuss them in chapters 1 and 2, before presenting my own account in Chapter 3.

Chapter 1 examines the *Rule-Following View*, according to which reasoning is essentially a rule-following activity. When you reason, you intentionally follow and not merely mechanically comply with the rules of reasoning. I present and raise objections against different versions of this view, looking at Paul Boghossian's (2014) "Taking Condition", John Broome's theory of how the requirements of rationality govern reasoning, and Philip Pettit's (1996; 2007a) meta-propositional account of reasoning. I conclude that none of them is fully satisfactory. But there is a common thread which runs throughout: they all agree that what makes a thought process count as reasoning is that the agent holds an attitude of endorsement towards the process. The discussion leads me to formulate three desiderata for a theory of reasoning that is based on this assumption. Lastly, I discuss McHugh and Way's *Fittingness View,* according to which reasoning is essentially guided by the aim of getting fitting attitudes. It is either implausible or best understood as a version of Broome's rule-following view.

[6] While practical reasoning (or deliberation) has mostly been discussed in moral philosophy and the philosophy of action, theoretical reasoning has been mostly discussed in epistemology. The recent trend to merge the debates on practical and theoretical rationality (Broome 2013a; Gibbons 2013; Worsnip 2016; Kiesewetter 2017; Lord 2018b) is increasingly changing over to the debates on reasoning in the practical and theoretical domain. These debates are still rather misaligned, and it is often unclear how concepts and theories used in one domain translate to the other.

Chapter 2 discusses the second major view predominant in the current debate on reasoning: The *Reasons View*. According to this view, reasoning is essentially responding to normative reasons. Many moral philosophers and epistemologists assume a version of the *Reasons View*, but few defend it in detail. I present and assess what I take to be the most elaborate versions of the *Reasons View*, namely those defended by Paul Grice (2001), Joseph Raz (1978; 2015), Antii Kauppinen (2018) and Johnathan Dancy (2018). My conclusion is that the defenders of the *Reasons View* do not provide an account of the nature of reasoning that is sufficiently explanatory. This is because they define reasoning in terms of responding to reasons, but do not give an account of what that consists in.

In **chapter 3**, I propose my novel account of reasoning. According to the *Approval Account*, reasoning is a transition in thought from some attitudes to another which the agent approves of. The attitude of approval is evaluative in nature. The standard against which reasoners evaluate their transitions when they approve of them is set by their (perhaps) implicit conception of what kind of response is permissible in a certain context and why. That is, the basis of the approval is the reasoner's (implicit) conception of normative reasons and how they aggregate to make certain beliefs or action permissible. Reasoning is thus an inherently normative activity. I explain what I mean by this, and how the A*pproval Account* meets the desiderata provided in chapter 1. Though normative reasons play an essential role in reasoning according to the *Approval Account*, the account itself does not claim that reasoning involves or is responding to normative reasons, the nature of which still remains to be explained. I turn to this task next.

Chapter 4 addresses question (II) *What is it to respond to normative reasons*? Put roughly, I argue that responding to reasons is a mental process that leads an agent from some attitudes to a response, where this process manifests the agent's implicit conception of normative reasons and is under the agent's (virtual) control. Treating a consideration as a

normative reason requires neither believing that it is a reason, nor having a 'sensitivity to normative properties.' When we respond to reasons, we express our implicit conceptions of what is right (or good) to do, and what is right to believe. If these conceptions are false, we treat things as reasons which are not reasons. On the basis of this account of what it is to respond to reasons, I argue that reasoning, understood as the *Approval Account* proposes, is one way of responding to normative reasons.

In **chapter 5**, I turn to question (III) *What is correct reasoning?* We must distinguish between different dimensions of evaluation of reasoning: correctness, goodness, and the aim of reasoning. I argue that reasoning is a correctness-fixing kind (Thomson 2008; Geach 1956). That is, the standard of the correctness of reasoning derives from the nature of reasoning – from what it is to reason. Correct reasoning is correct *qua* reasoning, just like a correct chess move is correct *qua* chess move. According to the *Approval Account*, you reason correctly whenever you manifest a correct conception of normative reasons for belief and action. Further, an instance of reasoning can be good *in some respect* or good *for* something, but not good *qua* reasoning. Lastly, I argue that the metaphor of the aim of reasoning is misleading on some interpretations, and false on others. Reasoning aims at the truth and the good only in so far as its standard of correctness pertains to truth and goodness.

I conclude that the three questioned I raised in the outset cannot be addressed in isolation. The capacity to reason is the capacity to respond to normative reasons. Without the ability for normative thought, without the ability to understand the difference between what is and what ought to be, we would not be the kind of reasoners that we are.

1. The *Rule-Following View*

1.1. What is Reasoning?

Many of our thought processes merely happen to us, without our active engagement. In the language of psychologists like Kahneman (2011: 20-22), they are manifestations of 'System 1'– thinking. Passive thought processes that happen "quickly, with little or no effort and no sense of voluntary control." When you walk to your favourite café while talking on the phone, for example, you think, but do not actively reason about the steps you take to get to the café, or which turn to take. When you see an old friend across the street on the way, you automatically recognise their face. Associative processes are another case in point. We are often prompted to think about things without choosing to do so, and may even be surprised or annoyed when we notice we are thinking about them. Though a great deal of our daily life is guided by passive, automatic information processing, we are not entirely subject to it. We have some control over our cognitive life. We have the capacity to reason. What I call 'reasoning' is a paradigm case of so-called 'system 2'- thinking. It is an "effortful mental activity" (Kahneman 2011: 20-22). It is the activity through which we form, revise or discard attitudes on the basis of others. For example, I could reason as follows:

> "It is 5 pm.
>
> If it is 5 pm, then Blackwell's is closed.
>
> So, Blackwell's is closed."

where each sentence expresses a belief. I acquire the belief that Blackwell's is closed on the basis of the belief that it is 5 pm and the belief that if it is 5 pm, then Blackwell's is closed.[1] Let us call the attitudes that I reason from the *premise-attitudes* and the attitude that I form

[1] That is, my belief that Blackwell's is closed is caused by the other two in the right kind of, non-deviant way. I assume that the basing relation is a causal relation as e.g. Turri (2011) argues, see Evans (2013) for critical response to this.

on their basis the *conclusion-attitude*. The contents of those attitudes are the *premises* and the *conclusion*, respectively.[2] They are that which we think about when we reason. We reason with attitudes, about their contents.

The "so" is not part of the content of the conclusion-belief, but indicates that I derive the conclusion through reasoning (Pettit 2007a; Broome 2013:23). But what does it mean to say that 'I derive' the conclusion? What is it to reason to the conclusion? After all, the first two beliefs could also simply cause me to believe that Blackwell's is closed. I could acquire this belief through passive 'system 1'- thinking. An account of the nature of reasoning explains the difference between, on the one hand, merely causal processes through which we acquire attitudes on the basis of others, and, on the other hand, instances of active reasoning. When I reason to the belief that Blackwell's is closed, I acquire that belief *because* I believe it is 5 pm and that if it is 5 pm, then Blackwell's is closed, where the 'because' does not only signify a causal relation between beliefs, but also an intentional one.

The question of the nature of reasoning has recently generated a considerable debate. Two families of views dominate the discourse: On the one hand, some authors have defined reasoning in terms of rule-following. At the centre of their accounts is the

Rule-Following Claim Reasoning is a rule-following activity.

On the other hand, there is the view that reasoning is best explained in terms of responding to normative reasons. The claim at the core of these views is the

Reasons Claim Reasoning is a way of responding to normative reasons for belief or action.

[2] I am here following Broome's (2013: ch. 13) terminology, and shall use it throughout the thesis.

The two views are largely seen as mutually exclusive. This chapter presents and critiques accounts of reasoning that fall into the first group. In **sections 1.2, 1.3, and 1.4,** I discuss Paul Boghossian's, John Broome's and Philip Pettit's versions of the rule-following view, respectively. In **section 1.5**, I extract a common feature of all of them: they all claim that what is distinctive of and essential to reasoning is that the agent in some way or the other endorses the thought process. This is indeed a common thread in many accounts of reasoning. **Section 1.6** discusses Conor McHugh and Jonathan Way's *Fittingness View* of reasoning. I argue that this is best interpreted as a cousin of Broome's rule-following account of reasoning. **Section 1.7** concludes that no version of the *Rule-Following View* is fully convincing and motivates the next chapter, which presents and critiques different versions of the *Reasons View*.

1.2. The Taking Condition

In his search for an answer to the question of what reasoning is, Boghossian (2014) turns to Frege, who says

> "to make a judgement because we are cognisant of other truths as providing a justification for it is known as inferring" (Frege 1979:5).[3]

Frege and Boghossian use the term 'inferring' to refer to what I call theoretical reasoning, but this terminological difference should not bother us. Boghossian notices that Frege's account of theoretical reasoning excludes the possibility of reasoning with false beliefs (for what we take to be true might not be true), as well as the possibility of incorrect reasoning (we might be mistaken about what justifies what). He therefore reformulates Frege's claim as follows:

[3] In the German original version, Frege uses the word "Rechtfertigungsgründe", which means "justifying reasons." He says, "to make a judgment because one is cognisant (conscious) of other truths as being justifying reasons, is known as inferring." Frege clearly understands reasoning, or inferring, or making a judgement, in normative terms. Boghossian, and even more so Broome, depart from the this conception, and conceptualise reasoning in purely non-normative terms. I argue in chapter 3 that this is where thy took a wrong turn.

thinking counts as reasoning if and only if the agent "*takes* his conclusion to be supported by the *presumed* truth of those other beliefs." This has led him to defend the much-debated

> "Taking Condition: Inferring necessarily involves the thinker *taking* his premises to support his conclusion and drawing his conclusion *because* of that fact" (Boghossian 2014:5).

The premises are the contents of the premise-beliefs. Reasoning from believing p to believing q, Boghossian says here, is judging that q *because* one takes the presumed truth of p to provide support for q. According to Boghossian, this 'taking' is what differentiates passive thinking from active reasoning. But what is it to take one's premises to support one's conclusion? Boghossian considers different interpretations of this 'taking', eventually concluding that it is best understood in terms of rule-following, which, he concludes, is unanalysable. I here briefly summarise the most influential points of Boghossian's arguments, using as an example modus-ponens reasoning, that is, reasoning from the belief that p, and the belief that if p then q, to the belief that q.

Boghossian first considers two doxastic construals of 'taking', that is, two interpretations in terms of belief. According to the first interpretation, the 'taking' is a meta-belief with normative content: a belief about one's belief in a normative proposition. In the example of modus ponens reasoning, the Taking Condition would entail that I count as reasoning if and only if I believe that [my belief that p and my belief that if p then q support the belief that q]. However, Boghossian argues that this proposal is implausible because it renders reasoning too demanding or sophisticated: it requires reasoners to have certain concepts that do not seem necessary for reasoning. For example, children can play games and thereby reason without having or deploying the concepts of believing or normative support.

According to the second doxastic interpretation, 'taking' is a first-order belief about entailment. (Philip Pettit calls these meta-propositional beliefs: beliefs about propositions (see 1.4)). In modus ponens reasoning, it is the belief that [p and if p then q entail q]. However, the Taking Condition thus understood would again be implausibly demanding; it would require reasoners to have the sophisticated concept of entailment.

According to the third doxastic construal of the 'taking' that Boghossian considers, the 'taking' is a first-order belief with conditional content, i.e. the belief that [if p and (if p then q), then q]. The problem Boghossian points out with this proposal arises when we ask what role the belief plays in the reasoning process. If it is a premise, then the Taking Condition says that you cannot reason to believing [q] from believing [p] and believing [if p then q] without the belief that [if p and (if p then q) then q]. This raises Lewis Carroll's (1895) regress problem. It implies that you would also need to hold the belief that [if (p and (if p then q) and (if p and (if p then q) then q)) then q] as a premise-belief in order to reason to believing q via modus ponens. The argument can be iterated infinitely, and thus implies that you would need an infinite number of premises to do simple modus ponens reasoning, which is false. So maybe the taking-belief in question is not a premise-belief, but a background condition. However, this leaves us to doubt why I count as reasoning in virtue of this taking-belief, if it is merely in the background and does not contribute to the thought process. Boghossian concludes that any doxastic construal of 'taking' seems to fail. He next attempts to conceptualise the 'taking' as an intuition, that is, an intellectual seeming or a "temptation to believe", as Ernest Sosa calls it. Boghossian rejects this proposal on the grounds that an intuition that some considerations imply another does not entail that the thought process in question is active. Any passive passage of thought may be accompanied by such an intuition. The Taking Condition, understood as an intuition, fails to meet the desideratum that the presence of the attitude of 'taking' should explain what makes the process active.

Because the doxastic and the intuitional accounts of the Taking Condition fail, Boghossian concludes we cannot understand 'taking' as an occurrent intentional state. Instead, he goes on, we should understand reasoning as rule-following. Boghossian attempts to analyse what it is to follow a rule as opposed to merely comply with a rule, for that seems to be the difference between a passive thought process that appears to be reasoning and an instance of reasoning. The first conception of rule-following he considers is this: you follow a rule only if you have an intentional state that represents the rule. To apply the rule is to grasp what the rule requires, subsume your premises under it, and realise that the rule's trigger-conditions are satisfied. You then draw the conclusion the rule prescribes. However, this complicated process seems to require reasoning itself. If following a rule requires reasoning, we cannot explain reasoning in terms of rule-following. Boghossian calls this the 'Inference Problem' for rule-following (Boghossian 2008), and because of it rejects this analysis. He next considers a dispositional construal of rule-following. To follow a rule R is to have a disposition to conform to R. But again, this that does not explain the activity of rule-following. A mere disposition to form the attitude upon considering some supposed facts could be a purely causal, passive process. He concludes that the concept of rule-following is unanalysable. Nevertheless, he claims, if you follow a rule, you treat your premises as reasons for believing the conclusion and thus satisfy the Taking Condition.

One problem with this argument is that it is not clear how treating the premises as reasons implies taking them to support the conclusion attitude. The latter seems to require that I treat the premises as giving me *sufficient reason* for this attitude. It is not enough to just treat them as reasons. The support that the Taking Condition appeals to seems to be a strong sort of support (see also Broome (2019)). But even if this could be solved, there is a more serious

issue with Boghossian's argument.[4] Rule-following does not entail treating things as reasons for belief. Boghossian claims that reasoning involves treating the premises as reasons and taking them to support the conclusion-attitude. But while 'taking p to support believing q' and 'treating p as a reason' have an inherent normative dimension, the rules of reasoning that Boghossian describes have no normative dimension. If, for example, I reason from believing p and if p then q to the belief that q, thereby following the modus ponens rule

"form the belief that q on the basis of believing p and believing if p then q",

it does not follow that I take anything to be support or a reason for anything. The rule I follow does not say anything about support or reasons. As Broome (2014a) argues, too, I can follow a rule without taking myself to have a reason for the conclusion. I can also follow a rule without taking myself to have any reason to follow it. The Taking Condition, referring to normative reasons and support relations, is not suitably interpreted in terms of rule-following if the rules are wholly non-normative in nature, and rule-following does not require normative thought. According to the Taking Condition, reasoning essentially involves treating considerations as normative reasons or responding to relations of normative support. But rule-following, at least as presented by Boghossian (and also by Broome), does not entail that the reasoner treats considerations as reasons.

Two intuitions seem to drive Boghossian's analysis of reasoning here. On the one hand, there is the natural thought that reasoning essentially involves taking the premises to *support* the conclusion. Reasoning is tracking relations of normative support; it is responding to the reasons we have to believe (or do) something. Let us call these claims versions of the

[4] McHugh and Way (2016b) also object to Boghossian construal of the Taking Condition in terms of rule-following.

Reasons Claim Reasoning is a way of responding to normative reasons for belief or action.

On the other hand, reasoning seems to be a rule-following activity. For example, modus ponens is a rule of reasoning, and in modus-ponens reasoning we follow that rule (at least if all goes well). These rules make no explicit mention of reasons, values or other normative concepts. Let us call this the

Rule-Following Claim Reasoning is a rule-following activity.

With his Taking Condition, Boghossian accepts a version of the *Reasons Claim*. He then commits to the *Rule-Following Claim* in his attempt to explain the Taking Condition. But, I argued, it is unclear how following a rule entails treating considerations as reasons or taking things to be reasons. It is an open question whether the *Reasons Claim* and the *Rule-Following Claim* are compatible at all, and whether an account of reasoning should accommodate one or both of them. The answer depends on how exactly we understand rule-following and responding to normative reasons. At the end of chapter **2**, I return to the question of which of the two claims, if any, should guide our analysis of reasoning.

To sum-up the discussion of Boghossian's paper, he argues that (i) reasoning requires 'taking' the premises to support the conclusion. What differentiates active reasoning from passive thought is that you satisfy the Taking Condition; it explains why the process counts as active. On the doxastic and intuitional construal of the Taking Condition, however, this turns out to be false. Boghossian thus argues that (ii) no non-circular analysis of rule-following can be given and that (iii) the Taking Condition is entailed by the fact that in reasoning, we follow rules. I have neither objected to (i) nor to (ii). But I have argued that the claim that reasoning is rule-following does not entail the Taking Condition. We are left with no plausible interpretation and defence of the Taking Condition, and thus with no explanation of what

reasoning is. In the heated debate that the Taking Condition has generated, most participants deny (i) and reach the same conclusion, amongst them McHugh and Way (2016b), Wright (2014), Broome (2014a), Hlobil (2014) and Siegel (2017). However, Hlobil (ms) and Kietzmann (2018) defend the Taking Condition by proposing their own interpretations of the attitude of 'taking'. I do not review this debate here, and only briefly respond to Hlobil's proposal.

Hlobil (2019) defends the taking condition, which he thinks requires the reasoner to take his inference to be good. A 'taking', he agrees with Boghossian, can neither be a belief nor an intuition. He thus makes a new proposal, according to which we should understand the act of reasoning analogously to the act of judging or believing. Believing is taking something to be true, it is the mental act of attaching doxastic force to a proposition. Analogously, Hlobil claims, reasoning is attaching inferential force to a structured collection of propositions, that is, an argument. Attaching inferential force to an argument is an act in which we relate to an argument, and that act "forces our attitudes towards the premises and conclusion into certain patterns" (Hlobil ms: 18). "The act of attaching inferential force can bring one to generate new acts of attaching doxastic force" (Hlobil ms: 20). In other words, it can bring one to believe the conclusion.

There are several problems with this approach. First, Hlobil says that attaching inferential force is not a propositional attitude. But if the act of attaching inferential force is analogous to the act of attaching doxastic force and attaching doxastic force to the proposition that p is holding a belief the content of which is p, then it is unclear why attaching inferential force to an argument is *not* holding a propositional attitude the content of which is the argument. In order to reason, one needs a representation of the argument, which is not the same as a set of representations that together form an argument. To say that reasoning requires relating to and thus representing an argument seems to postulate a first-order meta-propositional belief about

entailment relations between propositions, which, as we have seen, is over-demanding, and either raises regress problems (if understood as a premise) or lacks explanatory power (if not featuring in the reasoning). Another problem of Hlobil's proposal is this. He says that reasoning is the act of attaching inferential force, and that this act can bring us to form new beliefs, it can "bring one to believe the conclusion." But coming to believe the conclusion in a certain way *is* reasoning, not a consequence or effect of it. For these reasons, I remain sceptical of Hlobil's analysis of Boghossian's Taking Condition and conclude that we have no defensible account of it. We need something else.

1.3. Rational Requirements and Rules

According to John Broome (2013, 2014b, 2014a), reasoning is essentially a rule-following activity, and that this is what differentiates reasoning from mere mental causation. So far, he is in agreement with Boghossian. But Broome denies the *Reasons Claim* as well as Boghossians conclusion that rule-following is unanalysable. Broome gives a detailed analysis of what it is to follow a rule in reasoning. Following a rule, according to Broome, is being *intentionally guided* by a rule, which is manifesting a "guiding disposition" (Broome 2014b:632). This is a complex disposition to a) behave in a certain way, and b) for your behaviour to 'seem correct' to you. Seeming correct is an attitude you hold towards your behaviour. Because your behaviour seems correct to you, you count as following a rule, and therefore as actively reasoning.

To take an example, suppose you believe it is Sunday, and that if it is Sunday, the library is closed. Suppose a) you come to believe that the library is closed, and b) this seems correct to you, relative to the rule of modus ponens.[5] You thereby count as following the rule

[5] There is no condition in Broome's account that you form the conclusion-attitude *because* it seems correct to you to do so (which is a central feature of the Taking Condition and Frege's view.

of modus ponens. But how do a) and b) make it the case that you are *doing* something? It could just happen to be the case that you conform to the rule, and that this behaviour - possibly even unconsciously - seems correct to you. If that was so, you would not necessarily do anything. Broome's answer to this objection is that when you follow a rule (that is, when a) and b) are the case), you *do* so *intentionally*. When you follow a rule, he says, "you are always intentionally guided by the rule: you follow a rule because you intend to comply with it. The intention may not be deliberate; often it is a habit" (Broome 2014b: 631). He writes

> "An intention to do something is a disposition of a particular sort to do it. It is a disposition with a guiding role; it guides you towards doing it. An intention to comply with a rule is a guiding disposition to comply. […] To follow a rule is to manifest a disposition to behave in a particular way, and for this behaviour to seem right. Seeming right involves being ready to accept correction. This is just to be guided. So this disposition is nothing other than an intention or a habit. Following a rule manifests a disposition, and this disposition turns out to be an intention to comply with the rule or perhaps a habit of complying. The guidance involved in following a rule is nothing other than intentional guidance by this disposition."[6]

The dialectic is this. To follow a rule is to manifest a complex disposition, i.e. the disposition to comply with a rule and for this to seem correct to you. Next, Broome gives an analysis of intention. An intention is a sort of guiding disposition. When you intend to get a cup of tea now, this intention can guide you to act as you intend. It sets a standard of correctness in the sense that if you do not act as you intend, you are disposed to correct your behaviour. So, an intention is a guiding disposition. Next, Broome argues, the guiding disposition you manifest when following a rule in reasoning qualifies, on his account of intention, as an intention to follow the respective rule. Therefore, you follow the rule because you intend to follow it. It is therefore something you do.

[6] Broome here uses the notion of "seeming right." In later work, he corrects himself and uses the notion of "seeming correct", which is more in line with his overall theory of reasoning.

It is not obvious what exactly the analysis of intentions as guiding dispositions refers to. Searle (1983) distinguishes between, on the one hand, intentions we form through decision making prior to action and, on the other hand, the intentions with which we act when we act intentionally without having deliberately formed an intention prior to our action, which he calls intentions-in-action. According to Broome, the disposition we manifest when following a rule is an intention, so it seems he is interested in intentions-in-action. The intention to follow a rule in reasoning guides your reasoning in the same way in which an intention to act in a particular way guides your acting in that way, he says. This does not imply that we reason for the sake of applying rules. This is analogous to the fact that you can intentionally walk without walking for the sake of walking, but as a means to an end. In short, when you reason, you *intentionally* follow a rule, which is why you count as doing something. Following a rule is manifesting a "guiding disposition", namely the disposition to a) behave in a certain way, and b) for your behaviour to seem correct to you.

Broome's theory of reasoning, also as part of his larger theory of rationality, has generated an extensive discussion.[7] In this section, I add a few points to it. I first assess the concept of a guiding disposition and whether it accounts for the idea that reasoning is a rule-following activity in the way Broome argues. I then investigate how we best understand the rules we follow in reasoning according to Broome's view.

Seeming correct and rule-following

The view is that when you reason, your behaviour must *seem correct relative to a rule*. On one interpretation, this is obviously too demanding. To have this attitude of seeming, it appears, you must somehow compare your behaviour, for example, your forming the belief

[7] See, for example, Way (2013); Boghossian (2016); Pettit (2016); Southwood (2016b), Broome (2016, 2015), McHugh and Way (2016); and Fink (2016).

that the library is closed, with a rule, for example the rule of Modus Ponens. It is hard to see how you could do that in one attitude, without any reasoning. The attitude of seeming appears to require you to identify that your premises satisfy the triggering-conditions of the rule, and to verify that your conclusion is the sort of response the rule demands.[8] Of course, things might just seem correct to you, unqualified. But for them to seem correct to you *relative to a rule* is something else.

However, a close reading of Broome reveals that he has something else in mind when he says reasoning requires the process to seem correct to you relative to a rule. He says the following: Seeming correct is an attitude you hold towards your behaviour. It includes in an essential way the *disposition to correction*. This is a counterfactual disposition to lose your attitude of seeming correct if you were to check your reasoning and came to a different conclusion (Broome 2013: 238).[9] When you follow a rule correctly, you have the steady disposition "for a particular answer to come up frequently, and for it to seem right to you, were you to repeat the process several times" (Broome 2016: 3438).[10] It is this disposition which is the "rule-encoding state" (Broome 2016: 3439) and which explains your behaviour. If what you do seems correct to you, but you lack this steady disposition, you do not follow the rule correctly, and therefore do not reason correctly. Whether you follow a rule correctly depends on your dispositions.[11] Following a rule has an internal standard of correctness.

This conception of the attitude of something *seeming correct to you relative to a rule* is implausible. To say that a process seems correct to you *relative to a rule* is to say that the rule somehow affects whether or not the process seems correct to you. In some sense or the

[8] Boghossian (2016, Section 6) seems to make a similar point.
[9] The disposition to behave in a certain way that seems correct to you relative to a rule can be understood as a "linking belief." Or at least as implying a "linking belief." Such a linking belief is, according to Broome, not essential but necessary for theoretical reasoning (Broome, 2016).
[10] In that sense, the attitude of seeming correct has "double content consisting of the act and the particular rule" (Broome 2012:22).
[11] You would be following a rule, but not following it correctly, which is one form of incorrect reasoning on Broome's account (Broome, 2013: 240).

other, you measure your behaviour against the rule. The problem is that this is not compatible with Broome's conception of what it is to follow a rule. For what seems correct to you on his view is determined by sub-personal processes (Broome 2014a). The rule does not itself set the standard for the checking. Instead, the result of your checking establishes the rule! If you were to repeat the process and check several times, and it would steadily seem correct to you, then you would – because of that fact - count as following the rule that describes the process. Your steady disposition would encode the rule you follow: it describes the process that steadily seems correct to you. It is not the case that the rule sets a standard against which you measure your thinking when you check, but that your hypothetical checking, if it steadily led you to the same conclusion, would define the rule.[12] Your transition does not seem correct to you relative to a rule in the sense that your possessing the rule somehow affects whether or not the process seems correct to you and how you change your response if necessary. Instead, you count as possessing the rule in virtue of having a steady disposition for the transition complying with the rule to seem correct to you. The possession of the rule does not explain what seems correct and not, but rather what seems correct explains which rule you possess. In what sense then, are you guided by the rule you possess?

I believe my worry here is in line with Pettit's criticism of Broome's view of rule-possession and rule-following. Pettit writes that problems arise

[12] Broome thinks both parts of this sentence are true: the rule sets the standard against which you measure your thinking when you check, and your checking defines the rule. He commented: "You do a sum and get the answer 5. When you check, your rule of addition sets the standard against which you check. It's a genuine check, because it may cause you to change your answer. The rule is given by your steady disposition, which may differ from your disposition on the particular occasion that led you to the answer 5." What is crucial here that "it [that is, the rule] may cause you to change your answer." Broome evidently thinks that possessing the rule causes me to change my behaviour. That means that my disposition for some answer to come up *causes* me to change my answer 5 when checking. This causal element seems to be crucial to account for the idea of being guided by a rule. It may overcome the problem I raise here but is not explicitly included in Broome's account. The proposal relies on the idea that I can reason in violation of a rule in the sense that I can follow a rule incorrectly. This is my next topic.

> "if the proposal were that for a response to be right is for it to be such that it would steadily seem right: that is, seem right over many variations on the actual situation. It would be circular, after all, to characterize being right in terms of seeming right, even seeming right in a steady fashion. It would leave us with the question: How does a response seem when it seems right? And the answer cannot be: it seems that it would seem right in a steady fashion. That answer would leave us in the dark as to what property it is that a response seems to have when it seems right to you—when you see it as right" (Pettit 2016: 3382 -3).

The problem here is not with the nature of the attitude of seeming. Things seem right to us all the time. The problem lies with the idea that something seems right *relative to a rule*. We can accept that reasoning requires the process to seem correct to you, but not that it seems correct to you *relative to a rule*, nor that the manifestation of this obscure attitude constitutes rule-following.

Incorrectly following rules

This issue arises also when we look at the account from a different angle and consider incorrect reasoning. Broome argues he accounts for the possibility of incorrectly following a rule. According to him, you follow a rule incorrectly when you lack the steady disposition for the process to seem correct to you. But, if you lack that disposition, there is no rule-encoding state. In what sense then, do you follow a rule at all? You cannot follow a rule that is not there. So, unless you follow a rule correctly, you are not following any rule.

Broome thinks my analysis of his view here is mistaken. What he did not say in *Rationality through Reasoning* but meant to say is this: you follow a rule incorrectly when what you do seems correct to you, and you lack the steady disposition for the process to seem correct to you, but you have a steady disposition for a different answer to seem correct. Then

there is a rule, and that is the rule you follow. When you incorrectly follow rule R, what you do seems correct to you relative to rule R, even though it is not correct relative to rule R.

But, I wonder, why do you count as following rule R? Just because you have a steady disposition for R-compliant transitions to seem correct to you? Suppose you reason from believing p and believing if p then q to believing q* by incorrectly following a correct rule. According to Broome's definition, your transition then seems correct to you, but you don't have a steady disposition for it to seem correct to you. So, you don't follow the rule 'from p and if p then q derive q*'. Which rule do you follow?

Suppose you check your answer again and it doesn't seem correct to you. You change your answer to q. Now it seems correct to you again. Does that imply you checked *relative to the rule of modus ponens*? Do you follow the rule of modus ponens both in forming the belief that q* and in forming the belief that q? I don't see why you would count as following the rule of modus ponens in the transition to q* simply because, if you were to check it, you would reason to believing p in compliance with the modus ponens rule. You may have the steady disposition for modus ponens transitions to seem correct to you, but this disposition may have played no role in either the transition to believing p, let alone the transition to believing p*. (As far as the transition to q is concerned, this is the point made above.) Even if we grant that, upon checking your transition to q*, you follow the modus ponens rule correctly, this does not entail that your possession of that rule had anything to do with your transition to q*. The problem is not that we cannot determine which rule you followed when forming the belief that q*. The problem is that in order to count as reasoning, you must be following some rule.[13] But for any rule you may possess, it is unclear what role it plays, let alone how it guides an incorrect transition. In general, the problem is that according to

[13] Broome writes that on his account, it may be indeterminate which rule you follow *when you follow a rule correctly*. I agree that that is not a problem for his view (Broome 2013:241-2).

Broome's view, you may count as incorrectly following rule R, even if the state that encodes R, the steady disposition for transitions that comply with R to seem correct to you, may play no role in the incorrect transition.

In defence of Brome's account, one might want to argue that you follow the rule R incorrectly in a transition T only if (i) T seems correct to you relative to R, but you don't have a steady disposition for this to seem correct to you *and* (ii) you have the steady disposition for transitions that comply with R to seem correct to you *and* (iii) where you to check T, this steady disposition *would cause* you to lose the attitude of seeming correct in T and to form the response which R requires. Condition (iii) overcomes my above objection that the rule R doesn't guide your reasoning. The attitude of seeming correct relative to rule R, in Broome's framework, includes the disposition for correction. That disposition for correction is regulated by the disposition for transitions in line with R to seem correct to you. Unfortunately, however, this proposal doesn't get us far. It says that a transition seems correct to you *relative to rule R* if and only if the disposition to correct which is essential to that attitude is regulated by the *steady disposition for R-transitions to seem correct to you*. This is hopelessly circular and brings us back to exactly Pettit's and my worries which I presented above. We cannot define seeming correct in terms of being correct and vice versa. It all comes down to this: we cannot make sense of incorrectly following a rule because we cannot make sense of following a rule.

This leaves me sceptical about the project of analysing rule-following in terms of the attitude of seeming correct. Put differently, it leaves me sceptical that the attitude of seeming correct, assuming it is essential to reasoning, is best understood as grounding rule-following.

The rules of reasoning

Suppose the above problems can be answered. What would the rules of reasoning be? Though Broome says that the rules of reasoning may be particular or general, he only discusses rules that are general. They are schematic, abstract constructs. For example, the modus ponens rule of reasoning he proposes reads:

> "From <p; belief> and <if p then q; belief> to derive <q; belief>" (Broome 2013: 252).[14]

It is difficult to see how you could follow a rule like this, under Broome's conception of rule-following. Consider again the example from above. For simplicity, let us only consider only correct reasoning. You reason "It is Sunday. If it is Sunday, then the library is closed. So, the library is closed." The process seems correct to you relative to a rule, and you have a steady disposition for this process to seem correct to you. This disposition encodes the rule you follow. But how could that disposition encode the abstract modus ponens rule? What seems correct to you is the particular instance of reasoning, which complies with the particular rule

From <It is Sunday; belief> and <If it is Sunday, then the library is closed; belief> derive <The library is closed; belief>.

Even if you have a much more general steady disposition, which can be described by the schema "from <p; belief> and <if p then q; belief> to derive <q; belief>", that disposition might play no role in the reasoning. What seems relevant first and foremost is your steady disposition for your particular transition to seem correct. Maybe other instances of modus ponens also steadily seem correct to you, so that in some sense, you have a steady disposition for modus ponens, in general, to seem correct to you. But we can attribute this general

[14] Boghossian (2014) also proposes that reasoning is governed by general rules, not particular ones.

disposition to you because you have the particular dispositions, not vice versa. The fact that other instances of modus ponens reasoning also seem correct to you doesn't play any role in the instance of reasoning we are looking at. So even if we can attribute to you the general disposition, it is not clear in what sense you are guided by the respective general rule. And we do not want to say that you have a representation of the abstract rule in mind when you are guided by it. To say that you must recognise your behaviour instantiating that rule would lead to Boghossian's 'Inference Problem'.

Here is one more worry: According to Broome, when you reason, what you do seems correct to you relative to a rule, and because of that, you count as following that rule. But what guarantees that this is a rule of reasoning? For all we know, it may be any sort of rule that tells you what to believe or intend given what you believe or intend. Consider the rule "derive <this year's harvest will be rich; belief> from <April had a lot of rain; belief>", or the rule "derive <be forgiving and friendly; intention> from <the hurrying stranger pushed me out of their way; belief">" or the rule "derive: <he is innocent; belief> from <he is not proven guilty; belief>." None of these seem to be rules of reasoning but pragmatic or moral rules. Even if all of the above problems can be solved, and we can account for the idea that you are guided by rules, nothing guarantees that these are rules of reasoning. Even if we can make sense of the idea that what seems correct to you seems correct to you *relative to a rule*, this may not be a rule of reasoning. This is a problem, because reasoning is supposed to be following rules of reasoning, not just some rule.

To conclude: Broome argues that in reasoning, the process seems correct to you, relative to a rule. Because of this attitude of seeming correct, you count as following a rule. It makes the process count as something you do, an act of yours. I pointed out issues with the relational

construal of this attitude of seeming correct, and Broome's argument that this attitude grounds rule-following. I have then argued that the account does not accommodate the possibility of incorrectly following a rule in the way Broome claims. My last objection concerned the rules we allegedly follow in reasoning, and whether we can make sense of the idea that they are schematic, general rules. Overall, Broome's account does not satisfactorily explain the activity of reasoning in terms of rule-following. I next turn to Philip Pettit's different, though arguably similar account of reasoning.[15]

1.4. Meta-Propositional Beliefs

Pettit argues that what is essential to reasoning is the involvement of meta-propositional beliefs. Meta-propositional beliefs are beliefs in meta-propositions: propositions which attribute relations and properties to propositions. A meta-propositional belief, for example, is the belief that [the proposition that p is true], or the belief that [the truth of the proposition that p and the truth of the proposition that if p then q entail the truth of the proposition that q]. Meta-propositional beliefs are thus fundamentally different from beliefs in ordinary propositions in which only concrete objects and their properties figure. They constitute "propositional ascent." The ability to reason is the ability to ask oneself and form beliefs about certain propositional properties and relations, which robots and animals lack, and may

[15] Pettit (2016) argues that Broome's theory is equivalent to his own view in all relevant respects. Broome (2016, forthcoming) agrees that their views are similar, but exposes a difference. While Pettit believes that a linking-belief is essential to reasoning, Broome does not. A linking-belief, for Broome, is the first-order belief that the premises imply the conclusion. Pettit calls linking-beliefs meta-propositional beliefs. (He thinks Broome's linking-beliefs are meta-propositional in character. Broome denies this, arguing that his linking-beliefs as well as Pettit's meta-propositional beliefs are in fact just beliefs in conditionals, which he does not take to be meta-propositional.) According to Broome, a linking-belief is *not* essential to reasoning even though it is necessary for theoretical reasoning. It is an implication of another necessary condition for reasoning. Practical reasoning features no linking-beliefs. Again, Pettit holds that a linking-belief or meta-propositional belief is essential to reasoning. Independent of this disagreement, one finding is surprising: Pettit thinks his account is essentially the same as Broome's. But Pettit defends what I call the *Reasons Claim*, the claim the reasoning involves recognizing and responding to normative reasons (Pettit 2007a:238). Broome explicitly denies this.

be unique to human beings (Pettit 2007b). When we reason in accordance with modus ponens, for example, we

> "are able, however implicitly, to register that propositions of the form 'if p, then q' and 'p' entail the truth of the conclusion 'q'. We will form a belief in that entailment, a sensitivity to the connection involved. Such a belief is a belief about the propositions involved, not a belief in those propositions. It is a higher-order or meta-propositional belief: a belief in the meta-proposition" 'p' and 'if p, then q' entail 'q'" "(Pettit 2007a).

We possess the ability to form meta-propositional beliefs, according to Pettit, because "we have language and can use sentences as ways of making proposition into objects of attention" (Pettit 2007b: 500). As opposed to lower animals and robots, we are able to intentionally "check and channel" the processes through which we form beliefs and intentions "by forming the meta-propositional belief" and registering the relationship between the premises and conclusions of our reasoning (Pettit 2007a: 234). When I conclude q on the basis of considering p and if p then q to believing q, I "register their relationship." I count as reasoning in virtue of having the meta-propositional belief that p and the belief that if p then q entail q.

As presented above, Boghossian has argued that a meta-propositional belief cannot be necessary because it would render reasoning too sophisticated an activity. Pettit answers to this objection. He says the meta-propositional belief (the linking-belief) is held not *in sensu composito*, but *in sensu divisio*. We don't have to hold the belief in "a unified, once-for-all-cases sense." Instead, he writes

> "We may believe it just insofar as we are disposed on consciously assenting to any instance of the antecedent to assent in the same way to the consequent of the conditional. Specifically, we are disposed on assenting to 'A' to think 'so t', where 'the word ''so'' or some cognate' (Pettit 2007b:500) serves to mark the

> connection with the premise, 'A'. And surely most of us who are untrained in the theory of reasoning do hold our linking beliefs [i.e. meta-propositional beliefs] in that case-by-case way; we do not have the skill or the inclination to spell things out further" (Pettit 2016:3379).

On this view, you count as holding the meta-propositional belief that 'A entails t', for example, in virtue of being disposed to form the belief that t whenever you consider A, somehow seeing t as arising from A. But, one may ask, if the meta-propositional belief is merely a disposition, how can its presence make it the case that you count as reasoning? (Also: what is it to see t as arising from A?) In answering the over-intellectualising objection, Pettit seems to have gone too far: on the dispositional understanding of the meta-propositional belief he proposes, his account doesn't explain why reasoning is something you *do*. It is not clear how the meta-propositional belief gives me "a degree of personal control" over my process; a mere disposition doesn't "transcendent the non-intentional" (Pettit 2007b).

There is more that warrants this conclusion. Pettit says the meta-propositional belief is the result of a non-intentional process. He writes: "I must be able to rely on the attentional activity to generate a correct meta-propositional belief and, second, I must be able to rely on that meta-propositional belief having the required effect. Reasoning does not work in parallel to non-intentional rational processes, seeking the same end by different means; it exploits rational processes in a novel way, searching out extra inputs to impose as checks on how the process is operating." (Pettit 2007b:502). But if I just happen to have one or the other meta-propositional belief, how can having that belief make me count as actively reasoning in congruence with it, rather than just behaving in congruence with it? How can a *belief*, that I don't acquire intentionally impute intentionality to the thought process it is about? Pettit states that in reasoning, we "intentionally pursue the satisfaction of rational desiderata" (503). But

we certainly don't have to have the intention to pursue the satisfaction of rational desiderata when we reason. Pettit does seem to explain why reasoning is an intentional activity.

The final point on Pettit's account I would like to make here is that he, like Boghossian, just claims that his view accounts for the idea that reasoning is treating things as and responding to normative reasons. According to Pettit, "we human beings take account of reasons, treating them as reasons, when we are moved not just by those [premise-]beliefs, but also by the meta-propositional belief that 'p' and 'if p, then q' entail 'q'; that we treat them as reasons shows up in the use of the 'so q' or the 'therefore q'", says Pettit (2007a:236). (This also raises the question of *how* we are moved by the meta-propositional belief.) But again, since the meta-propositional belief has no normative content, it is not obvious in what sense the account accommodates the claim that in reasoning, we manifest a sensitivity to normative reasons. We seem to manifest a disposition to form beliefs in accordance with supposed entailment relations between propositions. To complicate matters even more, Pettit says that meta-propositional beliefs are reason-providing beliefs, giving us reasons for the conclusion. But if reasoning is responding to reasons, the reasons we respond to are plausibly provided by the truth of the premises, not the meta-propositional belief (especially not if the latter is a disposition). In sum, Pettit's account does not seem to explain why reasoning is active, nor why it is a way of responding to normative reasons.

1.5. Endorsing Transitions of Thought

Neither Boghossian's, nor Broome's nor Pettit's account of the nature of reasoning is fully convincing. But there is a common thread in them: all argue that in some way or the other, what makes a thought process count as reasoning is that the agent holds a certain attitude towards the process (that is, their behaviour) or towards the contents of the attitudes

constituting the process. The idea is that a thought process counts as reasoning only if you endorse it in some way (Broome 2013: 228). They say that in reasoning,

Boghossian	you *take* the premises to *support* the conclusion and drawing the conclusion because of that
Brome	your behaviour *seems* correct to you, relative to a rule
Pettit	you hold *a meta-propositional belief* that links the premises to the conclusion.

In the same fashion, Frege and Hlobil argue that

Frege (1979)	you *are cognisant of* the premises being justifying reasons

More views could be added to this list. It has been argued that when you reason

Anscombe (1974:2)	you *see* that the conclusion follows from the premise[16]
McDowell (1994:13)	you *have a standing reflection* "about the credentials of the putatively rational linkages that govern it"
Grice (2001:14)	you *think* of yourself as conforming to valid patters of inference
Audi (2004:126)	you "are guided by *a sense of* some relation of support between the former – call it the premise set – and the latter, the conclusion"
Raz (2015:159)	you "*realise* that some items (premises) support others (conclusions)"
McHugh and Way (2018:180)	you "*manifest a sensitivity* to the fact that a rule is fittingness-preserving"

[16] Anscombe also writes in her *Practical Reasoning* (1974:2) that one conceives the second proposition (i.e. the conclusion) "under the aspect *then* in relation to the first", which seems to be different.

Valaris (2017:194)	you *recognise* that the premises "constrain the way things might be so as to guarantee" that the conclusion is true.
Siegel (2017)	you have a "reckogning state" in which you reckon the premises to support the conclusion
Kauppinen (2018: 419)	you "have a positive attitude towards performing the act of forming" the conclusion-attitude
Hlobil	*attach inferential force* to a structured collection of propositions

The general idea is that as reasoners, we are conscious of and can think about the activities of our minds. We are self-reflective. As Korsgaard (2009) puts it, we have "a capacity to distance ourselves from them, and to call them into question." If we accept the assumption that all the above views appear to share, namely that a thought process counts as an instance of reasoning in virtue of the agent's holding a certain attitude towards it or that which it is about, then one desideratum for a theory of the nature of reasoning is to

D1 explain the content and nature of the attitude that is essential to and grounds reasoning.

As we have seen in the discussion of Boghossian's 'taking', Broome's 'seeming correct' as well as Pettit's 'meta-propositional belief', a second desideratum for a theory of reasoning is then that it

D2 shows how the presence of that attitude makes the process count as *active* and explains its role in the reasoning process.

Furthermore, the objections levied against the views discussed above show that a theory of the nature of reasoning should

D3 not render reasoning conceptually over-demanding, that is, not require reasoners to have complex concepts like that of a normative reason, belief, entailment relation, etc.,

D1 to **D3** are largely uncontroversial. I have raised objections to the views presented above on the grounds that they fail to meet these desiderata and referred to others who have argued in similar ways.

1.6. The Fittingness View

Besides Broome's *Rule-Following View* and the various versions of the *Reasons View*, the third view dominating the recent debate is the *Fittingness View* proposed by Jonathan Way and Conor McHugh (2016a; 2018). According to this view, reasoning is a goodness-fixing kind, more precisely, a functional kind. Following Thomson's (2008)'s theory of goodness-fixing kinds, McHugh and Way argue that reasoning has a constitutive aim or function in the strong sense that only activities regulated by this aim count as reasoning. The aim of reasoning, McHugh and Way argue, is to get fitting attitudes. Without offering a reductive account of fittingness, the view relies on the idea that fitting attitudes are, for example beliefs that are true and intentions that are permissible.

The aim of reasoning also sets a standard of evaluation: good reasoning is reasoning that reaches this aim, bad reasoning is reasoning that is regulated by that aim without serving it well. More precisely, McHugh and Way argue that good reasoning is fittingness preserving:

Put roughly, if the premise-attitudes are fitting, then, because of that, all other things being equal, the conclusion-attitude is fitting, too. I discuss this standard of evaluation in chapter **5**.[17] Let us accept it for now, for our focus here is on what reasoning *is*, not on what good reasoning is. Let us also assume that reasoning is a goodness-fixing kind (which I deny in chapter **5**), and that the aim of reasoning is to get fitting attitudes. The question then is: what exactly does "being guided by the aim of getting fitting attitudes" consist in? Does it explain why reasoning is active?

McHugh and Way say the following: You are *guided by* the aim "through the rules of reasoning that you follow." McHugh and Way agree with Broome that reasoning is a rule-following activity, but contrary to him, they believe that what is essential to reasoning is not that it is rule-following, but that it is guided by an aim. Following rules of reasoning, according to them, is to manifest your (imperfect) sensitivity to what will serve the aim of getting fitting attitudes. So, to be guided by the aim of getting fitting attitudes is to manifest your sensitivity to what will serve that aim. You can be sensitive to the fittingness-preservation of a certain transition without representing the transition as fittingness-preserving. You "thereby count as aiming at fittingness" (McHugh and Way 2018:180).

If you follow a pattern that is fittingness-preserving, you thereby manifest *your sensitivity to the fact* that it is fittingness-preserving. If you follow a rule that is not fittingness-preserving (that is, a bad rule), "your sensitivity to what will preserve fittingness is imperfect" (p.182), i.e. you have a bad reasoning disposition. Our dispositions to reason, whether good or bad, are governed by second-order dispositions that "suggest an overarching concern for fittingness" (p.183). These are dispositions regarding which rule to follow, when two rules

[17] I here use good reasoning synonymously to correct reasoning, and bad reasoning synonymously to incorrect reasoning. I argue in chapter 5 that there is a difference between the goodness and the correctness of reasoning, but this does not concern us here.

conflict, for example. Also, when you follow and comply with a bad rule and recognise that the transition is not fittingness preserving, you stop following the rule because of your higher-order disposition to engage in fittingness-preserving thought.

In sum, McHugh and Way propose a dispositional account of rule-following, and thus a dispositional account of reasoning. To reason is to be guided by an aim, which is to follow a rule, which is to manifest "one's sensitivity to the fact that a rule is fittingness-preserving", and this is to manifest a disposition. I wish to make four comments on this.

Rule-Following

First, the view is at its core, a rule-following view of reasoning. Though McHugh and Way claim that rule-following is not essential to reasoning on their view, it is certainly necessary. What is essential is that you are guided by the aim of getting fitting attitudes. But if what it *is* to be guided by the aim *is* to follow certain rules, then rule-following is not a mere consequence of an essential feature of reasoning, but part of that in virtue of which you count as reasoning when you reason.

The activity of reasoning

Second, McHugh and Way do not address, let alone explicitly answer the question of what makes reasoning an activity rather than a mere passive thought process.

They say that reasoning is *aim-directed*: it has a constitutive aim in the sense that only activities regulated by this aim count as reasoning. But it is not exclusive to activities to have a constitutive aim (indeed, most kinds which obviously have an inherent function or aim are not activities, such as knives or hearts), Indeed, we might think that there are mental activities which constitutively aim at getting fitting attitudes, such as meditation, but which neither are, nor involve instances of reasoning.

McHugh and Way's appeal to *rule-following* offers little support of their claim that reasoning is an activity. For they adopt a dispositional account of rule-following without addressing the well-known problem Boghossian raised: you can manifest a disposition in reasoning without manifesting your agency.

They might argue that reasoning is active because in reasoning, you manifest *your sensitivity* for fittingness-preservation. But this again is the manifestation of a disposition, and as such doesn't suffice to make your transition count as active. Also, being "sensitive to the fittingness-preserving character of a given rule" cannot be what it *is* to aim at fittingness because you can exhibit the latter without the former, namely when you reason badly: in those cases you are guided by the aim of getting fitting attitudes (otherwise you wouldn't be reasoning), but you cannot manifest a sensitivity for the fittingness-preserving character of the transition, because it is not fittingness-preserving. I conclude that McHugh and Way do not explain what makes reasoning an activity, and thus fail to meet **D2.**

Incorrect or bad reasoning

My third comment on the *Fittingness View* concerns the possibility of incorrect or bad reasoning. McHugh and Way claim that their view allows for this possibility: you reason badly if you incorrectly follow a rule, or if you correctly follow a bad rule.

Regarding the former case, to follow a (good or bad) rule *incorrectly*, they say, is to misapply the rule: the transition between your attitudes does not accord with the rule but the transition nonetheless triggers "your disposition to follow it [the rule]" (p.181). The last part of this sentence is puzzling. We may say that the disposition to follow a certain rule is triggered by some state of mind, so that you reason in a certain way, which is not in compliance with the rule, but which counts as following that rule. But *the disposition to follow a rule triggers the transition*, not vice versa. In virtue of what do you count as following the

rule, when you follow it incorrectly? How can you manifest a disposition to follow a certain rule incorrectly? I have argued above that these questions arise for Broome and that his answers are unsatisfactory. McHugh and Way do not give any answers.

There is a second way of reasoning incorrectly: when you correctly follow a bad rule. You are then still guided by the aim of getting fitting attitudes, but your sensitivity to what will serve this aim is imperfect, or so the story goes (p.181). But what is the sensitivity you manifest? *In virtue of what* are you still guided by the aim of getting fitting attitudes? We need an account of how the aim figures in your reasoning.

One could say that when you follow a bad rule, you are not sensitive to the fittingness-preserving character or the rule, for it is not fittingness-preserving, but *it seems to you* that it is. You manifest a sensitivity to fittingness-preservation because a certain transition *seems* fittingness-preserving to you. In good reasoning, what seems fittingness preserving is fittingness preserving. Since fittingness-preserving rules play the analogous role to correct rules in Broome's theory, we have now arrived at a rudimentary replication of Broome's account of what reasoning and rule-following is. We have just substituted Broome's idea of what a correct rule is with McHugh and Way's account of what a good rule is. But this is inessential to the account of what reasoning is. McHugh and Way, it appears, cannot deny that reasoning is rule-following, which essentially involves that what you do seems correct to you relative to a rule (or it seems to you to conform to a good rule). In other words, their account of reasoning in terms of guidance by an aim collapses into Broome's rule-following account. The views differ in what makes a rule good or correct, but again, this is a matter of evaluating reasoning, not a matter of the nature of reasoning. So again, the view cannot account for the possibility of incorrectly following a rule (which I objected to Broome's view, too.)

Higher-order dispositions

My fourth and final comment strengthens this conclusion. McHugh and Way argue that higher-order dispositions regulate your reasoning dispositions. Higher-order dispositions "suggest an overarching concern for fittingness." The problem is that it is mysterious what higher-order dispositions are, and how they regulate your reasoning. Consider an example they give:

> (Promise Rule) From believing you promised to ϕ, move to intending to ϕ.
>
> […] You might believe that you promised to meet Jenny for lunch, but also believe that by doing so you would leave someone to die. In that case, you will refrain from applying the Promise Rule and coming to intend to meet Jenny for lunch. Why? Because these are cases where following these rules will lead to unfitting attitudes. (McHugh and Way 2018: 175)

According to McHugh and Way, you do not use the Promise Rule here because it will lead to unfitting attitudes. Does that mean the rule would be a bad rule in this instance? Or would it simply be inappropriate to follow that rule, despite it being a good rule? In any case, you must have some kind of epistemic access to the fact that following the rule would lead to an unfitting attitude for this to direct your reasoning. Somehow you anticipate that following the promise rule would lead to the intention to ϕ, and that this intention would be unfitting. Because of that, you do not follow the Promise Rule. How could you do that without reasoning? This seems to be a case of meta-reasoning: reasoning about how to reason. Also, there are cases where following a certain rule would lead to unfitting attitudes, but no higher-order disposition could stop you from following the rule. Suppose you falsely believe that it is the first Saturday of the month, but you correctly believe that if and only if it is the first Saturday of the month, the dance class takes place. Following the modus ponens rule would lead to the unfitting (false) belief that the dance class takes place. How could any higher-

order disposition stop you from following that rule? It is simply not transparent to you that following modus ponens would lead to an unfitting belief. Something is going on in McHugh and Way's promising example, and it needs to be worked out what it is in order to explain the alleged role of higher-order dispositions in reasoning.

A straightforward analysis of the promise case is this. Reasoning from the belief that you promised to ϕ, *and* the belief that if you were to ϕ, you would leave someone to die, to the intention to ϕ does not seem correct to you, even if reasoning in accordance with the Promise Rule (i.e. reasoning from believing that you promised to ϕ to intending to ϕ) does seem correct to you. It is not that the Promise-Rule is incorrect, it is just does not apply in that case. The triggering conditions of the rule do not match the premises-attitudes. This is, again, in line with Broome's account.

Here is another of McHugh and Way's (2018:171,183) examples. Consider the

"(Means-End Rule) From intending to E, and believing in order to E you must M, move to intending to M."

They say that

> […] if you intend to *E* and think that in order to *E* you must *M*, but *M*-ing strikes you as a very bad idea, you will probably drop your intention to E rather than reasoning to the intention to M. Why […] do you not apply the […] Means-End rule? Because you see that doing so will lead you astray. You see that something has gone wrong with your existing attitudes, and that the way to make your attitudes fitting is to revise these attitudes rather than to follow where they lead.

According to this quote, it is your *seeing* that the rule will lead you astray that prevents you from following it. The formulation is ambiguous between two interpretations: Does the *rule* seem incorrect to you, or does it seem incorrect to follow the rule, despite it being a good or correct rule? It appears as though only the latter is true. Following the Means-End Rule appears incorrect in this case because the triggering conditions are not met. You do not only

intend to E and believe that in order to E you must M, but also believe that M-ing is a very bad idea. In other words, the rule simply doesn't apply. But this does not imply that you take the rule to be incorrect. The reasoning you do wouldn't seem correct to you relative to this rule, but it may seem correct to you relative to the rule you do follow, which leads you to revise your present attitudes. Again, we can make sense of the example without appealing to second-order dispositions, simply by appealing to the concepts constituting Broome's account. We then of course have to answer the objections I raised to this. McHugh and Way say that your seeing "that something has gone wrong with your existing attitudes" steers your reasoning in the case at hand. This is best understood as a belief about your attitudes (or their relation) which affects the direction in which you reason, and again, no appeal to second-order dispositions is necessary. As Broome says, the direction of your reasoning is, at least partly, determined by the robustness of your beliefs (Broome 2013:245).[18] I do not argue that we should accept Broome's theory. But McHugh and Way's theory, on its most plausible interpretation, collapses into Broome's theory in all relevant aspects.

1.7. Conclusion and Outlook

In this chapter, I have presented and critiqued versions of the view that reasoning is essentially a rule-following activity. I have argued that Boghossian's, Pettit's and Broome's account, as well as McHugh and Way's theory, are problematic in various aspects. But I have not denied in general the

Rule-Following Claim Reasoning is a rule-following activity.

[18] Note that the Promise-Rule example is different in this respect. Here, it is not a belief about your attitudes, but a belief about the world that affects what seems correct to you.

The thought that reasoning is rule-following has great intuitive appeal. Who could deny that there are rules of reasoning? It makes sense to say that we can follow modus ponens in theoretical reasoning, and that some rules of reasoning are correct while others are not. At least some reasoning, we might think with Broome, is like speaking a language, which is guided by the rules of grammar. And some reasoning is like doing arithmetic, which is also guided by well-defined, formal rules of mathematics. A lot of reasoning is guided by the inference rules of logic. There are also rules in the practical realm. A correct rule of practical reasoning, we might say, is the instrumental rule to take the means you believe to be necessary to achieve your ends. Whether or not your belief about the means is true, whether or not your ends are worth aiming for, if you follow the instrumental rule, you are reasoning correctly.

But another characterisation of reasoning comes equally naturally, summarised by what I called the

Reasons Claim	Reasoning is a way of responding to normative reasons for belief or action.

Despite his commitment to the *Rule-Following Claim*, Boghossian's starting point, for example, is the thought that reasoning involves the reasoner taking his premises to be reasons for his conclusion. Indeed, the considerations we adduce in reasoning are commonly thought of as reasons. In modus ponens reasoning, my reasons for believing q are p and if p then q (Pettit 2007a:236). When the detective reasons to a belief about the guilty party, he responds to the evidence he has, that is, to what he has reason to believe. And when we deliberate about what to do, we attend to the reasons we have for doing one thing rather than another and eventually form an intention on the basis of what those reasons on balance demand. The capacity to reason just is the capacity to track what we have reason to do and to believe (Pettit 2007a:240).

There seem to be at least as much initial intuitive support for the *Reasons Claim* as for the *Rule-Following Claim*. Depending on the kind of reasoning we consider, the one or the other characterisation seems more plausible. When we contemplate which job offer to accept, we seem to be weighing and responding to reasons for action. When we try to find the guilty party in a murder trial, we weigh and respond to reasons for belief, given by the evidence. What would the rules that guide us in these situations even be? In contrast, when we calculate how much the investment will pay back, or when we do logic or our tax statements, then it seems we follow rules, rather than reasons. So, should a theory of reasoning accommodate the *Rule-Following Claim*, the *Reasons Claim* or both? Given the strong intuitive support for both, a fourth plausible desideratum for a theory of reasoning would be to

D4 either accommodate or explain the mistake behind the *Rule-Following Claim* and the *Reasons Claim*.

In the next chapter, I present and assess views that put the *Reasons Claim* at their core: they define the nature of reasoning in terms of responding to normative reasons. We will then encounter a fundamental difference between the *Rule-Following View* and the *Reasons View*: the latter postulates, while the former denies a close relationship between reasoning and normativity.

2. Reasons View

2.1. Introduction

Many moral philosophers (and epistemologists) assume a close connection between the normative reasons we have to believe and do things and the activity of reasoning through which we form beliefs and intentions. Some even argue (or presume) that we can define one in terms of the other. Especially prevalent is the view that we can understand reasoning in terms of normative reasons.[1] Many accept the

Reasons Claim: Reasoning is a way of responding to normative reasons for belief or action.

The claim has great intuitive appeal. It is natural to say that when we reason about what to believe or what to do, the considerations we attend to are reasons to believe this rather than that, or reasons to do this rather than that. Reasoning just is the weighing of and responding to the reasons we have to believe or do certain things. Nevertheless, the *Reasons Claim* is not self-evident.

In this chapter, I look at accounts of reasoning that attempt to defend the *Reasons Claim*. They are versions of the so-called *Reasons View* of reasoning. In **section 2.2**, I first illustrate the widespread acceptance of the *Reason-View*. I then present and critique what I take to be the four most promising versions of the *Reasons View,* namely those developed by Paul Grice, Joseph Raz, Antti Kauppinnen and Jonathan Dancy. I consider each in turn in **sections 2.3 –**

[1] Reversely, some argue that we can understand normative reasons in terms of correct reasoning. Setiya (2014) defend this view for practical reasons, Alex Gregory (2016) for theoretical reasons and Jonathan Way (2017) and Asarnow (2017) for both, practical and theoretical reasons. In my manuscript 'Are reasons essentially premises?' I give detailed arguments against all these versions of this so-called *Reasoning View* of reasons, which I cannot re-state here. My view of reasoning, as we shall see, can be seen as another argument against their approach, for it implies that it is false.

2.6, respectively. I am here primarily concerned with their views on the nature of reasoning, not the standard of correctness of reasoning. None of them, I argue, gives a fully satisfactory answer to the question of what reasoning is. **Section 2.7** summarises and considers the *Reasons View* in the context of the larger debate on what reasoning is.

Section 2.8 compares the *Reasons View* and the *Rule-Following View*, which I considered in the previous chapter, with respect to what they imply about the relation between reasoning and normativity. The conclusion of this chapter, drawn in **section 2.9**, is that the *Reasons View* of reasoning is surprisingly under-theorised,[2] and that we should not attempt to give an account of the nature of reasoning in terms of reasons.

2.2. The Popularity of the *Reasons View*

Many authors assume (with varying relevance for their respective theories) a close relation between reasoning and normative reasons that can be seen as a commitment to the *Reasons View*. They do so without providing any substantive argument for the view, mostly because their writings do not target the question of what reasoning is, but other issues concerning the nature of normativity or the will. The following list of excerpts of writings on reasoning, rationality, reasons, moral motivation and related topics show the popularity of the view:

- John Pollock (1987:482) assumes that the constituents we use in reasoning are prima facie reasons and defeaters, and defines defeasible reasoning as reasoning in which "the premises taken by themselves may justify us in accepting the conclusion, but when additional information is added, that conclusion may no longer be justified" (Pollock 1987: 481).

[2] Kauppinen (2018) and McHugh and Way (2018) state this, too.

- Thomas Scanlon (1998: 55-56) argues that the features of our practical thinking "are better represented in terms of a framework of reasons than in terms of competing desires, for in deciding what to do, one considers "reasons in the standard normative sense." When you reason, you treat things as reasons for action and belief (Scanlon 2007).
- Christine Korsgaard (2009:30) argues that the faculty of reason is identified as the "active dimension of the mind" and the "ability to recognise and respond to reasons." The rational principles that describe the activity of reason are those that are applied by the person who is actively thinking, i.e. who is "trying to work out what to believe or to do." The principles "pick out the substantive considerations that we then regard as reasons" (Korsgaard 2009: 30).
- Christian Piller (2001:196) takes for granted that a "theory about reasons and their relations also provides a theory of correct reasoning. The relation is simple enough: a person reasons correctly if she follows the path of reasons (in the right way)."
- Robert Audi (2004: 119, 126) argues that "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." Reasoning is "guided by a sense of some relation of support between" the premise set and the conclusion.
- Mark Schroeder, when theorising about what counts as an agent's reasons for action, assumes that "when Ryan is reasoning well, the kinds of things about which he should be thinking are his reasons" (2007: 26). He thereby imposes the so-called *Deliberative Constraint* on reasons: something is a reason for an agent only if it can play a certain role in his reasoning. He assumes that (correct) practical reasoning is defined in terms of reasons, writing that "there are many explanations of why it may be that Ryan can

reason well without thinking about his reasons, or by thinking about only part of his reasons, but I still think that this is an odd thing to say" (Schroeder 2007: 28).

- Peter Hacker (2007) locates the feature that distinguishes humans from other animals in the ability to reason, which he takes to be "the ability to deliberate (weigh reasons)", and the ability to "form beliefs or resolve to act on the basis of the balance of reasons" (p. 203). He claims: "Human beings can reason from given premises to theoretical or practical conclusions. We take such-and-such to be a reason for things that things are thus-and so" (p. 3).
- Maria Alvarez (2010:362; 2018) in analysing the nature of practical reasoning, assumes that "what we will find among the premises, is statements that describe the thing wanted and give the agent's reason for wanting to do that thing (and, needless to say, many of these statements may be implicit in one's reasoning."
- Simon Blackburn (2010:7), despite his expressivist views, believes responding to reasons is essential to the activity of reasoning, saying that there is a "difference between taking one thing as a reason for another, and being subject to some kind of free association in which one thing leads to another, but not by a process of reasoning."
- Kurt Sylvan (2016:3.3), in describing what it is for some belief to be based on another, assumes that "Inferring p from q entails treating q as a normative reason to believe p."
- Nishi Shah (2006: 486), in arguing for the view that only evidence gives us reasons for belief, assumes that "reasoning is the process in which agents recognize reasons, and then ϕ on the basis of this recognition."

While none of the quoted authors attempt to defend the *Reasons Claim* in detail, Paul Grice, Joseph Raz, Antti Kauppinen and Jonathan Dancy develop systematic accounts of reasoning in terms of normative reasons. I look at each of their proposals in turn.

2.3. Paul Grice

Paul Grice (2001) starts his inquiry of the concept or group of concepts captured by the word 'reason' with the supposition that "reason is, of its nature, the faculty which is manifested in reasoning." It is the faculty "which equips us to recognise and operate with reasons" (p.5).[3] He conjectures a close relationship between normative reasons and reasoning. On the one hand, reasoning is drawing inferences from a set of premises to a conclusion, in which "these premises are, or are thought to be, reasons." On the other hand, "reasons (justificatory) are the stuff of which reasoning is made" (p.67). But he also holds that the standard of correctness of reasoning is set by formal logic. The principles or rules of reasoning, according to Grice, are non-empirical in character. It is the logicians' task to specify different kinds of inference rules. They have a descriptive version, in which they specify relations between propositions, and a prescriptive version, in which they specify relations between attitudes. Without specifying what exactly the inference rules we follow in reasoning are, he says we cannot formulate the prescriptive version of inference rules as requirements of the form "if you accept A, then accept B", for this would be too demanding on the reasoner; an idea which later appears in (Harman 1986). Leaving this issue aside, Grice argues that in reasoning, you *think* of yourself as conforming to valid patterns of inference (p.14): A reasoner has a belief or desire that his reasoning conforms to one or the other, but not necessarily any particular, valid inference rule (p.7).

[3] I am referring here to the Locke Lectures Grice gave at Oxford in 1979, titled 'Aspects of Reason'.

There are a few problems with this account of reasoning. The first is that it this over-intellectualises reasoning: you can reason without having representations of formal inference rules, and even without having the concepts necessary to have those representations (such as the concepts of proposition, or entailment). Second, the account leads to Lewis Carrol's (1895) regress problem: If I do some reasoning, and have the belief or desire that my reasoning conforms to a valid pattern of inference, then what is the role of this attitude in the reasoning? If it is a premise, then it seems there must be another belief or desire that the inference including this attitude as a premise is valid, which again requires a belief or desire … and so on. Alternatively, if we hold that the respective belief or desire is not a premise of the reasoning, but plays some other role, the question is what that role may be. The third issue with Grice's proposal is that the claim about the predominant role of formal logic in reasoning is not obviously reconcilable with the claim that reasoning is an operation on normative reasons. Formal inference rules do not necessarily give us reasons. As Grice notices, too, we may have no reason to form beliefs that formally follow from what we believe, for example. Also, the relation of "being a reason-for" does not hold between the propositions that are the constituents of formal inference rule. Neither are propositions normative reasons, nor are there normative reasons for propositions. It remains unclear in what sense we accept the claim that following a valid rule of inference *is* an operation on reasons.

I have objected to Boghossian's and Pettit's accounts of reasoning along similar lines in chapter **1**. They all seem to assume, without offering an explanation, the compatibility (indeed identity?) of the *Rule-Following* and the *Reasons Claim*. Despite his insistence on the role of inference rules in reasoning, Grice's account of reasoning seems to be first and foremost focused on the *Reasons Claim*, which is why I have placed the discussion of his view here. He defines reasoning as an activity that allows us to recognise and operate on reasons but does not explain what these notions mean. Overall, Grice does therefore not offer

an account of the nature of reasoning that explain what the activity consists in, and how it differs from passive transitions of thought.

2.4. Joseph Raz

Joseph Raz roughly agrees with Grice that the faculty of 'Reason' is the faculty which equips us with the capacity to recognise reasons *and* to engage in reasoning. He investigates reasoning in relation to normative reasons and what he calls 'Reason': our rational powers. Understanding their relationship is crucial for understanding normativity more generally (Raz 2015), and what it is to be a person (Raz 2011: 85). In his most recent publication on the topic, Raz gives a detailed characterisation of the activity of reasoning, which he claims is "nothing but a way of responding to reasons" (Raz 2015: 151). In reasoning, Raz argues, the agent takes the reasoning to establish conclusive reason for the conclusion.[4] This taking is "manifested in accepting that deviation from relevance and from the other rules governing arguments is a mistake, being willing to correct them" (p.148). Unfortunately, Raz does not say what the attitude of "taking" (or "accepting") exactly involves, nor how the rules governing arguments relate to normative reasons. If we interpret him as claiming that the attitude of taking grounds reasoning, and thus marks the difference between active reasoning and mere passive thought, then we are back at Boghossian's problem: we have seen in **1.2** that a doxastic, dispositional and a rule-following construal of the attitude of 'taking' seems to fail to yield a plausible account of the nature of reasoning. But maybe this problem is less relevant for Raz then it is for Boghossian. For Raz does not seem to think that taking the premises to be reasons for the conclusion grounds reasoning; it is not sufficient for reasoning:

[4] I don't think Raz intends to say that the reasoner takes his reasoning to establish conclusive reason, but that the premises as a whole establish conclusive reason for the conclusion-attitude. It is unclear then whether the 'premises' refers to the premise-attitudes or their contents.

"If F is a reason for the agent to R, not all ways of coming to R when taking F to be a reason for it are cases of reasoning" (Raz 2015:153).

Not every action done for a reason, he says, involves reasoning. Raz implies that you take something to be a reason, whenever you act for a reason, none of which requires reasoning. In reasoning, in contrast, the agent comes "to realise that some items (premises) support others (conclusion)." In other words: "Reasoning is a special case of responding to reasons, responding by discovering, by realizing that C, or that P supports C" (p.159). As anticipated in section **1.5**, Raz also proposes that what is distinctive of reasoning is that the agent holds a certain attitude towards the reasoning process, here a realisation. His view is therefore subject to desiderata **D1** and **D2**: it should lay out the nature and content of the realisation in question and explain how the presence of that attitude makes the process one of reasoning. Regarding the content of the realisation, Raz remains vague: what you realise is either the conclusion of the reasoning or that the premises support the conclusion. Suppose it is the latter. Then what is it to realise that? Raz says that realising is not coming to believe something; it is instead an experience, a Eureka moment. But how does this make the process one of active reasoning? The fact that a thought process is accompanied by a particular experience does not make it active, for the emergence of that experience is presumably not something that is in the reasoner's control. Any mental process that is not reasoning could involve such a Eureka moment.

But Raz implies at some places that what you realise in reasoning is not that the premises support the conclusion, but the conclusion itself. He argues that what one is led to when one realises something is a belief. Therefore, he concludes, the conclusion of reasoning is a belief, and reasoning cannot lead to action or intention. The second claim - that reasoning leads to belief - follows from the first only on the additional assumption that what one realises

in reasoning is the conclusion, not that the premises support the conclusion. But if we interpret Raz's account of reasoning as saying that reasoning essentially requires the reasoner's realisation of the conclusion, it becomes implausible. It then fails to explain in what sense reasoning involves the reasoner's acknowledgement of a link between what he reasons from and his response. It also fails to explain why reasoning is active: any kind of thought process that causes you to acquire the belief that p can cause you to realise that p. So we shouldn't assume that what you realise in reasoning is the conclusion, if we want this realisation to be that in virtue of which you count as reasoning. Without that assumption, however, Raz's arguments to the conclusion that the conclusion of reasoning is belief, is invalid. We thus have no reason to doubt that reasoning cannot lead to intention (or action).

In fact, Raz himself does not seem to think that what you realise in reasoning is the conclusion. He later gives a third interpretation of the attitude of realisation that he claims is essential to reasoning. What you realise, he says, is that the conclusion is a well-supported answer to the question one is considering. But it is not clear what the role of questions are in reasoning. Reasoning does not require asking yourself a question, and in a sense, every intention or belief you form trivially counts as an answer to some question.[5]

I sum, I stick to the following – I hope charitable – reading of Raz's account: reasoning is marked by the agent's *realising* that a certain relation holds between the contents of the premises-attitudes and the conclusion attitude. This realisation is manifested by a certain experience, a Eureka moment. But what exactly is the content of the realisation, and what its nature? And how does this realisation make the process one of active reasoning? Raz does not answer these questions, and thus leaves unexplained in what sense "Reasoning is the reason-guided mental activity of finding out how we should orient ourselves towards the world" (2015:161). He argues at length that 'Reason', i.e. our rational powers, allow us to

[5] I argue in chapter 5 that accounts of reasoning in terms of answering questions are not promising.

recognise and respond to reasons, and that they are therefore constitutive of being a person (Raz 2011). But he does not explain what these "rational powers" are. I find that deeply unsatisfying. To say that animals like pigeons lack 'Reason' and are thus neither subject to nor responsive to reasons in the way we are seems to describe a phenomenon to be explained, rather than explain it.

2.5. Antti Kaupinnen

Antti Kaupinnen (2018) develops one of the most detailed versions of the *Reasons View* (the other being Jonathan Dancy's). He first rejects Boghossian's Taking Condition for similar reasons discussed in chapter **1**, and argues that for a thought 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." (p.418).

He agrees with Broome that an attitude of *endorsement* towards the transition that is made is necessary for reasoning. Kauppinen is unsatisfied with Broome's construal of the attitude of 'seeming right'[6] and rejects that rule-following is necessary for reasoning. Unfortunately, however, Kaupinnen does not offer an alternative account of what the attitude of 'seeming right' consists in. Instead, he says that in practical reasoning, for example,

> "you somehow treat the contents of the inputs to reasoning, taken together, as a reason for the act or the intention" (p. 418).

To treat the premises as reasons is to

[6] Broome first called it the attitude of 'seeming right' (Broome 2013: 13), and later the attitude of 'seeming correct' (Broome 2014b).

> "have a positive attitude towards performing the act of forming the intention, 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 asking for and giving 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 reasons for the latter." (p.419).

This is a plausible starting point towards an account of what it is to respond to and treat things as reasons. In fact, the one I propose in chapter **4** is similar is some respects.[7] But, as far as the account of *reasoning* is concerned, we need an account of what the nature and content of the "positive attitude" in question is. We also need to explain the role it plays in the reasoning process and how it makes the transition counts as an activity of the agent. (Recall, these are the desiderata **D1** and **D2** that an account of reasoning should meet, or so I have argued.) Kauppinen, however, is silent on these issues and goes on to describe practical reasoning without reference to treating things as reasons. Practical reasoning, he says, is triggered by a state of being unsure or puzzled about something. This kicks off a search for reasons, or more generally, considerations which bear on the problem.[8] You then weigh those reasons, thereby *taking* the facts (or supposed facts you believe) to favour certain actions. You *see* that certain states of affairs favour something. You *register* when you have considered enough reasons and let the "deliberative gauge" point you in some direction. The weighing operation ends

[7] It seems Kauppinen intends to explain what it is to treat considerations as reasons in terms of the role this 'treating' plays in reasoning. But then treating something as a reason cannot explain what reasoning is, or the explanation is circular.

[8] Following Dancy (2004), Kauppinen accepts that these consideration include reasons, but also their modifiers and conditions. In brief, reasons are considerations that favour (or disfavour) certain responses. Conditions are considerations which enable (or disable) those considerations to be reasons. The fact that she is drowning is a reason for me to jump in, we might say, only on the condition that I can swim. Modifiers (intensifiers and attenuators) affect the weight of the reason. Both Dancy and Kauppinen agree that in reasoning, you the things you consider can play either of these roles.

when you *take* the balance of reasons to favour some particular action (p. 418). In response to this taking, you form the respective intention.

This picture raises many questions. What is the nature and role of my *taking* certain facts to favour an action in this reasoning process? Is it just to 'treat' them as reasons? And what is it to *take the balance of reasons* to favour a certain kind of action? This is substantially different from seeing one consideration or a set of considerations as reasons. The former requires a certain weighing of the reason, the latter does not. Is this taking necessary for reasoning? A satisfactory account should answer these questions and decipher the metaphor of the *"deliberative gauge"* that moves us in one or the other direction.

In sum, we see that Kauppinen is among the authors who argue that for some transition in thought to count as an instance of reasoning, the agent must hold some attitude of endorsement towards the transition (see the list in section **1.5**). Kauppinen gives an account of reasoning in terms of treating things as reasons. If you have the right kind of endorsement towards a thought process leading you from some beliefs to an intention, for example, you count as treating the contents of those beliefs as reasons, and because of that, as reasoning. However, Kauppinen does also not explain the nature and content of that attitude, nor how it makes the process count as an activity. The proposal, though generally plausible, fails to meet desiderata **D1** and **D2**.

2.6. Jonathan Dancy

In his "Practical Shape", Dancy (2018) defends what I believe is to date the most rigorously developed version of the *Reasons View* of reasoning. Two things set Dancy's approach apart: first, in his inquiry, practical reasoning takes priority over theoretical reasoning. Most theorists have been misled, Dancy argues, because they let their view of practical reasoning

be constrained by mistaken conceptions of theoretical reasoning. Second, he does not aim to give necessary and sufficient conditions for reasoning, but instead describes certain "focal" cases of successful reasoning and argues that a thought process qualifies as reasoning in so far as it is relevantly similar.

Dancy argues we should not think of reasoning in terms of inference, using the notions of premises and conclusions. Reasoning is deliberation on the basis of considerations, where considerations are whatever we start from or start with in practical as well as theoretical reasoning. In theoretical reasoning, we respond to these considerations with forming a belief. In practical reasoning, we respond to these considerations with intention or action. He writes:

> When an agent deliberates well and then acts accordingly, the action done is of the sort most favoured by the considerations rehearsed, taken as a whole – just as when an agent reasons well and then believes accordingly, the belief formed is of the sort most favoured by the considerations rehearsed (p.34).

The favouring relation is central to Dancy's account.[9] He takes it to be a relation between states of affairs, agents and responses. Drawing on his earlier work, he distinguishes between reasons, conditions under which considerations are reasons (enablers/ disablers) and modifiers of the strength of reasons (intensifiers/ attenuators). He says: "Reasoning is not just a matter of producing reasons and trying to work out which side is stronger. It is much more interesting than that. The favouring relation brings with it all sorts of supporting relations" (p.45). But because reasons are at least part of that which favours certain responses, Dancy's

[9] Dancy sometimes seems to use ‚reasons' and ‚favourers' interchangeably (e.g. p.38,45,90,141). Strictly speaking, if a reason favours a certain response, then the sense in which a set of considerations consisting of reasons, modifiers and conditions favours a response must be a different one.

view can be considered a version of the *Reason View,* even if he doesn't spell it out in terms of 'reasons', but in terms of 'favouring'.

The quoted passage above describes the relation between an action (or belief) and the considerations on which it is based *if* the agent reasons *well*. It is not an account of what reasoning is but sets a standard of assessment for reasoning. However, in Dancy's approach to the nature of reasoning, this understanding of good or correct reasoning plays a central role. For according to Dancy's focalism, events count as instances of reasoning "to the extent to which they are relevantly similar" to successful or focal cases of reasoning (p.7,106). These focal cases are to be understood in a normative, more precisely, evaluative way (p.108). They are cases of ideal reasoning which involve no mistakes. The focalist, says Dancy, recognises not only correct, but also flawed reasoning as reasoning because of its relevant similarities to the focal cases. The relevant similarity is given by the material theory of reasoning which "lays out the different ways in which a consideration can make a difference to the strength of a case" (p.46). In short, it seems we are to judge whether some chain of thinking is reasoning depending on the extent to which the considerations adduced play one of the recognised roles (reasons, conditions, modifiers) in favouring the response. Dancy's approach is thus different from the ones discussed before in so far as he first defines what good or correct reasoning is, and based on that gets an understanding of what the activity of reasoning consists in. He contrasts reasoning not explicitly with passive thought, but other ways of responding to reasons.

However, I see two major problems with this focalist approach. First, it is possible to reason to a response R on the basis of considerations which play no role in favouring R. Suppose I reason 'It would hurt them if I do it. No one will see it if I do it. So, I shall do it.' The focalist would want to recognise this as an instance of incorrect reasoning. But the two considerations adduced are neither favourers, nor modifiers, nor conditions. (One could argue

that the second would be an intensifier or condition if the first was a reason, but that does not help.) Nor is the response the one most favoured by the considerations. On what grounds then, could the focalist recognise this as incorrect reasoning? Secondly, some thought process might perfectly well resemble correct reasoning and be similar to focal cases in "the relevant respects" without actually being an instance of reasoning. It may not constitute an agent's *activity*. One might just happen to form the intention or belief that is most favoured by a set of considerations, without engaging in reasoning. The merely causal thought process would be correct reasoning *if* it was reasoning. Dancy agrees that reasoning is something we do, not something that just happens to us: "In reasoning I am as active as I ever am" (p.154). So if his focalism does not allow us to differentiate between passive thought and reasoning, this is a problem for the account.

The problem parallels the following issue: Suppose the beliefs in some considerations cause you to acquire another belief on their basis, and that the latter belief happens to be the one most favoured by those considerations. This response does not count *as a response to reasons*, it seems, unless you somehow recognise (or realise, in Raz's terms) that it is the one most favoured by the considerations on which it is based. There is a difference between reacting to how we take the world to be and responding to the reason we take ourselves to have. Analogously, deliberating well requires more than getting to the response most favoured by what we believe. We have to get to them *because* they are the ones most favoured, more so: because we somehow recognise, that this is so. Dancy does not seem to follow the authors discussed so far in holding that active reasoning, but not passive thought, requires that agent to hold a certain attitude towards the transition. Given these issues, maybe he should hold that reasoning requires the reasoner to somehow *recognise* or *realise* or *take it* that the response is the one most favoured. The account could then be added to the list in

section **1.5** and would be subject to desiderata **D1** and **D2**: it would have to explain what the attitude towards the process is.

The problems pointed out seem to stem from the focalist's focus on correct or good reasoning, rather than on what makes something reasoning in the first place. We must separate the questions of what reasoning is and what it is for it to be correct or good, whether or not we aim to give necessary and sufficient conditions. But which are we to answer first? On the one hand, it seems plausible that we must first know what reasoning is in order to say what correct or good reasoning is. On the other hand, one might argue that we cannot define the boundaries of what counts as reasoning unless we are geared to good or correct cases of reasoning. But, even if our intuitions as to what counts as correct or good reasoning inform our theorising about the nature of reasoning, this does not entail that an account of the nature of reasoning must be given solely and wholly in term of good or correct reasoning. In other words: what it is for reasoning to be correct tells us *something* about the nature of reasoning. But there is more to say about it.

In sum, I do not see strong reasons to prefer Dancy's focalist approach to reasoning over the alternative of first investigating the nature of reasoning, and then the standards of evaluation of reasoning.

But let us put this issue aside for a moment and look at what Dancy calls his material theory of practical reasoning. It can be summarised by two core claims: Practical reasoning is weighing practical reasons and, in reasoning, the response is the one most favoured by the considerations adduced, taken together. Dancy explains what it is for a response to be most favoured, but does not explain what it is to weigh reasons. Suppose you reason to the intention to ϕ on the basis of some considerations. Dancy explains how these consideration can play different roles "in making a difference to the strength of the case" for ϕ-ing (p.46): they can

be reasons, modifiers or conditions. But none of this says anything about the *weighing* of reasons. Dancy exclusively discusses cases in which all considerations adduced contribute to strengthening the *same* case. When we *weigh* reasons, however, the considerations adduced strengthen *different* cases, and we have to compare the strength of the favouring given for each to determine which is favoured most. Responding to a set of considerations all of which make the case for some sort of response requires "a sensitivity to reasons" as well as their modifiers and conditions (p. 46). But *weighing* reasons additionally requires comparing the normative support given for the different courses of action. How do we get "a sense of how to respond" (p.61)? Dancy notes that "the comparative aspect of deliberation" is hidden in his model (p.44). But it is essential to it, if reasoning is understood in terms of what response is most favoured. It seems we must understand the "sensitivity to reasons" as well as how we get a sense of how to respond in order to understand the activity of reasoning.

Dancy concludes that his account of practical and theoretical reasoning entails that we are normatively sensitive agents:

> Perhaps the bottom line here is a conception of ourselves as normatively sensitive agents, capable of responding to reasons as reasons wherever they are to be found. A reason is a reason for a response, and in responding we are active; we are not pushed about by reasons (as if by our desires, say), but respond to them as calling for this rather than that from us. The connection between reason and response is quite different from the connection between cause and effect. The former is normative" (p. 175).

Note that the last sentence of the quoted paragraph is puzzling. That there is a normative relation between a reason and a response does not explain the response *qua* response. Suppose my current situation normatively requires me to stay at home today. I think about the relevant features of the situation and decide to stay at home. Is my decision a *response* to the reasons I have? As argued above, it is a response to reasons only if the normative relation in question

plays some role in the explanation of my decision. I must somehow *manifest a sensitivity* to this relation, which again raises the question of what that is.[10]

2.7. Rules or Reasons?

All these versions of the *Reasons View* have a claim in common, namely that reasoning is defined in terms of responding to reasons or favourers, which in turn is explained in terms of the agent's having a certain attitude towards the thought process or that which it is about: a recognition that some considerations are reasons (Grice), a realisation that some things support others (Raz) or some other sort of positive attitude towards forming the conclusion-attitude in a certain way (Kauppinen). The challenge that all views face but fail to meet is to explain the nature and content of this attitude, and its role in the reasoning process, so that it makes the process count as an instance of active reasoning; as something you *do*. (Thus, they do not meet desiderata D1 and D2 formulated in chapter **1.5**). I argued that Dancy's account does not allow us to distinguish instances of reasoning from transitions in thought which resemble correct reasoning but are not instances of reasoning or *responses* to reasons *as* reasons, because it does not require the agent to 'recognise' the favouring relation involved (whatever that amounts to). If some form of recognition like this is necessary for reasoning, or at least present in focal cases, then Dancy's account raises the same question as the other three.

Defenders of the *Reasons View*, I conclude, do not provide us with a satisfactory account of the nature of reasoning. This is due to the fact that they explain reasoning in terms of

[10] In an earlier work, Dancy says that in reasoning, one responds to the consideration rehearsed *as together calling for or favouring it*"(Dancy 2014:6). I take the question of what it is to respond to a consideration *as* favouring a response to be the same as the question of the normative sensitivity. I answer it in chapter 4.

responding to reasons, without explaining what this consists in. Pointing to the notions of *treating* considerations *as* reasons, *recognising* reasons or *taking* things to be reasons is insufficient. They do not explain but merely rephrase what needs to be explained. I do not deny that reasoning essentially involves or even is responding to reasons. I do not deny the *Reasons Claim* but have to conclude that it has not been established yet. This conclusion is parallel to that of chapter **1**. There I have argued that the different versions of the *Rule-Following View* are also unsatisfactory. I have not denied the *Rule-Following Claim,* that is, the claim that reasoning is a rule-following activity, but argued that it remains unclear in what sense it is true.

Where does this leave us in the search for an answer to the question of what reasoning is? The debate I have reviewed is based on two implicit assumptions. First, most believe that reasoning must be defined in terms of either rule-following or responding to normative reasons. Those who defend the *Rule-Following Claim* or the *Reasons Claim* do so by developing accounts of the nature of reasoning in terms of rules or reasons, respectively. Second, the two alternatives are mostly assumed to be mutually exclusive. I believe we should give up both of these assumptions. We should not begin our inquiry of reasoning with the aim of defining reasoning either in terms of rule-following or in terms of responding to reasons and allow for the possibility that reasoning turns out to be a way of responding to reasons, a way of rule-following, or both. Before I pursue this approach, I wish to make salient a fundamental difference between the *Rule-Following View* and the *Reasons View*.

2.8. Reasoning and Normativity

According to the *Reasons View*, reasoning is the activity in which we exercise our capacity to detect normative features of the world, to "see" what the world requires of us and to steer

our minds in accordance with these demands. Suppose, for example, I reason as follows: "She is very ill. I have a free evening today. So, I shall make her dinner", where the first two statements express true beliefs and the last sentence an intention which I form on the basis of the two beliefs. According to the *Reasons View*, what is essential to my reasoning is that I 'see' the fact that she is ill *as a reason* for making her dinner. My recognition of the normative import of this fact is what guides my reasoning. I do not explicitly think about what is a reason for what, but respond to facts *as* normative reasons. Reasoning is normative in the sense that normative features of the world play an essential role in each and every instance of reasoning. The normativity involved in reasoning is genuine normativity, which is related to what we ought to do.

According to, for example, Broome's and Pettit's versions of the *Rule-Following View*, the nature of reasoning is independent of any genuinely normative features of the world and our capacity to respond to them. Reasoning is following a set of rules of logic, rules of coherence or rules of rationality. Following them correctly allows us to bring our minds in order, where this internal order may fit the world or not. Of course, how we reason may be correct or incorrect, whether or not it is correct is independent of what there is reason for us to do or believe. Reasoning, we may say, is normative only in the weak sense in that it involves a standard of correctness. In reasoning, you measure your behaviour against correct rules. But on this view, reasoning is not normative in the strong sense in that genuinely normative features of the world play an essential role in reasoning.[11] The rules of reasoning set a standard of correctness of reasoning, but they neither give us reasons nor do they derive from normative reasons.

[11] The question of what role normativity plays in reasoning here is different from the question of normativity in the context of meaning and rule-following that Wittgenstein scholars discuss. As Sluga (2011: 121 - 122) argues, Wittgenstein is right in saying that some rules are normative, but other are not.

Does this difference between the two views imply that the *Reasons Claim* and the *Rule-Following-Claim* are mutually exclusive? Not necessarily. For we can accept the *Rule-Following Claim*, i.e. the claim that reasoning is a rule-following activity but deny that the rules of reasoning are independent of normative reasons. The *Rule-Following Claim* allows for different conceptions of the rules of reasoning, just as it allows for different conceptions of rule-following. I suggest we give up the assumption that the rules of reasoning derive from abstract coherence requirements of rationality or logic. If the rules of reasoning are related to normative reasons, then we can accept both claims, and account for the intuitive force of both.

2.9. Conclusion

The view that reasoning is responding to normative reasons (*Reasons View*) has been assumed by many but defended by few. I have argued that Grice's, Raz's, Kauppinen's and Dancy's elaborate versions of the *Reasons View* are all unsatisfactory. They analyse reasoning in terms of responding to reasons, treating things as reasons and recognising reasons, without explaining sufficiently what these notions refer to. The *Reasons View*, I conclude, remains surprisingly underdeveloped.

In the next chapter, I develop an account of reasoning which takes as a starting point neither the *Rule-Following Claim*, nor the *Reasons Claim*, but the common ground shared by all accounts that we have seen in this and the previous chapter: that a thought process counts as an instance of reasoning if and only if the agent holds an attitude of endorsement towards it. This attitude, I argue in the next chapter, is the attitude of approval. In chapter 4, I turn to the second question raised in the introduction of the thesis: *(II) What is responding to normative reasons*? This chapter has illustrated the need to address it anew. I develop a novel account of what it is to respond to normative reasons, which goes further than the ones I critiqued

here. On the basis of this account of responding to normative reasons, I argue that reasoning – understood as I suggest in the next chapter - is a way of responding to reasons.

3. The Approval Account

3.1. Introduction

Reasoning is the mental activity through which we form, revise or discard attitudes – like beliefs and intentions – on the basis of others. In this chapter, I propose a novel account of the nature of this activity, the *Approval Account*.

In chapters 1 and 2, I presented and critiqued a number of accounts of reasoning and argued that their common denominator is the following claim: What is essential for a thought process to count as an instance of reasoning is that the reasoner holds a certain attitude of endorsement towards it. Furthermore, I have argued that theories of reasoning that are based on this claim are subject to the following desiderata. They should

D1 explain the content and nature of the attitude that is essential to and grounds reasoning.

D2 explain its role in the reasoning process and show how its presence makes the process count as *active,* and

D3 not render reasoning conceptually over-demanding, that is, not require reasoners to have complex concepts like that of a normative reason, belief, entailment relation, etc.

This claim that reasoning is essentially a thought process the agent endorses is also at the core of the *Approval Account*. It is therefore subject to **D1** – **D3**. The aim of this chapter is to present the *Approval Account* and to show how it can meet these desiderata.

In section 3.2 I state the *Approval Account*. Put roughly, reasoning is thinking you approve of. What is essential to reasoning is that you hold an attitude of approval towards the

thought process as you experience it. I explain in detail what the nature and content of this attitude is. The attitude of approval, I argue, is evaluative in nature. To approve of your behaviour is to evaluate it against a certain standard. Because this standard involves the normativity of 'reason' and 'ought', reasoning is an inherently normative activity. In **section 3.3** I explain in detail the normative nature of the attitude of approval, and thus of reasoning. **Section 3.4** argues the approval that is essential to reasoning is a special kind of approval: it is based on the reasoner's implicit conception of normative reasons. **Section 3.5** explains the role of the attitude of approval (and that which it is based on) in the reasoning process. I distinguish between explicit premise-attitudes, implicit premise-attitudes and non-involved attitudes. **Section 3.6** summarises the discussion and shows how the *Approval Account* meets desiderata **D1**, **D3** and **D4. Section 3.7** further defends the claim that reasoning essentially involves genuine normativity against the view that reasoning is following rules of coherence that are independent of normative reasons and what we ought to believe or do. **Section 3.8** argues that, according to the *Approval Account*, reasoning is an intentional mental action. It thus meets desideratum **D2**. **Section 3.9** argues that the account can accommodate instrumental reasoning, before **Section 3.10** briefly concludes.

I here consider only reasoning that leads us to acquire attitudes. If there is reasoning that causes us to give up attitudes, for example, I do not consider it here. Further, I focus on theoretical reasoning that leads you from some beliefs to another, and practical reasoning that leads you from some beliefs to an intention. I apply the account to practical reasoning from intentions in section **3.9**, where I consider instrumental reasoning. There are other kinds of reasoning that I cannot deal with here, in particular, hypothetical reasoning[1] and reasoning

[1] The mental state of supposing that p is different from that of believing that p. When you suppose that p, p is not "regarded as really true" (Velleman 1996: 708). You accept the truth of p for the sake of argument, but not in and of itself. To apply the *Approval Account* to suppositional reasoning would require an understanding of the nature of the mental state of supposing that p. How does an agent see the content of a

with degrees of beliefs or credences.[2] Maybe we can also reason from and to other kinds of attitudes, such as fears and hopes, for example.[3]

The claim that at least some practical reasoning leads to an intention is uncontroversial in the contemporary debate.[4] However, in what follows, I assume that *all* practical reasoning leads to intention. This claim is controversial, in particular because some defend the arguably Aristotelian thesis that practical reasoning can lead to action.[5] I cannot enter into this debate here and shall only briefly defend my assumption with two points.

First, I believe practical reasoning cannot 'lead to' action in the same sense in which theoretical reasoning 'leads to' belief, because reasoning is a *mental* action and does therefore not include physical movements. If an instance of reasoning causes you to raise your arm, the physical movement that is part of this action cannot be part of the reasoning. The action can therefore not be part of the reasoning, even if it is caused by the reasoning.

---

supposition? What is the first-person contents of a supposition? A further question is what the conclusion-attitude of reasoning on the basis of the supposition that p is. Is it a belief "under the supposition of p" or anther supposition, or a belief in a conditional (Green 2000:378)? Reasoning on the basis of suppositions is different from reasoning on the basis of beliefs. Space does not permit me to examine it further here.

[2] I cannot give a full account of reasoning on the basis of credences here. But the investigation begins as the investigation of reasoning on the basis of beliefs. Suppose you have some degree of confidence that p. Considering p, how do you see p? You do not see it as being true in the same way as you do when you fully believe that p. Neither do you see it being true with probability p. Credences are not beliefs that are probabilistic in content. How you experience the content of a credence explains how you are aware of the degree of your confidence in p. This is analogous to the case of beliefs: when you reason on the basis of a belief that p you are aware of how you relate to p because you see p as being true. We don't track probabilistic relations between the attitudes we reason with directly, which would require higher-order attitudes. Instead, how we hold contents in mind determines how we see those contents, and thus how we treat them in reasoning. Reasoning commits you to believing that if the premises, if they were as you perceive them, were all that was relevant to your response, then they would (partly) make it the case that you have sufficient reason to respond as you do. This guarantees that in reasoning on the basis of credences, the conclusion-attitude "fits" the nature of the premise-attitudes, and their content. You track what you reason about, and how you relate to it.

[3] See e.g. Dancy (2018: ch. 7) on reasoning with doubts and hopes.

[4] Most of the theories discussed accept this claim. Only Raz (2015; 2011: 123) denies that there is reasoning to intention, arguing that practical reasoning leads to the belief that you ought to do something. I discuss his view in chapter 2.

[5] See e.g. Dancy (2018: 34, 62, 63), Fernandez (2016). According to Anscombe (see 'On Practical Reasoning, in Raz's 'Practical Reasoning' (1978)) and Broadie (1968), Aristotle thought that practical reasoning leads to action.

Second, I take it that practical reasoning leads to either so-called prior intentions or intentions-in-action. What Searle (2010: 83) calls prior intentions are intentions you form in advance of your actions, they are plans for the future. You can, for example, reason to the intention to book the train tonight. Because you do not execute this intention now, the intention is a plan for the future, a prior intention. Intentions-in-action, in contrast, are the intentions with which you act when you act intentionally. If you now decide to stand up, you act with the intention-in-action to stand up. Many questions arise regarding the relationship between prior intentions and intentions-in-action, but they need not bother us here. (Does my prior intention to book the train ticket cause or translate into an intention-in-action when the time has come, and I book the ticket?). We can assume, following Searle, that an intentional action is a bodily movement caused by the agent's intention-in-action to perform it. Thus, when you reason to an intention which is an intention-in-action, you reason to something that is an essential aspect of an action. For example, when I reason to the intention to close the window now, and immediately get up to do so, the conclusion of my reasoning is something that is part of the action it leads to. But that does not make the action (part of which is the bodily movement) the conclusion of the reasoning.[6] The reasoning leads only to an intention. So, when I speak of practical reasoning, I speak of reasoning that leads to an intention, where this intention may be an intention-in-action or a prior intention. With that, I can now present the *Approval Account*.

3.2. The Approval Account

The account of reasoning I am going to propose reads:

[6] I am here defending the thesis that practical reasoning leads to intention against an objection that Dancy (2018: 167-9) raises against Broome (2013: 250). Dancy insists that we cannot separate the action from the intending. I insist we cannot conflate the intending with the action. In any case, I suspect that my arguments here would not be affected by calling practical reasoning to intention-in-action reasoning to action.

Approval Account: Agent N reasons from attitudes $A_1(a_1)$ … $A_n(a_n)$ to $A_C(a_C)$ if and only if, and if so because, N immediately *approves* of his/her own forming A_C on the basis of $A_1(a_1)$ … $A_n(a_n)$ as experienced by him/her.

Suppose you form the belief that Blackwell's is closed on the basis of the belief that it is 5pm and the belief that if it is 5pm, then Blackwell's is closed. As we have seen in chapter 1 and 2, many have argued that you form the belief that Blackwell's is closed through the activity of reasoning if and only if you endorse the process leading to it in some way. I argue that this endorsement consists in your immediate approval of this transition. You count as reasoning in so far as you immediately *approve* of your forming the belief that Blackwell's is closed on the basis of believing that it is 5pm and that if it is 5pm, then Blackwell's is closed. Reasoning is a process which takes you from some attitudes to another. The approval is part of that process, and its immediacy is crucial. Reasoning is not the approval of past transitions. **Section 3.8** explains the role of the approval in the process in detail. To say that you approve of your forming of the conclusion-attitude in this way is to say that you hold the attitude of approval towards this transition as you experience it. In this section, I first explain the content or object and then the nature of the attitude of approval.

3.2.1. The Content of the Attitude of Approval

According to the *Approval Account*, reasoning requires the approval of the transition as you experience it. How do you experience the transition? Let us begin with the premises. When you reason, you think about the contents of the premise-attitudes, not the attitudes themselves. The contents of attitudes like beliefs are propositions. A proposition is a meaningful

description of a state of affairs, and it can be true or false. A true proposition describes an obtaining state of affairs. A fact is a state of affairs that obtains.

Different true propositions can describe the same obtaining state of affairs in different ways. The propositions then have, or are, different meanings.[7] Attitudes that have propositions as their contents are thus intentional attitudes, they are about something beyond themselves.

For any such attitude A(*a*), where the letter 'A' stands for the type of attitude, and the letter 'a' for the content, the nature of 'A' determines how the content 'a', appears to you. It determines how you are aware of it.[8] And it thus determines how you treat it in reasoning. I call the content of a propositional attitude as it appears to you, or as you experience it, the *first-person content*. In reasoning, you operate on *first-person contents*, not plain propositions. For example: Believing that 'it is 5pm' is taking 'it is 5pm' to be true. When you reason on the basis of the belief, you treat its content as being true. The proposition – the content of the belief - appears to you as a fact, and you treat it as one. The first-person content of beliefs are thus facts. From a third-person perspective, the content of a belief is of course merely a supposed fact, for the belief may be false. But from your perspective it is not merely supposed. The proposition appears to you as a fact. As Frege (1979) says at the beginning of his 'Logic', reasoning is "making a judgement because we are cognisant *of other truths* as providing a justification for it" (emphasis added).

How then do you experience the transition which leads you to acquire the conclusion-attitude? In theoretical reasoning, you take the conclusion to be true because you take the

[7] Propositions may be understood as meanings with reference, or references with meaning.

[8] Broome (2013: 251-3), argues that in reasoning, you are not just conscious of a content and an attitude, but "you must be conscious of the pair in a particular way that recognizes the connection between them. You must be conscious of the proposition as believed by you" (253). Thus, he says, you reason about the marked contents of attitudes, which are pairs consisting of an attitude and a content, for example <belief, p>. I agree on the first part. But I find it hard to make sense of the idea that we reason about such abstract objects as marked contents.

premises to be true. Some beliefs cause another, so that the latter is based on the former. This could also happen sub-personally, without you doing anything. You would then not be reasoning. For the transition to count as reasoning, you have to approve of it. But you approve of it only as you experience it. And you do not experience it as your coming to believe something because you believe something else. You do not take the reflective perspective which would allow you to form such meta-beliefs. From your first-person perspective as the reasoner, what seems to happen is that some facts compel you to take a proposition to be true. You come to believe the conclusion because of the premises. Of course, what appears to you as facts may not be facts, and the correct causal description of the process (from a third-person perspective) is that you come to believe something because of other beliefs. But this is not how the process appears to you. Theoretical reasoning requires that you approve of the transition as you experience it. That is, it requires that you approve of your *coming to believe a proposition because of some facts*. In the example, what you approve of is your coming to take 'Blackwell's is closed' to be true because 'It is 5pm' and 'If it is 5pm, then Blackwell's is closed'.[9] Analogously, practical reasoning from beliefs to an intention requires you to approve of the transition as you experience it. You approve of your committing to a certain course of action (now or in the future) because of some facts.

As we have seen in chapters 1 and 2, defenders of the *Rule-Following View* and defenders of the *Reasons View* argue that reasoning requires an attitude towards the process. The views differ not only in what the attitude is, but also in what the exact content of this attitude is. Regarding the latter, I would now like to contrast my view with alternatives. On the view I propose, the reasoner neither approves of a relation between the contents of the attitudes

[9] The 'take' has a narrow scope. You come to take ['Blackwell's is closed'] to be true because of some facts.

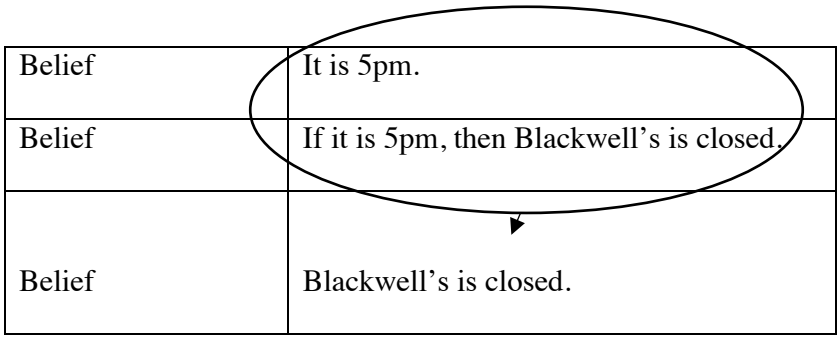
involved, nor of a relation between the attitudes themselves, nor of a relation between the contents of the premise-attitudes and the contents of the conclusion-attitude.

I follow Dancy's (2018) example and use grids to represent exemplary instances of reasoning. This helps to illustrate the differences between alternative views and mine. For each, the left column shows the attitude type, and the right column the content of the respective attitude (a proposition). The first two lines represent the premise-attitudes, the third the conclusion-attitude. The circle and the arrow illustrate the transition, process or relation that the reasoner allegedly has an attitude towards when deriving the conclusion.

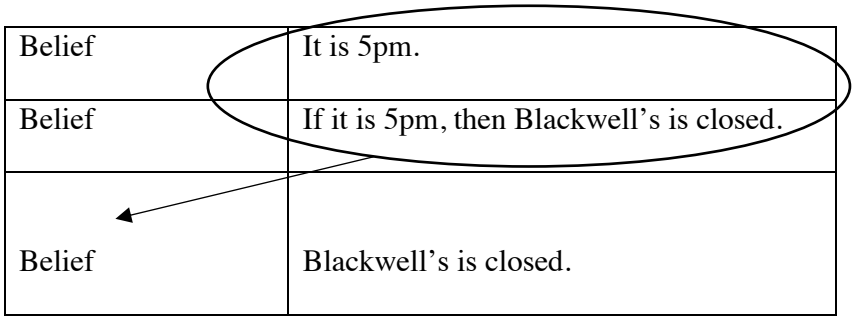
The view I propose is not that you approve of a relation between attitudes. You do not approve of your coming to believe that Blackwell's is closed because you believe that it is 5pm and you believe that if it is 5om, then Blackwell's is closed. This view could be represented as follows:

Belief	It is 5pm.
Belief	If it is 5pm, then Blackwell's is closed.
↓ Belief	Blackwell's is closed.

Nor do you approve of the relation between the propositions that are the contents of the attitudes constituting the reasoning process.



A third possibility would be to say that you approve of the relation between, on the one hand, the propositions that are the bare contents of the premise-attitudes and, on the other hand, the conclusion-attitude. Dancy (2018:44) argues that a relation between the considerations you adduce and your response (the favouring relation) is the central relation in reasoning. He represents it as follows:



On my view, you do not track a relation between contents, nor a relation between attitudes, nor a hybrid of those. What you approve of is your behaviour viewed from the first-person perspective. Two differences between my view and the three above are noteworthy. First, on my view, the starting points of the transition that you approve of are the contents of the premise-attitudes *as you experience them*. From a third-person perspective, we describe the contents of your beliefs as propositions. But from the first-person perspective, you experience them as facts. To illustrate this in the grid, we have to add a column, which represents these first-person contents of the premise-attitudes. In the case of beliefs, they are facts. The second

difference between my view and the three above is that you do not approve of these contents as standing in a certain relation to the conclusion-attitude or its content. Instead, you approve of *your coming to believe* something because of them. You approve of a kind transition in thought as prompted by certain facts (which is difficult to illustrate precisely).

Belief	It is 5pm.	Fact: It is 5pm
Belief	If it is 5pm, then Blackwell's is closed.	Fact: If it is 5pm, then …
Belief	Blackwell's is closed.	

This last illustration might give the impression that the nature of the conclusion-attitude of an instance of reasoning is a function of the first-personal content. This is not what I want to suggest. It would entail that for two instances of reasoning, if the first-personal contents of the premise-attitudes are of the same kind, they must lead to the same kind of conclusion-attitude. But you can reason from some beliefs to another belief or to an intention.

For the rest of this chapter, unless I say otherwise, when I speak of premises, I mean the first-person premises. Until section **3.9**, I only consider cases of reasoning on the basis of beliefs. In these cases, the first-person premises are the premises if they were true.

3.2.2. The Nature of the Attitude of Approval

According to the *Approval Account,* reasoning requires that you approve of your behaviour as you experience it (that is, for reasoning on the basis of beliefs, your acquiring of an attitude on the basis of some facts). At the end of this chapter, I argue that your approval is intentional, which is why it makes the process you approve of count as something you do, an act of yours.

We can understand the attitude of approval as a propositional attitude. It would thus be more precise to say that you approve of the (supposed) fact that you behave in this and that way. It seems a mere grammatical question whether or not you can approve of something you do not believe exists. Maybe you can approve of things as you suppose or imagine them to be. Maybe you can only approve of something you take to be the case. If that is so, then your approval of p would entail that you believe that p. If you approve that p, then p is presupposed. The attitude of approval would then have neither a world-to-mind, nor a mind-to-world direction of fit, but a presupposed direction of fit (Anscombe 1957: §32; Searle 2010:29). If p is false, your approval is not incorrect but misguided, just as your shame would be misguided if you were ashamed of your big nose but did not have a big nose (Searle 2010).

It seems plausible to me that in the case of reasoning, the attitude of approval does presume a belief. When you approve of your behaviour, you take what you approve of to be the case. Further, in reasoning, your approval is not misguided; the belief is true. This belief is not a higher-order belief about the conclusion-attitude, but a form of practical knowledge you have in forming that attitude (Anscombe 1957). When you act intentionally, you normally know what you are doing without observation. If reasoning is an intentional mental action (as I argue in section **3.8**), then you know what you are doing when you reason. On the assumption that knowledge that p entails the belief that p, we can ascribe to you a true at least implicit belief about the transition to the conclusion-attitude. It is necessary but not essential to reasoning; it is not that in virtue of which you count as reasoning.[10] In any case, however, the approval that is essential to reasoning is the approval of your actual behaviour as you experience it. Suppose you form the belief that q on the basis of the belief that p, but what you approve of is forming the belief that q because of z. Your actual behaviour does then not

[10] This is analogous to Broome's (2019). argument regarding the "linking belief", which he takes to be necessary but not essential for theoretical reasoning.

count as an instance reasoning despite your approval, for you do not approve of *it*, but some other transitions which you in fact do not make.

The attitude of approving is *evaluative* in nature. To approve of p is not only to represent that p, but to take a certain evaluative stance towards p. Approving is evaluative in nature, but not (necessarily) in content. In the examples of reasoning above, there is nothing evaluative about that which you approve of. We can describe your behaviour from the first-person point of view in fully non-evaluative terms.[11]

Approving is not to be understood in cognitivist terms; it is not believing that your behaviour meets a certain standard, that it is good, correct, or fitting. You may have no such belief in reasoning. Neither is the approval simply a feeling or an emotion. The evaluation is more substantive than that: it is intentional. When you reason, you choose to approve of the transition in thought. It is because of this intentional approval that the transition counts as an instance of active reasoning (see **3.8**). Our emotional reactions are not intentional in the same way, and simply liking some mental or bodily change does not make it something you do. (That you like getting goose bumps does not make the process which causes them something you do. You are still passive with respect to the process that brings about the goose bumps.)

To count as reasoning, the attitude of approval must be a special kind of approval. For example, you do not count as reasoning when you approve of a transition because it is, say, pragmatically beneficial, aesthetically pleasing, displays spontaneity or courage.[12] The approval must be an evaluation of the transition as you experience it. This is, in effect, an evaluation of the relation between the first-person premises and the conclusion-attitude. The attitude of approval manifests this kind of evaluation, when based on a certain set of attitudes.

[11] Sometimes we reason on the basis of beliefs about what we ought to do, what is good, just or right to do. In those cases, what you approve of is normative in nature, of course. The point here is that the content of the approval need not be normative.
[12] I thank Niko Kolodny for pressing me on this point.

This basis forms the standard or norm against which you evaluate your behaviour. In a loose sense of the word, you measure your behaviour against this standard. But the attitude of approval is not relational, and you do not approve of your behaviour *relative to a certain standard*.[13] You simply approve of your behaviour, where this approval is based on a certain set of attitudes and manifests your agential claim of the process.

Because the attitude of approval is an evaluation of your behaviour against a certain standard or norm, it counts as normative in the pre-theoretic sense of the word. As Wedgwood (2018) reminds us, dictionaries define 'normative' as that which involves or pertains to some norm or standard. There are different kinds of norms or standards, and the kind of norm or standard that something normative pertains to determines the sense in which it is normative. The question before us is therefore: What kind of norm or standard does the evaluative attitude of approval, which is essential to reasoning, pertain to? With that, we ask again the question raised at the end of chapter 2: In what sense is reasoning normative?

In chapter 2, we have seen two possible answers to this question. On the one hand, that which is normative may be understood as that which is related to what you ought to do or believe, or what is good or what is right. This is the sense which has dominated ethics (Wedgwood 2018).[14] Let's say that what is normative in this sense is genuinely normative. If we say that reasoning is genuinely normative, then we say that something genuinely normative plays some essential role in each and every instance of reasoning. The standard against which you evaluate your behaviour with the attitude of approval is set by genuinely normative properties, relations or facts. On the other hand, something may be said to be

[13] I have objected to Broome's account of reasoning that it requires the attitude of "seeming correct relative to a rule, which is relational in a complex way in which the attitude of approval is not.
[14] Whether the "normativity of all that is normative consist in the way it is […] related to reasons" (Raz 1999: 27) or whether it consists in the way it is related to 'ought' is irrelevant here. See note 27.

normative in the sense that it has some standard of correctness. The spelling of a word may be correct, but that does not necessarily imply that you ought to spell the word in that way. A move of chess may be correct, but you may have no reason to make that move. If we say that reasoning is normative in this sense, then we simply say that there is correct and incorrect reasoning. On a popular version of this view, the correctness of reasoning may be defined by the rules of reasoning, which do not make reference to claims about genuine normativity (see chapter 1). The norm or standard that the approving pertains to would be the one set by the correct rules of reasoning. Is the standard that the evaluative attitude of approval pertains to that of rules or that of genuine normativity?

3.3. The Normativity of Reasoning

I argue that the attitude of approval pertains to genuine normative standards. In short, the proposal can be summarised as follows.

The approval that is essential to reasoning is a special kind of approval: it is based on your implicit conception of what constitutes having sufficient reason to hold the conclusion-attitude. That is, your implicit conception of what makes it the case that a certain response is permissible. In short: it is your conception of permissibility-making. This is a relation of constitution or ground: some things (facts, or property instantiations) make it the case that and explain why there is sufficient reason for you to respond to them in a certain kind of way.[15] Suppose, for example, you reason to the belief that q from believing p and r. You approve of your coming to believe that q because of p and r. You count as reasoning if and

[15] I use the grounding relation or 'making it the case' - relation here as a placeholder for an explanatory relation, the investigation of which is a different project, see e.g. Frankfurt (1971) and Fine (2012). I do not commit to a particular view about the relata of the grounding relation, which may be facts, properties, or property instantiations. Because it seems to best fit ordinary language, I state my grounding claims in terms of facts: some facts ground other facts. Here some facts ground a normative fact, namely that it is permissible to believe or do something.

only if this approval is based on your implicit conception of what constitutes having sufficient reason to believe q.

To approve of a mental transition in the way that constitutes reasoning commits you to believing that *if the first-person premise were all that was relevant* to whether or not q is true (or to whether or not you ought to ϕ), they would at least partly make it the case that you have sufficient reason to take q to be true (or ought to ϕ).[16] This is a belief that a relation holds between the premises and the normative status of your response, on the condition that the premises are all that is relevant. The normative status of the response is that there is sufficient reason for you to form it (which is to say that you have sufficient reason to form it). The relation is that the premises at least partly make it the case that this is so. The inferential relation is therefore one of constitution or ground: it is about what grounds a permissible response to the world as we take it to be. In reasoning we track this relation, because we are guided by our conception of what it consists in. (Or at least we seem to track it. It may be an illusion that there is something to be tracked.)

I begin by explaining the claim that reasoning, i.e. the approval of your behaviour, rationally commits you to certain normative beliefs. Specifying the contents of the normative beliefs elucidates the evaluative nature of the attitude of approval itself.

The normativity of approving, and thus of reasoning, is genuine normativity. What guides your reasoning is your conception of normative reasons and how they aggregate to make a response sufficiently supported by reasons, i.e. permissible. You do not need an explicit representation of these normative properties and their grounds in order to reason. An

[16] Harman (1986: ch 5) makes a similar statement, albeit pertaining to theoretical reasoning only: "When one infers that something is so, one is committed to a claim that there is no significant chance that one's conclusion is false, given one's reasons for it."

implicit conception suffices. **Section 3.4** explains what this conception consists of. **Section 3.7** offers further support for the claim that reasoning essentially involves genuine normativity.

Commitments to Normative Beliefs

Reasoning rationally commits you to normative beliefs in the following sense: it is structurally irrational to reason in a certain way while holding certain normative beliefs. Suppose you reason as in the example of modus ponens reasoning about Blackwell's. Now suppose you deny that the fact that it is 5pm together with the fact that if it is 5pm then Blackwell's is closed gives you any reason to believe that Blackwell's is closed. We could criticise you on rational grounds.[17] Something seems incoherent. It comes close to a Moorean paradox to reason from p to q while believing that p is no reason to believe q. Searle (2010:10) argues that speech acts are ipso facto expressions of their corresponding intentional mental state: "It is 10pm!" expresses the belief that it is 10pm. Similarly, instances of reasoning are expressions of commitments to normative claims.

One might object that reasoning can only commit you to *not* having certain normative beliefs, just like believing that p commits you to not believing that not-p. Modus ponens reasoning rationally requires you to not believe that p and if p then q are no reason to believe q, for example. You are fully rational in so far as you are agnostic about reasons for believing q. That is correct. Being a reasoner does not require you to employ deontic concepts like that of a reason or ought. Nevertheless, if you had to make a judgement about whether or not the (supposed) facts that p and if p then q give you reason to believe q, you would, in virtue of your reasoning, be rationally required to either revise your reasoning or judge that these facts give you, at least under certain conditions, some reason to believe q. You are subject to a

[17] In section 3.9, I argue that this is so even for instrumental reasoning.

wide-scope requirement: rationality requires of you that [if you reason from A(*a*) to A(*c*) then you do not believe that *a* is no reason to hold A(*c*)]. Thus, being a reasoner rationally commits you to certain normative beliefs in the sense that you are under the jurisdiction of such wide-scope coherence requirements. You are subject to criticism on rational grounds if you reason in certain ways and at the same time hold certain normative beliefs. Importantly, reasoning does not require you to attend to or employ any normative beliefs. This would implausibly over-intellectualise reasoning (and thus go against D3 of the desiderata listed in chapter 1).

Having Sufficient Reason

So, what is the content of those normative beliefs that reasoning commits you to? It is not only that at least some premises are reasons for the conclusion, as suggested in the previous section. It is the belief that premises (at least partly) make it the case that you have sufficient reason for the conclusion-attitude, if they were all that was relevant. In this section, I explain the normative status that reasoners attribute to the conclusion-attitude: that they have sufficient reason for it.

Let's consider theoretical reasoning first. Reasoning from some beliefs to the belief that p commits you to believing that the premises make it the case that there is sufficient epistemic or truth-related reason for you to hold the conclusion-attitude.[18] What is it to have sufficient epistemic reason to believe something? Facts that give you sufficient epistemic reasons to believe that p are facts that have a certain epistemic property, for example,

[18] Theoretical reasoning does not commit you to any beliefs about pragmatic reason for belief. Reasoning is responding to truth-related reasons, as I will argue in chapter 4. We respond to pragmatic reasons for belief differently, not through reasoning on the basis of those reasons to the beliefs which they are reasons for. Pragmatic reasons to believe that p are not weighted against epistemic reasons for belief (see Reisner (2008)), but against reasons to have a true belief about p, see note 21.) Pragmatic reasons to believe that p can be thought of disabling conditions for epistemic reasons for believing not p.

a) they are sufficient evidence that p (i.e. evidence that surpasses a certain threshold), but not sufficient evidence that not p,[19] or

b) they entail that p but do not entail not-p, or

c) they ground that p (like the fact that the rose is crimson grounds the fact that it is red).[20]

That a), b) or c) is true of a set of facts is however, not sufficient for these facts to be sufficient epistemic reason for you to believe that p. This would be implausible, for we do not have any reason, let alone sufficient reason, to believe everything that follows from our beliefs. Nor do we have necessarily sufficient reason to believe what we have sufficient evidence for.[21] We should avoid cluttering our minds with trivialities (Harman 1986:12). Instead, facts that have certain epistemic properties are epistemic reasons to believe that p only if you have sufficient practical reason to have a true belief about p.[22] Epistemic reasons are conditional on value-driven reasons to have true beliefs. You have sufficient epistemic reason (or ought) to believe that p only if you have sufficient value-driven reason (or ought) to have a true believe about p.

Now, practical reasoning to the intention to ϕ commits you to believing that the premises make it the case that you have sufficient reason to ϕ. In the sense in which I use the

[19] If you have sufficient evidence to believe that p and sufficient evidence to believe that not p, then you ought to withhold judgment, which is different from merely not believing not-p and not believing p. Friedman (2013) explain in detail why suspension of judgment is, or at least involves, a genuine attitude which cannot be reduced to the absence of belief. Some deny that you can at the same time have sufficient evidence for p and sufficient evidence for not p (e.g. Lord 2017), but I allow for this possibility.

[20] We could also formulate this condition in terms of an epistemic property that the belief that p has if it is based on beliefs in the respective facts (for example, the content of the belief is entailed by the contents of the beliefs it is based on, etc.).

[21] Having sufficient evidence for p is not the same as having sufficient reason to believe that p, see also Leite (2007). I am assuming here a non-normative sense of evidence. See Kelly (2007) for an objection to this.

[22] See Steglich-Petersen (2011) for a discussion of reasons to have beliefs about p. In contrast to him, I do not believe epistemic reasons to believe p are hypothetical instrumental reasons. The respective facts are reasons *conditional* there being sufficient reason to have a belief about p. These need not depend on the agent's aim, but can be given by, for example, the fact that you already have a false belief about the matter, the fact that we argue about it, or the fact that it is your duty as a citizen to have a belief about it. I thus avoid the objections raised by Kelly (2003).

term here, to have sufficient reason to perform a kind of action ϕ does not necessarily imply that you ought to ϕ, though it may be that you ought to ϕ.[23] To have sufficient reason to ϕ is for it to be permissible to ϕ. I am here following Parfit (2011a:33), who writes that "our reasons to do something are sufficient when these reasons are not weaker than, or outweighed by, our reasons to act in any way of the other possible ways."[24] Permissibility can thus be understood in terms of ought: it is permissible to ϕ if and only if it is not the case that you ought not to ϕ. If you ought to do something, I assume, you are also permitted to do it.

In reasoning, you do not commit to believing that how you respond is how you ought to respond. How you reason may only be one of the possible ways in which you could correctly reason on the basis of your premises. How you ought to respond, all things considered, is a different question. It may even be that you ought not to respond at all. Maybe you ought not to ponder about the premises but think about something entirely different. Maybe you ought to derive another conclusion from them.

But, one might object to this, reasoning seems to involve not merely the recognition that there is some support for the conclusion-attitude. It seems to involve the recognition that the conclusion-attitude is in some sense *required* by the premises. This idea runs through most of the prominent accounts of reasoning. For example, Boghossian (2014:5) argues that "inferring necessarily involves the thinker *taking* his premises to support his conclusion and drawing his conclusion *because* of that fact." The support relation he has in mind is plausibly a strong sort of support that is "as strong as implication" (Broome 2019). Broome argues that in reasoning, your behaviour seems *correct* to you relative to a rule, not just good enough.

[23] Skorupski (2010:38) uses "sufficient reason to ϕ" in this sense, meaning that you are required or obliged to ϕ.
[24] As Dancy reminds us, this is a misuse of the word 'sufficient', for it is used as a comparative notion, not one referring to the passing of an absolute threshold.

The conclusion-attitude is required by the rule you follow (even if you are not required to follow that rule). Pettit (2018:48) argues that reasoners "think the truth of the premises which they believe as *forcing* them to acknowledge the truth of the conclusion" [emphasis added].[25]

The *Approval Account* captures this idea for theoretical reasoning also for practical reasoning. It implies that if the premises make it the case that you have sufficient reason to believe p, then, insofar as you form some belief about p, you ought to form the belief that p, rather than the belief that not-p. The premises do not require you to form the belief that p in the strong sense that you ought to reason to the belief that p, rather than a belief about q. The sense in which the premises require believing p is explained by the fact that they require you to form the belief that p, rather than the belief that not-p. If you form any belief about p, it ought to be the belief that p.

Also, in practical reasoning, the conclusion seems required in a certain way. The premises don't just give me some reason to do this rather than that. We can explain this impression as follows: first, in many cases, the premises, if they were all that was relevant, would indeed make it the case that you ought to ϕ, because they make it the case that you have sufficient reason to ϕ but not that you have sufficient reason to Ψ. If this is not the case and the premises give you sufficient reason to ϕ and sufficient reason to ϕ, then we have a case of equipollence: you have equal reason for both. If you now form one or the other intention on the basis of the premises, there is no sense in which you are required to choose this course of action over the other option. But neither should there be, if you have equal reason for both (see also chapter 5). In sum, we can capture the thought that in reasoning, we

[25] Here is a difference between mine and Dancy's (2018: 34) view. He argues that when you reason well from some beliefs to some response R, your responding with R is of the sort most favoured by the considerations rehearsed, taken as a whole. Reasoners take their response to be the one most favoured. I think that's too weak: a certain response might be the one most favoured by the considerations adduced, while being very weakly favoured. Suppose a set of considerations favours believing p more than it favours any other belief, while it gives only a very weak, certainly not sufficient, reason for believing p. According to Dancy's theory, the reasoning to the belief that p would nevertheless be correct, which is implausible.

seem to be forced to draw this rather than that conclusion. Yet, reasoning commits us only to believing that if the premises were all that was relevant, they would make it the case that we have sufficient reason to respond as we do.

I here seem to introduce a difference between practical and theoretical reasoning. Theoretical reasoning commits you to believing that the premises make it the case that you have sufficient truth-related reason to hold the conclusion-attitude (the belief). Practical reasoning, however, commits you to believing that you have sufficient reason to *do* what the reasoning leads you to intend to do. The reasons are not about the conclusion-attitude (the intention) itself. This apparent difference between practical and theoretical reasoning should not irritate us. We can resolve it by noting that if you reason to the intention to ϕ, and the premises (at least partly) make it the case that you have sufficient reason to ϕ, then you have sufficient reason to intend to ϕ. In general, when you reason to the intention to ϕ, your premises are reasons to intentionally ϕ.[26] I said practical reasoning commits you to believing that the premises make it the case that you have sufficient reason to ϕ, but it would have been more precise to say they make it the case that you have sufficient reason to *intentionally* ϕ. Intending to ϕ is part of what it is to intentionally ϕ; intentionally ϕ-ing is partly constituted by the intention to ϕ, where this intention is what Searle calls an intention-in-action. Reasoning, I have assumed above, can only lead to an intention, where this intention is either an intention-in-action or a prior intention that in some way leads to an intention-in-action, if the time comes and all goes well. If reasoning leads to an intention-in-action, the premises are reasons for that intention

[26] This does not exclude the possibility that, for some ϕ, you have sufficient reason to ϕ but not sufficient reason to intend to ϕ. You ought to keep breathing, for example, but you have no reason to intend to keep breathing, for you will do it anyway. But you also cannot reason to your breathing. You can reason to intentionally breathe in a certain way, on the basis of considerations which are reasons to intentionally breathe in this way.

because they are reasons to intentionally ϕ. If reasoning leads you to a prior intention, say the intention to call your mother tomorrow, then this action is presumably up to you in the sense that you would not ϕ unless you intended to ϕ, so that the reasons to ϕ are also reasons (of a somewhat instrumental nature) to intent to ϕ. In any case, therefore, practical reasoning, just as theoretical reasoning, commits you to believing that the premises make it the case that you have sufficient reason to respond as you do.

Given the Premises

I have claimed that an instance of reasoning commits you to a belief about the relation between the premises and your response on the condition that "*the premises were all that was relevant*." Let me explain this qualification.

First, note that reasoning does not involve your approval of the conclusion-attitude all things considered. It does not commit you to believing that you have, all things considered, sufficient reason for your conclusion-attitude. This is plausible, since it may not be the case that you ought to hold that attitude, even on your own standards. (You may even believe that you ought not overall hold it.) You may at all times be ready to revise the conclusion-attitude as well as your reasoning if you became aware of additional relevant considerations, and you may even at all times believe that you will probably do so, taking it to be likely that you missed some or the other relevant aspect of the situation. However, when you do respond to a set of premises with an attitude, say the belief that q, and approve of this transition, you manifest a belief that the premises are sufficient ground for you to respond in some or the other way. If you believed that something was missing that would make a difference to how you are permitted to respond, you would be hesitant and not approve of the transition to any conclusion. You would search your memory and call to mind further facts or gather more information about the situation. If you draw a conclusion, however, you manifest a belief that

you took into account all considerations that are relevant to whether or not you have sufficient reason to believe p (i.e. ought to believe that p).[27]

This belief may or may not be the result of a reasoning process. The gathering of evidence and collection of relevant considerations, as well as their assessment as sufficient, may happen entirely sub-consciously and habitually. Your readiness to form this belief, i.e. your readiness to draw a conclusion, depends on the perceived stake and complexity of the matter you reason about, the sense of urgency to respond and the cost of searching for more information. You may be ready to give it up easily, if just one more supposedly relevant aspect of the situation enters your sphere of attention.

This belief, that the premises are sufficient for you to form some belief, is an implicit premise-attitude of the reasoning, a background assumption, if you like. It plays a different role than what we may now call the explicit premise-beliefs, the contents of which you think about in reasoning. I will say more about this distinction in reasoning in section **3.5**. The point here is that because the conclusion you draw is premised on this implicit assumption, your evaluation of the transition is conditional, and the belief that the approval of this transition commits you to is not the belief that, all things considered, you have sufficient reason to believe q, but the belief that if *the premises were all that was relevant*, they would make it the case that you have sufficient reason to believe q.

Tracking what makes what permissible

Your reasoning rationally commits you to believing that the premises, at least partly make it the case that you have sufficient reason to respond as you do, that is, that it is permissible to

[27] Kauppinen (2018) argues that „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." It will become evident in chapter 4 that, on my view, while *reasoning* does require the belief the that you have considered all that is relevant, reason-responsiveness does not. You can respond to a consideration as a normative reason, without believing that you have considered all that is relevant.

respond to the situation in the way you do. To approve of your behaviour is thus to assess a certain relation between the first-person premises and the normative status of the conclusion-attitude. Again, the first-person premises here are the (supposed) facts you reason from, assuming you reason on the basis of beliefs. These facts at least partly make it the case that the conclusion-attitude is sufficiently supported by reasons. This is a claim about a grounding relation, or a relation of constitution. It is a relation about the nature of things, how they are in the world; the relation of 'making it the case'. If some facts ground or make it the case that p, then we can explain p by pointing to those facts. The grounding relation corresponds to an explanatory relation we rely on in our practices of explaining why things are so and so. In the context of reasoning, the relevant grounding relation is that of "making it the case that it is permissible for you to believe that p or to ϕ." The (supposed) facts you reason from (i.e. your first-person premises) ground your conclusion-attitude's having a certain normative property, namely the property of being sufficiently supported by reasons. They explain why it is permissible for you to respond as you do.

If you have sufficient reason or ought to believe p (or to ϕ), we can explain why that is so. Some of the facts which make it the case that you have sufficient reason or ought to believe that p (or to ϕ) are reasons to believe that p (to ϕ).[28] Not all facts are reasons, some are the conditions or modifiers of reasons. The relation of being sufficient reason to believe that p (or to ϕ) is not the relation of being a reason to believe p or to ϕ. What you have sufficient reason to believe or do depends on the reasons you have and how they aggregate.

[28] We can accept this and stay neutral on the question of whether reasons or 'oughts' are the fundamental features of normativity. On the one hand, one may hold that the facts in question are reasons because they (at least partly) make it the case that you ought to ϕ or believe that p. This is a version of the view that the property of being a reason can be reduced to "ought", which e.g. Broome (2015b; 2013: ch.4) defends. Kearns and Star (2009; 2013) offer an alternative theory of how reasons reduce to "ought." On their view, a reason for you to F is evidence that you ought to F. On the other hand, one may hold that the facts make it the case that you ought to ϕ or believe that p because they are reasons. This would be to defend the primacy of reasons (see e.g. Scanlon 1998: ch 1; Raz 1999:67; Dancy 2004; Schroeder 2007: 81; Skorupski 2010; Raz 2011:85; Parfit 2011a: ch. 1).

The relation that guides reasoning is not that of full grounding. The premises may merely be partial grounds, and thus partly make it the case that you have sufficient reason for the conclusion-attitude. In reasoning, we only pick out some salient features of the situation, as we do when we explain phenomena. For simplicity, I do not always restate the qualification that the first-person premises *at least partly* ground the permissibility of the conclusion-attitude. The central point is that the standard against which you evaluate your behaviour is set by your implicit conception of the grounding of the permissibility of certain beliefs or intentions.

In sum, I have argued in this section that reasoning – more precisely, approving of a transition in thought - rationally commits you to believing that the first-person premises, if they were all that was relevant, would make it the case that you have sufficient reason to respond as you do. These normative beliefs shed light on the evaluative nature of the attitude of approval. In chapter 5, where I develop an account of the correctness of reasoning, we will see that these normative beliefs that reasoning commits you to are correct whenever the reasoning is correct. Indeed, it would be implausible if the account entailed that correct reasoning commits you to false normative beliefs. The attitude of approval, which I argued is essential to reasoning, is based on your implicit conception of what makes it the case that you have sufficient reason to do or believe something, i.e. your implicit conception of normative reasons and how they aggregate to make a certain response permissible or required. I have said little about the nature and content of these implicit conceptions, so I shall turn to this next.

3.4. An Implicit Conception of Normative Reasons

What is the nature of this implicit conception? We can understand an implicit conception as an attitude or set of attitudes that an agent holds towards an object, so that the contents of these attitudes constitute a particular characterisation of the object.[29] Implicit conceptions may play different roles in our psychological lives. For example, implicit conceptions are arguably influential in the understanding, learning and application of concepts. Peacocke writes:

> "In mastering the concept 'cube', taken as an observational concept, a thinker must have an implicit conception with a content which includes this: that cubes are closed figures formed from square sides joined at right angles along their edges. Not all examples will be so trivial. In the case of any philosophically interesting concept, the question of the content of the implicit conception underlying it will be highly substantive" (Peacocke 1998: 52).

On this view, one's implicit conception of something influences judgments about that thing. Importantly, "the implicit conception does not explain the judgement under a merely syntactic description", but under its content-involving description (Peacocke 1998: 69). Implicit conceptions influence judgements at a sub-personal level, for example as a result of perceptual information being combined with background information and the content of the implicit conception. The content of the implicit conception is not explicitly represented. The respective beliefs are implicit; they may thus not be directly available to consciousness and only with effort retrievable. But we can try to articulate the implicit conceptions that influence our judgements. It requires skill and reflection on a well-chosen set of judgments. Some

[29] In contrast to Bengson (2018), I do not assume that a conception of an object X necessarily consists in having one particular attitude towards X (such as the attitude of 'conceiving of' something in a certain way).

implicit conceptions are easier to make explicit than others. It is much more straightforward to articulate the implicit conception involved in the understanding of the concept 'cube' than to articulate the implicit conception involved in our judgements about what is fair, for example. As Peacocke notes, a significant portion of moral and political philosophy is concerned with making our implicit conception of notions such as 'fair' explicit (see quote above). There are several advantages of making our implicit conceptions explicit (Peacocke, 1998: 52-53). It allows us to gain a general statement, a definition of some sort, and to explain and justify our judgements in particular cases to ourselves and others. Note that the content of the implicit conception related to the concept 'cube' is about that which makes an object a cube – that it has six equal-shaped sides, etc.

Analogously to the observational concept 'cube', I suggest, our implicit conceptions of 'normative reasons' are sets of beliefs the contents of which give a particular characterisation of that which makes something a reason. We can manifest this implicit conception without making an explicit judgment involving the concept 'normative reason'. Imagine the following conversation:

> "We should go to the dinner party tonight. There will be good food, and we can see Luna again."
>
> "But it will rain on the weekend!"
>
> "Dhhhh?"

The bewilderment expresses the thought that the rain next weekend has nothing to do whether or not we should go to the dinner party tonight; it is not a reason not to go. It expresses a judgement about normative reasons, without explicitly using the concept 'normative reason'. This judgement about normative reasons, I argue, is based on an implicit conception of what

makes a consideration a reason to do something in a certain context just as a judgement about cubes is based on one's implicit conception of what makes an object a cube.

The example of the dialogue shows that one can manifest one's implicit conception of normative reasons without making use of the concept 'normative reason'. What is true of arguments between people here is true of arguments with oneself, too. In reasoning, we manifest our implicit conceptions of normative reasons. For when one reasons, and approves of the transition in thought, this approval is based on one's implicit conception of normative reasons. Reasoning is a way of expressing our implicit conceptions of normative reasons. An explicit belief that some consideration is a reason is neither necessary nor sufficient for reasoning. (In comparison, to count as mastering the concept 'cube', it is not sufficient to believe that some object is a cube.) Manifestation of a conception of normative reasons is manifesting a conception of facts as reasons *to do* or *to believe* something, for example. One sees a fact as standing in a certain relation to certain responses (and to oneself). Because normative reasons have this relational character, manifesting a conception of normative reasons is not necessarily judging a fact to be a reason, but responding to a fact with another attitude that is not about the fact responded to, but that which it is a reason for.

The question before us now is: what is the content of this implicit conception of normative reasons that is necessary for reasoning? I said above it gives a particular characterisation of that which makes something a reason. So what makes a fact a normative reason? Answering that question is making explicit our implicit conception of normative reasons. I shall not give a substantive answer here. However, I shall make some remarks about the structure of the answer. In chapter 4, I explain in greater detail what the content of our implicit conceptions of normative reasons is. For now, the following rough sketch shall suffice.

The assumption is that if some consideration is a normative reason for N to ϕ or to believe that p, then there is something which makes it so. Therefore, if a fact is a certain kind of reason in some context, then we can explain why that fact is this kind of reason.[30] To begin with, note that the full explanation of why f is a reason for N to ϕ is not identical with the reason-constituting fact f (Bader 2016). What we may call the necessitating base (the reason-constituting fact, the modifiers and the conditions) makes it the case that f is a reason for N to ϕ. Further, we can explain this ‚making it the case'– relation. We can explain why certain facts are (conditional) reasons. That is what we are interested in here.

First, consider reasons for action. A reason for a certain kind of action is something that makes the action right, at least to some degree. To say that the fact that she is drowning is a reason to jump in the pond is to say that that fact makes it right to jump in the pond, at least to some degree. The question of why a fact is a reason has then to do with what kinds of acts are right. Put roughly, a fact is a reason for a certain course of action because it makes that action one that is right. There are different views about the rightness of acts. On the view that the rightness of an act is determined by the goodness of the outcome, which Rawls (1971: 22) calls a teleological view, a reason to do something makes an action right to some degree because it makes the performance of the action good to some extent. An explanation of why some fact is a reason, then involves claims about the goodness of states of affairs.[31] But we

[30] This, I believe, is compatible with the view that the concept of a reason is basic and not analysable in terms of other concepts. That a predicate is not analysable in terms of other predicates does not imply that instantiations of the property that the predicate picks out cannot be explained.

[31] Though I stay neutral on the question of whether the good is prior to the right or the right prior to the good, it I deny at least some elements of what Scanlon (1998: 97) calls the buck-passing account of value. My argument about reasoning, and my argument about the psychology of responding to normative reasons, can be seen as arguments against Scanlon's conception of the metaphysics of normative reasons. When he says that "to be good or valuable is to have other properties that constitute such reasons", he seems to exclude the possibility of explaining why a fact is a reason to do something by reference to the value of the act (or its outcome). However, I do agree with him that "it is not goodness or value itself that provides reasons but rather other properties that do so." That something is good is not a reason. Instead, as Scanlon says, "the reasons are provided by the natural properties that make a thing good or valuable." I also agree with his claim that "these natural properties provide a full explanation of the reasons we have", if understood as the claim that all reasons are natural facts. Interestingly, Scanlon (1998:23) also says that "Every action that we take with even a minimum of deliberation about what to do reflects a judgement that

need not accept this theory of the right, and may hold that the right is determined independently of what is good. We can thus say that the content of an implicit conception of normative reasons for action includes beliefs about rightness or beliefs about goodness.

Next, consider reasons for beliefs. If a fact is an epistemic reason to believe that p, then we can explain why that is so. One part of this explanation is that the fact indicates the truth of p (where the notion of 'truth-indication' needs elaboration). For example, the fact may entail the truth of p, probabilify p, or make it the case that p. The content of an implicit conception of normative reasons for beliefs thus includes beliefs about entailment, evidence or probabilification (I borrow the term from Dancy, 2018).

In sum, I have argued above that in reasoning, you approve of a transition in thought, where this approval is based on your implicit conception of what makes what kind of response permissible. This conception consists in a set of beliefs, the contents of which give a particular characterisation of permissibility-making: they describe that which makes it the case that some facts make a certain response permissible. (This is analogous to an implicit conception of what a cube is, the content of which describes that which makes something a cube.) The facts which make a certain response permissible are reasons (and their conditions and modifiers). Therefore, a conception of permissibility-making is a conception of normative reasons and how they aggregate to give you sufficient reason to believe or do something. It consists of beliefs which characterise what makes a fact a normative reason, and what makes a response permissible. Throughout the discussion, I have assumed that this conception is implicit. But it need not be. The beliefs which make up this conception may be explicit beliefs

a certain reason is worth acting on", where these judgments resemble Frankfurt's (1971) "second-order desires." The judgement that a reason is worth acting on is, I take it, the judgement that a consideration is worth acting on, i.e. that a consideration makes it worth to act in a certain way. This sounds very much in line with my proposal.

about what makes a fact a reason for you to do or believe something, and what makes it the case that you have sufficient reason to do or believe something. It would be thus more precise to say that reasoning requires an *at least* implicit conception of normative reasons and how they aggregate.

The beliefs which make up one's implicit conception of normative reasons are essential to reasoning because they are the basis of the approval in virtue of which one counts as reasoning in the first place. I will now explain what their role is in the reasoning process.

3.5. Different Kinds of Premises

I have appealed twice to the role of implicit premise-attitudes in reasoning. First, I argued that when you reason, the conclusion of our reasoning is premised on the implicit assumption that the considerations you adduce are sufficient to respond in one way or the other. Second, I claimed that the attitudes which constitute your conception of normative reasons are implicit premises-attitudes.

The recent debate on reasoning almost exclusively deals with extremely simplified examples of instances of reasoning. It is assumed that any instance of reasoning can be neatly presented by listing the premise-attitudes and the conclusion-attitude. When we draw a conclusion, we do so on the basis of a clearly discernible set of attitudes, as if an instance of reasoning took place in a corner of the mind that hosts all and only the attitudes constituting that instance of reasoning. We further assume that the attitudes in that space play only one of two roles: they are either premise-attitudes or they are the conclusion-attitude. This simplification may be misleading. In addition to the premise-attitudes on which the conclusion-attitudes is based, which I call the explicit premise-attitudes, we must allow for implicit premises-attitudes which play a different role in the reasoning process.

The distinction between explicit premise-attitudes, implicit premise-attitudes and the attitude of approval is a functional one; they all play different roles in the reasoning.[32] First, the explicit premise-attitudes cause the conclusion-attitude 'in a right kind of' way (that is, non-deviant way), so that the latter is based on the former. They are part of the transition that the agent approves of. The contents of the explicit premise-attitudes are that which the reasoner thinks about and operates on in the reasoning. Second, the attitude of approval is not a premise at all, but an attitude towards the explicit premises as experienced by the reasoner and the conclusion-attitude. Third, there are implicit premises-attitudes. In contrast to the explicit premise-attitudes, the reasoner does not operate on the contents of these, which we may call the implicit premises.

One example of an implicit premise is the consideration that the explicit premises are all that is relevant to whether you ought to believe p or ϕ, which I argued is part of every reasoning process. This is a premise, but not one that the reasoner operates on.[33] Another example of implicit premises are the attitudes that together constitute the reasoner's implicit conception of normative reasons. These are the basis of the approval. Because the transition counts as an instance of reasoning in virtue of this approving, an explanation of the conclusion-attitude (as the conclusion-attitude of an instance of reasoning) involves not only the explicit premises but also the implicit ones. Nevertheless, the conclusion-attitude is not based on the implicit premise-attitudes in the same direct way in which it is based on the

[32] The distinction between explicit and implicit premise-attitudes is, I believe, very close to the distinction between foreground and background attitudes that Pettit and Smith introduce (Jackson, Pettit, and Smith 2004: ch. 13). In decision making, they argue, "a desire for S figures in the background if and only if it explains the agent's choice of options. It figures in the foreground if and only if the agent reaches the choice via the recognition that he has the desire and that the option has the desirable property" (p.171 – 172). On my view of reasoning, the agent does not operate on the implicit premises, they are in the background. The reasoner operates the contents of the explicit premise-attitudes. The reasoner recognizes that he has these attitudes.

[33] If it was, the consideration that the explicit premises, together with the consideration that they are all that is relevant, would have to be *another* implicit premise, leading to an infinite regress of implicit premise-attitudes.

explicit premise-attitudes. The agent holds the conclusion-attitude of an instance of reasoning in light of the explicit premises but not directly in light of the implicit premises. The *Approval Account* thus avoids the regress problem discussed in chapter 1 - that which guides the reasoning is not a premise of it.

Some exemplary instances of reasoning we discuss in philosophical inquiries are *enthymematic*: they feature more premises than stated (Broome 2013: 259). We say that the premises that are not relevant for the discussion at hand are left "implicit." This is a different sense of "implicit." Premises that are implicit in this sense may play the role of what I call explicit premises or the role of implicit premises, and it may also be indeterminate which role exactly they play. On my view, all examples of reasoning we discuss in philosophy are enthymematic.

Lastly, we must distinguish implicit premise-attitudes not only from explicit premise-attitudes, but also from what I call non-involved attitudes. The latter play no role in the reasoning process, they neither affect nor are directly affected by it. Are there such non-involved attitudes? An instance of reasoning does not involve all of the things the reasoner has in mind. It would be absurd to say that every conclusion you draw is, in some way or the other, influenced by all other attitudes you hold. Even though reasoning is a more encompassing mental activity than we acknowledge in our discussions, there are some beliefs and intentions of yours which play no role when you are reasoning about whether to pick the standard fries or the sweet potato fries. Maybe more attitudes play a role than we would ordinarily expect, but to say that all are involved seems unreasonable. I assume that for most instances of reasoning, there are non-involved attitudes.

In sum, the conclusion-attitude of an instance of reasoning is based on the explicit premise-attitudes, but neither on the attitude of approval, nor on the implicit premise-attitudes, even though both are essential to the reasoning process.

3.6. Meeting the Desiderata

So far, I have proposed the *Approval Account*, according to which reasoning is a thought process that the reasoner endorses in a certain way, they approve of it. The account is subject to desiderata **D1** – **D3**, as I defined them in the preceding two chapters. It should

D1 explain the content and nature of the attitude that is essential to and grounds reasoning

D2 show how the presence of that attitude makes the process count as *active* and explain its role in the reasoning process

D3 not render reasoning conceptually over-demanding, that is, not require reasoners to have complex concepts like that of a normative reason, belief, entailment relation, etc.,

Section **3.2** of this chapter has explained the content and nature of the attitude of approving and thus addressed **D1**. In reasoning, you approve of your behaviour as it appears to you. When you reason on the basis of beliefs, this is your acquiring of a belief or intention on the basis of some facts.

If the attitude of approval is of the right kind, it makes the process you approve of count as something you do. You intentionally derive the conclusion. I have not yet explained in detail how the approval imputes intentionality in the thought process, but will do so in section **3.9**, arguing that the *Approval Account* meets **D2**. The right kind of approval is that which is based on your implicit conception of what makes what kind of response permissible. Approving of a transition in thought is measuring it against the standard set by this conception. Reasoning does require the capacity for normative thought, the ability to

distinguish between an 'ought' from an 'is', but no sophisticated concepts like that of a 'belief', 'normative reason' or 'entailment'. None of what the *Approval Account* requires of reasoners seems in conflict with **D3**.

In chapter 1, I presented two prominent characterisations of reasoning, one in terms of normative reasons, and the other in terms of rule-following. Because both have high intuitive appeal, a fourth desideratum for a theory of reasoning would be to

> **D4** either accommodate or explain the mistake behind the *Rule-Following Claim* and the *Reasons Claim*.

Can the *Approval Account* meet **D4**? I have not defended the *Reasons Claim*, which states that reasoning is a way of responding to normative reasons for belief and action. Unlike the defenders of the *Reasons View*, I have not given an account of what reasoning is in terms of 'responding to reasons', or in terms of 'treating things as reasons', 'taking things to be reasons' or 'seeing things as reasons.' My account is thus not subject to the objection I raised against accounts of reasoning that appeal to these notions, namely that they are explanatorily insufficient until these notions are explained. Of course, according to the account, normative reasons play an essential role in reasoning. I have argued that reasoners are guided by their implicit conception of what constitutes having sufficient reason for a certain kind of response. It seems plausible that when you manifest this conception in reasoning, you count as treating the considerations which are the premises of your reasoning as normative reasons. If this is correct, the *Approval Account* entails that reasoning is a way of responding to normative reasons. But do you count as responding to reasons whenever you reason? For the rest of this chapter, especially the next section, I assume that the answer is 'yes'. However, this is a speculation. In chapter 4, I develop an account of what it is to respond to normative reasons.

This allows me to establish the claim that reasoning –approving of transitions in thought - is responding to normative reasons.

Guidance by Reasons as Following Rules

This leaves the *Rule-Following Claim*, which states that reasoning is a rule-following activity. How we understand the relation between reasoning and rule-following depends on what we take rule-following and the rules of reasoning to be.

Regarding the first, one might think this gets us straight to Wittgenstein's rule-following problem. But it is not obvious that the question of rule-following in the context of reasoning is the question of rule-following that Wittgenstein (2009: §144-§148) raised in his 'Philosophical Investigations'. According to Pettit (2016:3381), it is not. In response to Broome, he argues that the rule of modus ponens reasoning, for example, is "complex in the sense that following it presupposes following more primitive rules governing when you ought to assent to 'if p then q' and 'p'." Wittgenstein's rule-following problem, however, "arises sharply with these primitive, indefinable rules that you are required to follow in the course of following a more complex rule like modus ponens." The question of what the rule-following activity of reasoning consists in, Pettit concludes, is thus not Wittgenstein's question of basic rule-following. Indeed, many of Wittgenstein's remarks on rules and rule-following in the 'Philosophical Investigations' make me doubt that his notion of rule-following is the one that the defenders of the *Rule-Following View* of reasoning are after. For example, Wittgenstein says that "a rule stands there like a signpost" (§85), in need for interpretation. But the rules of reasoning cannot be in need of interpretation, at least not if that interpretation requires reasoning. In some places, it seems, Wittgenstein is concerned with the kind of basic rule-following which Pettit refers to, the kind that is required for reasoning. At other places, Wittgenstein seems to be concerned with rule-following that requires reasoning. In sum, the

relation between rule-following in reasoning and Wittgenstein's rule-following problem remains to be established.

Let us put this aside for a moment and look at the rules themselves. The question at hand is whether the *Approval Account* can accommodate the thought that reasoning involves rule-following. How could we understand the rules of reasoning on this account? The "rules" we are guided by in reasoning are our heuristic approximations of what makes what kind of response permissible. They are rules of thumb about what is the right thing to do or believe, given certain considerations. We are guided by these rules, because our conception of them determines which transitions in thought we approve of, i.e. how we reason. I explain more in what sense our conception of normative reasons guides us in section **3.8**. The normativity of reasoning is genuine normativity. But our implicit conceptions of the grounding of the permissibility of certain kinds of responses can be understood as setting up rules which guide us in reasoning. Furthermore, we must notice that some instances of reasoning that we describe as rule-following are cases where we use reasoning to follow rules. We reason, for example, when we follow the rules of an operation's manual, the rules set by a recipe, or the rules of making a tax statement. These are not rules of reasoning, but rules that we reason about in order to follow them.

I cannot pursue the question of rule-following further here. But I hope to have illustrated that the *Approval Account* can capture the intuition behind the *Rule-Following View*. It can do so by denying that the rules of reasoning are rules of coherence or rationality, and that the sense in which we follow them is necessarily the sense in which we follow Wittgenstein's rules.[34]

[34]Pettit (2007a:243) makes similar suggestions, saying that treating reasons "as constraints on your performance will amount to following them in the sense in which you might be said to follow rules. You will have to understand or identify them, being able to recognise what they require over an indefinite range of cases.

3.7. Three Features of Reasoning

In sections **3.2** – **3.5**, I have argued that reasoning requires the evaluative attitude of approval, which pertains to genuine normativity. I now further defend the genuinely normative nature of reasoning by arguing that this view, but not its rivals, can explain why reasoning involves

i) a claim to universality across reasoners,

ii) the disposition to justify your conclusion-attitude by reference to the premises, and

iii) a claim to universality across instances over time.

The *Approval Account* can explain these three features of reasoning, precisely because it holds that the normativity of reasoning is genuine normativity. In contrast, the view that the normativity of reasoning is that of abstract rules of coherence or rationality cannot explain why reasoning seems to involve features i), ii), and iii). To make this argument, I compare the *Approval Account* to Broome's *Rule-Following View*. According to this view, what is essential to reasoning is that your behaviour seems correct to you relative the rule of reasoning you intend to follow (see chapter 1). The rules of reasoning derive from coherence requirements of rationality and do neither rely on nor entail any claims about normative reasons.

Let us first turn to i), the claim to universality across reasoners. If you reason and have an attitude of endorsement (on my account: the attitude of approval) towards your transition, you measure your reasoning against a standard you implicitly take to apply to all reasoners. It is not just something about your aims and desires, though these may play a role. You would approve of someone else's transition of this kind and expect them not to be perplexed if you expressed your reasoning out loud.

An explanation of this is available on the view I propose. Reasoning is based on an implicit conception of what makes it the case that you have sufficient reason to do or believe something. You conceive of these normative facts or laws not as something special about you, but as applying to any person with relevantly similar characteristics. The grounding of the permissibility of a response sets an interpersonal normative standard. If we reason from believing p to believing q, and thus take p to be a reason to believe q, for example, we take p to be a reason to believe p for anyone who is in a relevantly similar situation to us. If p makes it permissible to believe q for us, it makes it permissible for anyone who is relevantly similar, too.[35]

Equally, what makes it the case that you have sufficient reason to ϕ would also make it the case that someone else, in a relevantly similar circumstance, would have sufficient reason to ϕ. This is uncontroversial if the reasons in question are moral reasons. It is the Kantian idea that when we choose to act, we must make it our maxim to act in that way. Hume (1751: section 9.6.) expresses this thought when he says that someone who uses the language of morals to criticise others "expresses sentiments in which he expects, all his audience are to concur with him." Pettit (2018:197), in reference to Hume, says that what is morally desirable is what is multilaterally desirable, that is "desirable in a way that transcends the perspectives of different persons." Nagel (in Korsgaard 1997: 202) makes a similar point. He writes that "The reflective self [the one who thinks whether impulses comply with reasons] is in its nature more universal than the original, unreflective self, because it achieves its self-conscious awareness by detaching from the individual perspective." Scanlon (1998: 73) calls this the "universality of reason judgements." If you judge that g gives you a reason to ϕ, you must judge that anyone else who stands in the same relation to g and ϕ-ing has a reason to ϕ.

[35] Skorupski (2002:121) defends the same point in detail.

But the claim to universality across reasoners is not only involved in moral reasoning. It is also involved in prudential reasoning, which is not about what we owe to each other but what we owe to ourselves – to put it in Scanlon's terms. We make decisions on the basis of considerations about what is good for us, we expect others to understand and accept our reasoning just as we expect them to understand our reasoning on the basis of moral considerations - at least as long as we are not violating any moral demands. This expectation may be conditional in character. We may allow for the possibility that others disagree about what is good for us, which may be due to the fact that they have a different conception of what is good for them. But again, this is true of moral reasoning, too. Practical reasoning involves a claim to universality across reasoners who we expect to share our conception of moral and prudential reasons for action. The point is that when we reason and this involves a claim about what would be morally or prudentially right for us to do, we assume we can explain our decisions in ways we cannot if we merely act on our desires. In some way or another, you expect others to be able to re-enact your reasoning.

In sum, reasoning involves a claim to universality across reasoners because it is guided by our implicit conception of normative reasons, according to which those reasons are interpersonal – they apply to all persons under the jurisdiction of normativity.

Defenders of the *Rule-Following View* can argue that the standard against which you measure your behaviour, the correct rules of reasoning, applies to all reasoners. But this does not explain i), because it does not entail that a reasoner recognises this universal scope of the rule he follows. When you follow a rule and your behaviour seems correct to you relative to that rule, you may not conceive of the rule as applying impartially to all reasoners with relevantly similar characteristics. That your behaviour seems correct to you according the rule you chose

to follow makes reasoning a solitary exercise, wholly independent of other persons. There is no recognition that the rule also applies to others, no claim to universality across reasoners.

A second feature of reasoning is that it involves ii) the disposition to justify your conclusion-attitude by reference to your premises. Suppose someone reasons in accordance with modus ponens. If asked "why do you believe q?", she will be disposed to answer "because p and if p then q!." Her answer will not merely be an explanation of why she believes q. It communicates an attempt to justify her belief that q. Indeed, should she want to only explain and not justify her belief, she would have to make an extra effort. Analogously, when you reach a decision through reasoning and are asked why you do what you do, the premises of the reasoning seem to be what allows you to justify your decision. Reasoning involves a disposition to rely on the premises to justify one's conclusion-attitude to others.

The claim that the normativity of reasoning is genuine normativity can also explain this observation. With the attitude of approval, one evaluates the transition according to one's implicit conception of what makes what kind of response permissible. Because of that, one responds to premises as normative (or "justifying") reasons for the conclusion-attitude, which is why one is disposed to use them as justifying reasons when asked to defend one's conclusion-attitude.

The *Rule-Following View* has no such immediate explanation of the phenomenon. Why would a reasoner see his premises as justifying his conclusion-attitude? On the view that reasoning is following abstract rules of coherence, this seems a peculiar coincidence. To justify your conclusion-attitude is not only to show it to be in accordance with *some* jus, that it is in line with some rule or standard (Velleman 1996), but to show it to be in accordance

with a standard impartially applicable to all, where that is common knowledge. The former does not imply the latter.

A third feature of reasoning is that it involves iii) a claim to universality across instances and over time. The standards against which you measure your reasoning is taken to apply to all instances of reasoning, not just a particular one. If you reason in similar ways in similar situations, you approve of the respective transition on the same basis.

The defender of the *Rule-Following View* may hold that the rules of reasoning are general rules, like the general rule of modus ponens reasoning: "From p and if p then q derive q", which apply over time. I have argued in chapter 1, however, that if the rules of reasoning are general in this way, it is unclear how reasoners can follow them.

The *Approval Account* can accommodate iii) as follows. When you approve of a transition, you recognise that the standard according to which you evaluate the transition – the grounding of the permissibility of certain responses - applies across instances and over time. The features of the situation which make a certain kind of response permissible also make this kind of response permissible in a relevantly similar situation. This is just to say that if some fact is a reason, it is a reason in a relevantly similar context. Our conception of normative reasons involves this generality. We conceive of a reason as "an instance of a familiar, universal type", as Pettit (2018: 111) puts it. The view that the normativity of reasoning is that of normative reasons entails that reasoning involves a claim to universality across instances because, "when our choices to act are based on our conception of what we have reason to do, they will entail that we accept some general normative principle" (Wallace 2006:89).

I conclude that the view that reasoning is governed by rules that are independent of genuine normativity cannot account for the observation that reasoning entails a claim to universality across reasoners, a claim to universality across instances, and a disposition to treat the premises as inter-personally justifying the conclusion-attitude. The view that the normativity of reasoning is the genuine normativity of reasons and ought, however, can explain these features of reasoning. I take this to speak in favour of the *Approval Account* and for the claim that the normativity of reasoning is genuine normativity.[36] If we deny normativity, we deny ourselves the capacity to reason. Unless we believe that some things matter, there is nothing to guide our reasoning.[37]

3.8. The Intentionality of Reasoning

I now turn to desideratum **D2**. A theory of reasoning should

> **D2** show how the presence of that attitude makes the process count as *active* and explain its role in the reasoning process.

I have argued that your approval of the transition to the conclusion-attitude is essential to reasoning. The attitude of approval plays no role in the causation of the conclusion-attitude but is a person-level acceptance of the process leading to it. Now we may ask: How does this story explain why reasoning is an activity on the agent's part? When you reason, you do so intentionally, and because of that, the reasoning is an activity on your part. So the question

[36] My claim that reasoning is normative is to be distinguished from Wegdwood's (2006) claim about the normativity of reasoning. His account differs from mine in that it is a purely causal, more precisely, dispositional account of reasoning, and in that it relies on the causal efficacy of the normative.
[37] Gibbard (2002:159-160) reaches a similar conclusion, despite his non-cognitivist view. As deciders, he says, as agents who reason to intentions to do things on the basis of what they believe, we commit to there being a property about what we are to do.

is: what is the intention with which you reason? The *Approval Account* is meant to give an account of what it is to reason and must therefore explain what the intention is in virtue of which the transition counts as a mental action on the reasoner's part, rather than mere mental causation on a sub-personal level.

An "action requires some event or state that owns its occurrence to the agent" (Velleman 2000: ch.6). In the case of the mental action of reasoning, this is the agent's approving of his transition. The approving is that which is essential to and grounds the reasoning, and that which makes the process an activity of the agent's part. I do not say that the attitude of approving makes a causal contribution to the formation of the conclusion-attitude.[38] Rather, it changes the quality of the process.

Approving is a mental state, and not an action. However, assuming the attitude of approving is a mental act.[39] Approving of your behaviour is voluntary, and under your control. It is, to use Wittgenstein's understanding of voluntariness, "marked by the absence of surprise" (Philosophical Investigations, §628). In that respect, approving is different than mere liking, for example. In reasoning, you intentionally approve of the process, and in virtue of that, count as intentionally reasoning. What is the content of the intention with which you reason? You do not intend to form a particular conclusion, but you intend to form that conclusion-attitude which allows you to approve of the process as a whole. In short, you intend to approve of the process and reason with that intention. This is not an intention prior to the action, but an intention-in-action (Searle 2010: 83). What exactly is this action?

[38] By contrast, Vellemen (2000:134) has argued that agents who act identify with their motive for action, where this identification somehow "makes a causal contribution to their operation"; arguably also Frankfurt (1971).

[39] I borrow the expression from Pettit (2018: 213). Not all attitudes are assumed actively. In particular, I don't think we can plausibly hold that the meta-propositional beliefs, which Pettit claims to be essential to reasoning, are assumed actively (cf. section **1.4**). This is why they do not, like the attitude of approval, impute intentionality to the process.

We can explain what the mental action of reasoning is using Searle's (2010: 94) framework that I have already referred to above. According to his theory, a physical action is the condition of satisfaction of the agent's intention to perform the action. When I intend to raise my arm, the satisfaction condition of that intention is that I raise my arm because I intend to do so. Thus, an action is a bodily movement that is caused by the agent's intention-in-action to perform it. The intention-in-action may but need not be caused by a prior intention. Searle (p.103) claims that mental actions are "formally isomorphic" to physical actions, with mental events taking the place of bodily movements. The mental action of approving (or: the mental action of assuming the attitude of approval) thus consists of the formation of the attitude of approval (mental event) as caused by the intention to approve. Next, the mental action of reasoning is the process leading from the premise-attitudes to the conclusion-attitude, together with the agent's approval thereof, which is caused by the intention to approve. The intention with which you reason is the intention to approve of the process. This does not cause the conclusion-attitude, but it makes it the case that it counts as a conclusion-attitude of *reasoning*. This is so because the intention to approve causes the approval and that makes the process count as an instance of reasoning. In effect, in Searle's framework, the mental action of reasoning *is* the mental action of approving.

You do not form the conclusion-attitude as a means of securing your own approval. This is so for two reasons. First, you do not *perform* the process that causes you to acquire the conclusion-attitude. It happens to you. In theoretical reasoning, for example, the formation of a belief is not something you do, but something that happens to you. There is something involuntary in our responses to epistemic reasons, even if we recognise and realise what follows from what (Parfit 2011a: 48-49). This involuntary part is the very acquisition of the belief itself, which manifests your drawing of a conclusion only in virtue of your immediate

but post-hoc approval of the underlying causal process. When Max ponders over a mathematical problem, and suddenly thinks he got the answer, shouting "that's it", the phenomenology of the event is that the answer 'came to him', while at the same time we would want to say *he* derived the answer. We can intentionally reason, which leads us to new beliefs, but we cannot at will form this or that belief.

In practical reasoning, the formation of the intention itself is also not something you do unless you immediately approve of it. That which brings you into the state of intending is not under your agential control.[40] Think about the phenomenology of more complex cognitive decision-making processes. Consider a case in which you weigh the reasons for and against a certain course of action. You consider this and that and eventually pick an option, not at random, but through reasoning. The experience is that something tips you over. You make the choice, but it was suggested by your mind. The formation of the attitude (intention) itself is not necessarily something you do. You acquire it, and that only counts as something you do in virtue of your post-hoc approving of the underlying causal process. Libet (1999) argues and gives empirical evidence for the claim that the conscious awareness and sense of choice in what we call voluntary acts is preceded by an involuntary, unconscious brain process. Libet calls the electrical signal manifesting this brain process a "readiness potential." His interpretation of the findings is that a free will could veto or block the process which is initiated unconsciously. He does not deny the possibility that we have free will but conceives of free will as having a controlling role, rather than an initiating role. Libet's finding that we can "control ourselves" in this way, I believe, supports my proposal. The mental act of reasoning, marked by the agent's approval of a process, is preceded by the process itself. It would be incorrect to say that you form the conclusion-attitude as a means of securing your

[40] In Korsgaard's (1997: 234) Kantian terms: "Inclination is proposed, but it is the person herself who disposes."

own approval, because in a sense, *you* don't form the conclusion-attitude at all. When we reason about what is the case and what to do, we depend on the suggestions of our minds. The suggestions we take are the judgements and decisions we make.

There is a second reason why it would be misleading to say that you form the conclusion-attitude as a means of securing your own approval. Even though the approval is an aim, it is not that for the sake of which you reason. Your aim is to believe and do what is right to believe and do. That is, your aim is to believe and do what seems to you the right thing to believe and do. You intend to form attitudes in a way that accord with our conception of what is right to believe and intend, and what makes what right to believe and intend. This is just another way of saying that you intend to form attitudes in a way that you approve of. Because of that, you count as intending to approve of the process, without reasoning for the sake of approving. Here is an analogy: When I give a talk, I intend to do it in a way that my audience finds engaging and interesting. I intend to engage my audience. But this does not mean that I give the talk in order to engage my audience. I intend to give a talk in a certain way. Analogously, when I reason, I intend to form attitudes in a certain way, namely in the way that I approve of.

The last point I wish to make regarding the intentionality of reasoning and the role of the implicit conception concerns the notion of guidance. In presenting the *Approval Account*, I have in passing claimed that what guides the reasoner is his conception of what constitutes a permissible belief or action; or in short, your conception of normative reasons. How exactly does this conception *guide* you? Part of what it is to be guided by a standard is to act because of it (Broome 2014b). When you reason in a certain way, your reasoning is partly explained by your conception of normative reasons. It is the basis of your approval. You count as reasoning only if you approve of the process. But you would count as approving of the process

in the right kind of way if your approval was properly based on, and thus in line with your conception of normative reasons. If you do not approve of a transition in this sense, you do not count as reasoning. You would be reasoning, however, if you approved of the forming of a different response to the premises, one that you recognise as being sufficiently supported by normative reasons partly because of the premises.

3.9. Instrumental Reasoning

I have so far only considered reasoning where the premise-attitudes are beliefs. I now extend the account to reasoning on the basis of intentions, in particular, to instrumental reasoning. The paradigm case of instrumental reasoning takes you from an end that you intend, and a belief about the necessary means for achieving that end, to the intention to take that means. For example, you might reason "I shall go to London on Monday. I can go to London on Monday only if I buy a bus ticket. So, I shall buy a bus ticket", where the first sentence expresses an intention, the second sentence a belief, and the last sentence the intention you form on the basis of that belief and your initial intention. There are other kinds of instrumental reasoning. For example, you might reason to the intention to take a means you believe to be best or sufficient to achieve your end, rather than to one that you believe to be *necessary*.

According to the *Approval Account*, a passage of thought counts as instrumental reasoning if and only if you approve of the transition as you experience it. How do you experience an intention of yours, when you reason on its basis?[41] To intend something is to commit to doing it if the opportunity arises. If you intend to ϕ, you perceive this intention as a commitment to

[41] I do not take a side on the debate between Broome (2013: 23), who describes the contents of intentions as propositions (to intend to ϕ is to intend that you yourself will ϕ) and Hornsby (2016), who describes them as infinitives.

ϕ. Sometimes when you attend to an intention to ϕ you are prompted to reconsider it. This reconsideration might be prompted by the awareness of an additional reason not to ϕ or considerations that cast doubt on the (strength of) the reasons you take yourself to have for ϕ-ing. However, when you reason on the basis of an intention to ϕ, you do not reconsider that intention. You treat the commitment to ϕ as given, and the question of whether to ϕ as settled, just as you do not doubt whether p when you reason from the belief that p.

Instrumental reasoning, like any kind of reasoning, commits you to certain normative beliefs. Reasoning from the intention to ϕ and the belief that Ψ-ing is a necessary or sufficient means to ϕ-ing, to the intention to ϕ commits you to believing that the first-person premises, if they were all that is relevant to whether or not to Ψ, would make it the case that you have sufficient reason to Ψ. So instrumental reasoning commits you to believing that all other this equal, your commitment to ϕ, together with the fact about ϕ-ing and Ψ-ing, make it the case that you have sufficient reason to Ψ. Again, it would be implausible if correct reasoning committed you to false beliefs. So, is this belief true whenever the reasoning is correct?

"Obviously not", one might think. Your commitment does not give you any reason to Ψ, ever. And the fact that you can ϕ only if you Ψ is itself not a reason to Ψ, and it cannot be turned into one by your intending to ϕ. You cannot bootstrap reasons into existence at will (Bratman 1987: ch.1). An intention to ϕ neither is nor gives you a reason to Ψ (Broome 2001). So, the belief that instrumental reasoning commits you to according to *Approval Account* is always false! [42]

But I argue that the answer is 'yes'. The normative belief which instrumental reasoning commits you to is true whenever the reasoning is correct. Because of the force and

[42] Instrumental reasoning is seen as a problem for the *Reasons View* of reasoning because of the bootstrapping argument, see e.g. McHugh and Way (2018).

widespread acceptance of the bootstrapping objection, it seems I cannot leave that argument until chapter 5, where I discuss the correctness of reasoning. My account of what reasoning is would be implausible if it entailed that intuitively correct instrumental reasoning commits one to false beliefs.

To begin with, if your instrumental reasoning manifests a conception of normative reasons according to which your commitments give you reasons, then this conception, and consequently your reasoning, is not correct (given we reject bootstrapping).[43]

However, we must notice that even if we cannot bootstrap reasons into existence at will, at least some of our intentions do make a difference to what we have reason to do. They activate conditional instrumental reasons. Suppose you have sufficient reason to ϕ and that you can ϕ only if you Ψ. This fact is an instrumental reason for you to Ψ, but only on two conditions: first, that you intend to ϕ, and second, that you have sufficient reason to ϕ. Again, the intention is not a reason, but a condition on which a fact is an instrumental reason. The condition that this fact is a reason only if you intend to ϕ rules out the so-called 'actualist challenge' against instrumental reasons, which can be stated as follows: Suppose you ought to take the 12:10 plane and can only do that if you go to the airport now. Do you have a reason to go to the airport now? Suppose you have no intention to board the 12:10 plane (because you are terrified of flying), even if you were at the airport in time to do so. It seems without the intention to board the plane, i.e. if you actually would not board the plane, you have no reason to go to the airport.[44] However, if you intend to board the plane, if you commit to do it if you can, there is no ground for this objection. The fact that going to the airport is necessary to board the plane gives you sufficient reason to go to the airport. So, we cannot

[43] I will argue in chapters 4 and 5, that you reason correctly whenever the conception of normative reasons you manifest in reasoning is correct.
[44] Jackson and Pargetter (1986) and Broome (2001) argue this way, see Kiesewetter (2015) for a discussion.

bootstrap reasons for action into existence as we wish, but we can activate instrumental reasons by intending to do what we have sufficient reason to do. Our powers to create reasons are constrained by the reasons we are given. The instrumental reason that I cannot ϕ unless I Ψ is a reason for me to Ψ if I intend to ϕ and have sufficient reason to ϕ.[45] If ϕ-ing is not permissible, then the fact that Ψ-ing allows me to ϕ is not a reason to Ψ. So, if you do some instrumental reasoning, and thereby manifest a conception of normative reasons according to which your intention to ϕ can be a condition for some fact to be an instrumental reason (which is plausibly correct), your reasoning is correct. You then implicitly believe that you have sufficient reason to ϕ. The instrumental reasoning commits you to believing that your commitment to ϕ, together with the fact that you can ϕ only if you Ψ, makes it the case that you have sufficient reason to Ψ, given that you have sufficient reason to ϕ. Your reasoning does not commit you to false beliefs.

Does this imply that your instrumental reasoning is incorrect unless you believe that you have sufficient reason to ϕ, implicitly or explicitly? As long as we reject bootstrapping, my answer is 'yes'. In defence of this controversial position, I shall make the following two remarks. First, you can temporarily and unjustifiably believe that you have sufficient reason to ϕ, correctly reason to taking the means to ϕ, while at the same time holding beliefs which entail that you ought not to ϕ. You would then be irrational and live in denial about the fact that you ought not to ϕ, but your instrumental reasoning would be correct. As I argue in chapter 5, the correctness of reasoning is a local assessment, taking into consideration only

[45] For simplicity, I restrict my discussion to instrumental reasons for taking necessary means. More generally, we could define an instrumental reason as the fact that there is a positive probability, conditional on one's Ψ-ing, that one's Ψ-ing or some part of one's Ψ-ing helps to bring about that one ϕ's (Niko Kolodny 2018). Though I accept this proposal, I doubt Kolodny's claim that if you have *some* pro-tanto reason for ϕ-ing, and Ψ-ing helps me to ϕ, then you have *some* reason for Ψ-ing. Instrumental reasons are conditional on the end being permissible, i.e. on non-instrumental reasons in a similar fashion in which prudential reasons are conditional on moral reasons on Kant's view (Bader 2015). Korsgaard (1997) makes a similar point in her argument against the Humean claim that all reasons are instrumental.

the attitudes constituting the reasoning. If your instrumental reasoning is correct, the beliefs which entail that you ought not to ϕ play no role in it. You may even believe that you ought not to ϕ and still do some correct instrumental reasoning on the basis of the intention to ϕ as long as this belief does not figure in the reasoning. This seems plausible, for if the belief about what you ought to do plays a role in the reasoning, the reasoning would be incorrect. The reasoning expressed by "I shall get drunk! I can get drunk only if I go to the pub. I really ought not to get drunk. So, I shall go to the pub" is not correct instrumental reasoning.

The second remark in defence of my proposal is that many of the instances we call "instrumental reasoning" are not instances of reasoning at all, but descriptions of passive transitions of the mind. They are manifestation of a well-functioning 'system 1', the fast, intuitive, passive kind of thinking that happens in many different kinds of minds. These processes do not involve the agent's endorsement or approval, they are not something the agent consciously *does*. Often when we intend an end, and have beliefs about the means to get it, we automatically acquire the intention to take those means. We are just compelled to do what is needed – whether or not we believe we have sufficient reason for our end. No reasoning is required. Indeed, many animals which we would not necessarily grant the capacity for active, person-level reasoning (of the kind I have been discussing here), undeniably have the ability to engage in instrumental "reasoning" – thinking that leads them to take the means necessary for their ends. These processes that lead them - or us - to take the means that are believed to be necessary for intended ends, can be evaluated against a standard of rationality. They make us more rational, we may argue, because they bring us to satisfy the requirement of rationality that demands us to not at the same time intend an end, believe a means is necessary for achieving that end and not intend the means.[46] In this context, Kant's

[46] For discussions of this state-requirement of instrumental rationality, see for example, Brunero (2010); Rippon (2010); Schroeder (2004); Broome (2013:9.4). Indeed, instrumental rationality and instrumental

claim about instrumental willing seems plausible, though it goes further than my proposal. Kant writes: "Whoever wills the end, wills (in so far as reason has decisive influence on his actions) also the means that are indispensably necessary to his actions and that lie in his power."[47] Kant believes this proposition is analytic, that is, willing an end just is willing the means to it. I don't go that far, but Kant's view, too, is that no reasoning is needed to come to intend or to will the means, if one fully intends the end.

In sum, I believe the claim that instrumental reasoning is correct only if you believe that the end you intend is permissible is less implausible than it might seem at first sight. We must not forget that agents, who reason to the intention to take the means to ends they do not even implicitly believe to be permissible, are plausibly irrational. Second, we must not overestimate the importance of instrumental reasoning and underestimate the ubiquity of rational, automatic instrumental thinking. With that, I can conclude that the *Approval Account* can accommodate instrumental reasoning, and its correctness, in a plausible way.

I have here assumed that instrumental reasoning is reasoning on the basis of intentions. Before I end this chapter, let me briefly defend this assumption against two alternative approaches.

Dancy argues that the first premise of instrumental reasoning, which captures the end the reasoner pursues, should be seen as a belief, not an intention. The content is something like "that action would be useful, pleasant, etc", in which "value is already involved or imputed" (2018:118). It goes well or is correct because the consideration that the end is worth achieving, together with the consideration that a means is necessary (or sufficient or partial)

---

reasons (see Kolodny (2018) for an overview) have received much more attention than instrumental *reasoning*.

[47] See Kant's Groundwork for the metaphysics of morals (2002:417). Korsgaard's (1997: 234 - 236) interpretation of Kant's instrumental principle "If we will an end then we ought to will the means to that end" implies that my above argument is also broadly Kantian. For if we 'will' (as opposed to desire) an end, then we act on an object of volition, under the adoption of a maxim to act. And only then does the instrumental principle apply. This seems to fit my proposal that instrumental reasoning is correct only if we believe we have sufficient reason to do what we intend to do.

for achieving that end together favour taking the means. But this view has the odd implication that my reasoning to taking some means to ϕ on the basis of believing that ϕ-ing would be worthwhile (and a belief about the means) is correct even if I do not intend to ϕ - and even if I intend to not ϕ! Something essential gets lost when we understand the ends which we seek to implement through instrumental reasoning as Anscombian (1957:§37) "desirability characterisations."[48]

Alvarez (2010:5-7) argues that "I want to φ" expresses your intention to φ. The premises of a practical inference that leads to an action jointly show what is good about that action. My wanting to φ, however, does itself normally not contribute to goodness of φ-ing. Therefore, it is not a premise of the reasoning, a goal and that for the sake of which I reason. Alvarez assumes that in good practical reasoning towards an action, each premise contributes to the goodness of the action, which is too restrictive. Not every premise is a reason, some are conditions and modifiers. Seeing that, we need to conclude that the intention is not a premise, but that it is a condition for another premise to be a reason, which is my proposal. Further, it is misleading to say that we express an intention by "I want to ϕ", which seems to express a mere desire. Our intentions (in Kantian terms, our wills) matter more to what we have reason to do than our desires or inclinations. If I *decided* to have pasta rather than pizza for dinner, this seems to have more bearing on whether or not I ought to boil water than my mere *desire* to have pasta. I do not deny that sometimes an intention prompts me to reason in a certain way without functioning as a premise of that instance of reasoning. But we can make decisions on the basis of commitments we have previously made, and when we

[48] I read Anscombe here as describing practical reasoning from premises which are reasons for the practical conclusion-attitude, where the agent sees these considerations as reasons *because* of he wants certain things. The reasoner does not consider his wanting something as one of the reasons; it is not a premise. Anscombe, in a sense, makes all practical reasoning about one's own "wants", or desires. If we follow Dancy's interpretation of instrumental reasoning, it seems we are either similarly pushed to believe that all practical reasoning as instrumental reasoning, or to deny the distinctive nature of instrumental reasoning as reasoning which is – as opposed to other practical reasoning – about our own ends.

deliberate about what to do, we take into consideration what we have already committed to do.

3.10. Conclusion

In this chapter, I have defended the *Approval Account* of reasoning, according to which you reason from attitudes $A_1(a_1)$ … $A_n(a_n)$ to $A_C(a_C)$ if and only if, and if so because, you immediately *approve* of your forming A_C on the basis of $A_1(a_1)$ … $A_n(a_n)$ as experienced by you. You intentionally approve of the process, which implies that reasoning is an intentional action. The attitude of approval that is essential to reasoning is normative in nature: it is an evaluation of your transition according to your implicit conception of the relation of 'permissibility-making'. Approving of a thought process cannot be interpreted simply as a belief about it, or the belief that something is a reason for a certain response. Though reasoning rationally commits you to beliefs about reasons, beliefs about reasons are neither necessary, nor sufficient for reasoning. You need an implicit conception of what makes it the case that some facts allow a certain response. You need a conception if what constitutes having sufficient reason to do or believe something. The attitude of approval is not a premise of the reasoning, but rather manifests a person-level, evaluative endorsement of it.

Reasoning is thus an inherently normative activity. The capacity to reason is the capacity to control our mental life according to our conception of what is right to do and believe. It is what Korsgaard (2010) calls the capacity for normative self-governance: "the capacity to assess the potential grounds of our beliefs and actions, to ask whether they constitute good reasons, and to regulate our beliefs and actions accordingly."

However, I have not yet established the *Reasons Claim*, which states that reasoning is a way of responding to normative reasons. This is the task of the next chapter. There I give

an account of what it is to respond to normative reasons, which will allow me to argue that reasoning – understood as I propose here – is a way if responding to normative reasons.

4. Responding to Normative Reasons

4.1. Introduction

If there are normative reasons for us to do and believe things, then we must have the capacity to respond to them. Normative reasons, I continue to assume, are facts that count in favour of certain responses.[1] The fact that he is starving is a reason to feed him. The fact that it would be good for your health is a reason to go for a long walk in the woods. As humans, we do not merely react to the world as we perceive it, we also respond to the reasons we have to do, believe, hope, fear, or admire things, for example. We can respond to normative reasons.

But what exactly is it to respond to normative reasons? The notion of "responding to normative reasons" is a philosophical term of art.[2] It is widely used in philosophy but, I believe, mostly refers to the following phenomenon. Suppose Tom sees an elderly lady who stumbles and falls on the other side of the street. He crosses the street to help her. Whether Tom reflects on the situation or acts intuitively, he responds to the normative reason he has, seeing not only how things descriptively are (street, women, falling, ...) but also what this situation demands from him – that there is a reason for him to do something. He is not just behaving or reacting in a certain way but responding to the facts *as* reasons. Tom acts for the normative reason he has to cross the street.[3] But it is possible – indeed probable - that when Tom crosses the street, he only considers the non-normative features of the situation and acts on their basis. He may not have to have explicit beliefs about what is a reason for him to do

[1] See Introduction of this thesis. To the notes there, I must add that when I say 'normative reason' I mean 'objective reason' in the sense in which Schroeder (2007:14; 2008), amongst others, defines an objective reason in contrast to a subjective reason. A subjective reason for agent N to ϕ is a consideration which N believes and which, if it was true, would be an objective reason for N to ϕ. In this sense. These subjective reasons are Parfit's (2011a: 35) apparent reasons, and not to be confused with Parfit's (2011a: 45) subjective reasons (which relate to the agent's desires).

[2] I am borrowing this expression from Kolodny (2005).

[3] I assume that 'responding to a normative reason for action', 'responding to a fact *as* a (or qua) normative reason to act' and 'acting for a normative reason' all mean the same. Lord (2018), Mantel (2017) and Cunningham (2019) use the latter notion.

what. So why does Tom's action count as a response to *normative reasons*, rather than as a reaction to the non-normative *facts* of the matter (given that is what he thinks about)? How does normativity enter his mind so that he counts as responding to the normative features of the world? We need an account of what it is to respond to normative reasons that answers these questions. My aim in this chapter is to propose one.

It is rather surprising that few attempts have been made to provide an account of responding to reasons, for the notion appears in a wide range of philosophical debates.[4] To begin with, many philosophers assume that reasoning is a way of responding to normative reasons (the *Reasons Claim*). In chapter 2, I have assessed a number of accounts that define the activity of reasoning in terms of responding to reasons and argued that they are all unsatisfactory. Their shared problem is that they do not sufficiently explain what it is to respond to normative reasons. These accounts appeal to notions such as '*treating* things as reasons', '*taking* considerations to be reasons', or '*seeing* facts as favourers, but then do not explain what these notions exactly refer to. We need an account of what responding to normative reasons is if we wish to argue that reasoning consists in responding to normative reasons. To be clear, I don't deny that responding to reasons is '*treating* things as reasons', '*taking* considerations to be reasons', or '*seeing* facts *as* favourers.' But saying that alone does not explain anything, and only describes what needs to be explained in different terms. We need to dig deeper.

Indeed, my main motivation for developing an account of normative reasons is to advance our understanding of the activity of reasoning. However, the account is relevant for other debates, too. Here are some examples. First, recently many have claimed that the ability to respond to normative reasons is distinctive of human nature. Parfit (2011b: 620), for example, writes: "We would not be the animals that can understand and respond to reasons. In a world

[4] Appendix II argues that reference to a 'sensitivity to normative reason', as well as dispositional accounts of responding to normative reasons seem unsatisfactory.

without reasons, we would act only on our instincts and desires, living as other animals live." A precise understanding of what the latter amounts to should thus further our understanding of what kind of creatures we are.[5] Similarly, some have explained personhood, the will, or rational agency by reference to the ability to respond to normative reasons.[6] Furthermore, it is widely accepted that actions are morally worthy whenever they are "performed for the reasons *why* they are right" (Markovits 2010: 202) that is, when they are responses to the normative reasons that make them right. Pettit (2018: 283) argues that I create the goods resulting from robust moral performance only if I "act out of responsiveness to the relevant reason – to act, as we say, out of the disposition of friendship or respect." Understanding what it is to act for a normative reason (or to respond to a normative reason) thus allows us to understand what is required for an action to be morally worthy. An analogous case can be made in the theoretical domain. Suppose you hold beliefs the contents of which entail the truth of p. If you do not believe p for those reasons you have to believe p, you are not fully justified in believing p. As epistemologists say: you are propositionally, but not doxastically justified in believing that p (e.g. Evans 2013; Silva 2015).[7] So the account I propose may help to better understand what doxastic justification exactly requires.

Above I briefly described the phenomenon of responding to normative reasons. Summarizing the questions raised there, the task of an account of responding to normative reasons is to explain the

[5] See also Scanlon (1998: 23), Hacker (2007: ch.7), Pettit (2007b) Raz (2011: 98), Dancy (2018:175).
[6] For example, Raz (2002: 21; 2011: 86) argues that responsiveness to reasons is constitutive of being a person. Wallace (2015) argues that "the capacities constitutive of rational agency involve a capacity not merely for responding causally to the reason-constituting facts […] but for responding to one's acknowledgment of their normative significance for one's own situation." Wallace (2006) also writes that agents "who are equipped with what we might as well call a will—that is, with the capacity for self-determination independently of the desires and dispositions to which they are merely subject—can retain the capacity to comply with reasons that they knowingly act against."
[7] See Turri (2010) for an objection to this orthodox understanding of doxastic justification

Normative Leap: At least some of our actions and beliefs are not mere reactions to how things are, but responses to facts as normative reasons to do or believe things. These responses do not need to be accompanied by explicit normative beliefs involving the normative concepts of 'reason' or 'favouring', for example. Yet, the (supposed) normative properties of that which we respond to are essential to the nature of our responses.[8]

My aim is to propose an account of what it is to respond to reasons that can explain the *Normative Leap*. Put roughly, I argue that responding to reasons is a mental process that leads an agent from beliefs about the world to a response, where this process manifests the agent's implicit conception of normative reasons and is under the agent's (virtual) control.

In section 4.2, I argue that responding to a consideration f as a reason to believe or do something is a mental process through which you come to believe or intend something in light of f. **Section 4.3** argues that – even if we have no false beliefs – we sometimes treat things as reasons which are not reasons. Some of the reasons we respond to are merely *supposed* normative reasons. What is essential to count as responding to practical or theoretical reasons is that you manifest an implicit conception of rightness (or goodness) or epistemic value, respectively. **Sections 4.4 and 4.5** explain what these implicit conceptions consist in. However, even if you manifest such a conception, you may still not count as responding to normative reasons. **Section 4.6** argues that for a process to count as a response

[8] One might object that the *Normative Leap* is a phenomenon of certain theories about motivation and /or normative reasons. I address this worry in Appendix I of this chapter. I argue that the question of how we get from the non-normative to the normative in our responses to the world arises (i) for both 'internalists' and 'externalists' about normative reasons, (i) also on a Humean Theory of Motivation, even if it is supplemented with the doctrine of the Guise of the Good about desires, and (iii) even if all responses to practical reasons are responses to instrumental reasons.

to normative reasons, it must be at least under your agential control. In **section 4.7** then, I finally turn to the second aim of this chapter and investigate the relation between reasoning and responding to reasons. I will show that reasoning – understood as I suggested in the previous chapter - is a way of responding to normative reasons. The capacity to reason is the capacity to respond to normative reasons. **Section 4.8** briefly considers the question of what it is to correctly respond to normative reasons. Incorrectly responding to reasons is responding to merely supposed reasons. **Section 4.9** shows how the proposed account explains the *Normative Leap* and considers some implications. Our (implicit) conceptions of what constitutes a normative reason for belief or action explain why we treat considerations as reasons in the way we do. **Section 4.10** briefly summarises and concludes the discussion.

Before I begin, three preliminary remarks are necessary. The first sets the scope of my analysis. We sometimes think about what would be right to do and respond accordingly. A paradigmatic case of responding to normative reasons is doing what one believes one ought to do. If I decide to apply for the job because I believe I ought to, I am not only responding to normative reasons, but I am arguably even responding correctly (Broome 2007). I do not deny that in some sense, deciding to do what one believes one ought to do is responding to normative reasons. But this is a difference sense than I am interested in here. When I decide to ϕ because I believe I ought to ϕ, then I do not *treat* or *see* any consideration as a reason. I do not respond to a non-normative fact *as a* normative reason. There is no unexplained *Normative Leap*, for it is obvious how normative thoughts are involved: I respond to a normative belief. That process does not involve treating or seeing any considerations as reasons. Acting on an explicit normative belief – like the belief that you have some or sufficient reason to ϕ - is different from acting or believing for specific considerations one *sees* as normative reasons. The latter, we may say, is responding to normative reasons *de re*,

the former responding to normative reasons *de dicto*. I am here only concerned with responding to normative reasons *de re*.

The second is a clarification. Sometimes we respond to *a* normative reason, sometimes to multiple normative reason*s*. The latter is just a more complex version of and involves the former. Whether you respond to one or several considerations *as normative reasons*, the question that the *Normative Leap* raises is the same. When you see that it starts raining and based on that consideration alone run outside to take the laundry in, then you respond to *a* normative reason. You do not necessarily respond to only one consideration. The considerations you respond to include that which we ordinarily call the reason, as well as its respective conditions and modifiers (Dancy 2004). For simplicity, I leave conditions and modifiers out of the discussion. You respond to reason*s* when you take more aspects of the situation into consideration. In simple cases, all considerations you respond to favour the same kind of response; they are reasons for ϕ-ing (or believing p), together with their respective conditions and modifiers. In complex cases, you respond to a set of considerations some of which favour some kind of response, others another. You "weigh" opposing reasons against each other. I here focus on simple cases of responding to normative reason*s*, turning to cases that involve the "weighing" of reasons in chapter 5. Sometimes I drop the "normative" and speak of responding to reasons.

Finally, let me limit the scope of the discussion once again. I consider only how we respond to reasons for action (reasons to ϕ, where 'ϕ' stands for a kind of action), and how we respond to epistemic or truth-related reasons for belief. An epistemic or truth-related reason for believing p indicates the truth of p. If there are pragmatic reasons for belief (reasons for belief that are related to the goodness of having a belief), then how we respond to these reasons is

not my concern here.[9] I also do not consider reasons for emotions or other responses. In any case, we may not assume that we respond to all kinds of reasons in the same way. The assumption that we can give a unified account of responding to reason for action and epistemic reasons for belief is substantial enough.

I continue here to assume that reasons for action and belief are facts. But I do not make any claims about which facts are or can be reasons for an agent to do or believe something. Maybe facts are normative reasons only on some condition, say, the agent stands in a certain epistemic relation to the respective fact,[10] or that the agent has certain desires related to the fact.[11] As we shall see in section **4.8**, my account of responding to reasons does not rely on the realist assumption that there are normative reasons, but only on the metaethical claim if there are normative reasons, these reasons are facts. Maybe I explain an illusion – the illusion that we respond to reasons we have to do or believe things.

4.2. Believing and intending in light of considerations

I now begin to develop my account of responding to normative reasons. A natural starting point is the thought that when you respond to a consideration as a normative reason, then you respond in light of that consideration. For example, when Tom crosses the street and thereby

[9] Some have denied the existence of pragmatic reasons for belief on the grounds that we cannot respond to them in the "right way": we cannot believe that p on the basis of practical reasons for believing p (Shah 2006; Kelly 2002; Way 2016). But we *can* respond to pragmatic reasons for belief, just not in the way in which we respond to truth-related reasons, and this itself does not imply that pragmatic reasons for belief are not reasons, as for example Piller (2016) argues.

[10] For example, Gibbons (2010; 2013:135) perspectivism, the view that what you have reason to do depends on your perspective. Worsnip (2016:347) argues that reasons for belief and action are facts that are epistemically available to us. In a similar manner, Lord (2010; 2018: ch. 3,4) gives an account of "possessed reasons", and Kiesewetter (2017) of "evidence-relative" reasons; Hawthorne and Magidor (2018) and Littlejohn (2018) require knowledge. I do not need to consider this debate, because I am concerned with responding to consideration as normative reasons (whether they are normative reasons or not).

[11] My arguments are compatible with subjectivist theories of normative reasons "according to which reasons for acting are all provided by, or depend upon, certain facts about what would fulfil or achieve our present desires or aims", or desires we would have had we considered all the relevant facts (Parfit 2011a: 45), but I do not consider such views.

responds to a normative reason, he crosses the street in light of the reason he has. In this section, I defend this claim and argue that we should think of responding to reasons as a mental process. I consider an objection that Errol Lord has recently raised against the kind of approach I take.

4.2.1. Believing and intending …

One response to an epistemic reason for believing that p is to form the belief that p. For example, we may respond to the fact that the newspaper says that global warming leads to great suffering with the belief that this is so. We can respond to a reason to believe that p with another attitudes, the belief that not-p, suspension of judgment, or a credence (a degree of belief), for example. I leave these cases aside for now, and consider cases in which we respond to a reason for believing that p with the belief that p.

Moving on to reasons for action, one way of responding to a reason to ϕ is by forming the intention to ϕ, where ϕ stands for a certain kind of action. When I consider the fact that it is my mother's birthday tomorrow and thus form the intention to call her tomorrow, I respond to a reason to do something by forming the intention to do it. This is what Searle (1983) calls a prior intention; it is a plan for future action. But of course, not all responses to practical reasons are such prior intentions. Suppose you see the child drowning in the pond and upon realisation of this fact decide to jump in. Before you reach the pond, you fall. You fail to do what you take your reasons to require of you; you do not to help the child. But you do not, it seems, fail to respond to the reasons you take yourself to have. If you did not fall and reached the pond, you wouldn't have exercised your capacity to respond to normative reasons any differently. The point is that you may respond to a reason to ϕ by forming the intention to ϕ while failing – for no fault of your own – to execute that intention (Broome 2007). In the

example here, you manifest the intention to save the child by intentionally trying to do so. If all went well and you would have succeeded in saving the child, your response to reasons would be an intentional action. When you act in response to reasons, you are aware of what you are doing and in control of what you do; you act intentionally. The intention with which you act is an intention-in-action. But you also count as responding to the reason you have if you are prevented from executing this intention successfully. Thus, for whether or not you count as responding to reasons, the intention-in-action is the essential part of the action, not the bodily movement.[12] Therefore, I consider the responses to practical reasons to be intentions, where these are prior intentions or intentions-in-action. (Again, one can also respond to a reason to ϕ with another attitude, but I leave these cases aside until later.) Our responses to truth-related reasons for belief and to reasons for action can thus be understood as mental attitudes, in particular beliefs and intentions.

4.2.2. … in light of …

Suppose it is your mother's birthday, and that this fact is a reason for you to call her. Suppose you do decide - that is, form the intention - to call her. A natural thought is that this intention is a response to the normative reason you have for calling your mother only if you form it *in light of* that reason, namely that it is her birthday. This requires, first of all, that you believe that it is her birthday. However, if you believe that it is your mother's birthday and decide to call out of an impulse to continue your disagreement from last night, your intention to call is not a response to the reason you have to call. It is just a coincidence. You intend to call your mother *in light of* the fact that it is her birthday only if your belief that it is her birthday stands in a particular relation to your intention. It must cause the intention in the right kind of way,

[12] This parallels my brief defence of the claim that practical reasoning leads to intentions (section 3.1).

so that the intention is based on the belief. Generalising from this example, we can say that one must intend to ϕ in light of the reason one has to ϕ to count as responding to that reason. More formally: agent N's intention to ϕ is a response to fact f as a normative reason only if N intends to ϕ *in light of* f.

An analogous argument can be given for responding to reasons for belief. Suppose the fact that it is your mother's birthday is a reason to believe she is waiting for your call. Again, you may respond to this reason by forming the belief that she is waiting for your call. The belief counts as a response to the reason only if you form it in light of it. Generalising again, we can say that one's belief that p counts as a response to an epistemic reason f only if one believes that p *in light of* that reason (f). This is to say that one's belief that p must be caused by the belief that f in the right kind of way. The belief that p must be based on the belief that f.

The starting point of this analysis could have also been stated as the claim that responding to normative reasons requires acting, intending or believing *for those reasons*. If Tom responds to the normative reason he has to help the lady Tom acts for that reason. Analogously, if you call your mother as a response to the reason you have for calling her, you must do it for that reason – namely that it is her birthday. But what is it to do something for a reason? The word 'reason' in this question means 'motivating reason', not 'normative reason'. Regarding the latter example, some would argue that your motivating reason for calling your mother is the fact that it is her birthday. Others would argue your motivating reason is *your belief* that it is her birthday.[13] Luckily, we do not need to enter this debate here, for we can avoid the notion of a motivating reason entirely. You may call the fact that f or N's belief that f N's motivating

[13] Dancy (2002:ch.5) argues motivating reasons are real or supposed states of affairs, denying psychologism, the view that an agent's motivating reasons are mental states, which arguably goes back to Davidson (1963), see also Smith (1994:ch.4) and Bittner (2001: 65). Non-Psychologism has been defended by, for example, Raz (2002: ch.2), Darwall (2003), Schroeder (2007: 14), Alvarez (2009), Hyman (1999; 2015); McDowell (2013), Broome (2004) and Hornsby (2008).

reason.[14] The point is that responding to a normative reason f requires that one acts or believes *in light of* that reason f.

4.2.3. … considerations

Now, we have seen that to respond to a fact f as a normative reason, one must believe that p. So, strictly speaking, when we respond to reasons, we respond to things we believe. In the cases discussed above, what was believed was true. But this need not be so, as the following example shows.

Suppose you falsely believe that there is egg in the cupcake I baked for you and thus decide not to eat it, knowing you are allergic to eggs. Is your decision a response to a normative reason? One might think it cannot be a response to a reason, because there is no egg in the cupcake, so there is no reason for you not to eat it. You cannot respond to a reason that is not there. But we must note that the mental process which leads you from your beliefs about the world to your response is the same whether or not your beliefs are true. If your belief is false, we can say with Parfit (2011a:35) that you respond to an *apparent reason:* what you respond to would be a reason if it was true. When you respond to apparent reasons, you manifest the same kind of capacity in the same kind of way as you do when you respond to actual normative reasons. You are mistaken about the non-normative features of the world but that

[14] The theoretical analogue to motivating reasons are what Scanlon (1998:18) calls operative reasons. Again, for N's belief that p to count as a response to normative reason f, N must believe that p in light of f. You may call the fact f, or the belief that f the agent's operative reason to believe that p. It is largely assumed in contemporary epistemology that if my operative reason for believing p is f or the belief that f, then my belief that p is *based on* the belief that f (Sylvan 2016b). It is currently proposed to analyse this basing relation in terms of treating things as normative reasons for belief (Sylvan 2015; Sylvan and Lord forthcoming). In relation to this trend in epistemology, my account *RNR*– where it appeals to treating things as reasons for belief - could be seen as an account of the basing relation. But I do not understand it as such. I use the term 'basing' simply to describe a causal relation between attitudes that exclude deviant causal chains, as also e.g. Anscombe (1974) does. Most epistemologists understand *treating* as a disposition or a competence. I mention dispositional accounts of responding to reasons in Appendix I, but I will have to find space for detailed replies to the most recent proposals elsewhere.

does not impair your response to what you take to be true. In the example, there is no egg in the cupcake, so you respond to an apparent reason not to eat it. The nature of the responding is the same whether you respond to apparent or actual reasons. Responding to normative reasons - actual or apparent - does therefore not require true beliefs.

What we respond to are things we believe. I call something that is (explicitly and consciously) believed a consideration. A consideration is thus a proposition that is taken to be true. It may be true or false. A consideration that is true corresponds to a fact. Though we speak as if we respond to facts, we respond to considerations that may or may not correspond to facts. If the considerations we respond to are false, we respond to merely apparent facts. This is what happens when we respond to apparent reasons: we respond to considerations that would be reasons if they were true. In general: If agent N intends to ϕ or believes that p in light of consideration f, N may count as responding to normative reasons even if N's belief that f is false. If f is false, N responds to an *apparent* normative reason.[15]

However, in the example I discuss in this chapter, I assume that the consideration responded to are true. If you respond to consideration f, f describes a fact. I do so only to make the exposition of the account easier, it is no substantive assumption.

To conclude the discussion, we can think of responding to normative reasons for action or belief as a mental process, leading from beliefs about the world to intentions or beliefs. The capacity to respond to normative reasons is thus a mental capacity, independent of our abilities to perceive the world as it is, to know things, or to move our bodies.



[15] Hornsby and Hyman (1999; 2011) argue that acting (or believing) for a reason – better: acting for a fact - requires knowledge. Again, I am not concerned with this notion.

Understanding the nature of responding to normative reasons is understanding the nature of the process represented by the arrow. It may be one of active deliberation or reasoning, but it may also be more habitual and automatic. No active reasoning is needed to respond to normative reasons (Raz 2015), while a response to reasons is not merely mechanical either (Blackburn 1998: 5). In what follows, I propose the following account of

> *Responding to Normative Reasons (RNR):*
>
> Agent N's belief that p / intention to ϕ is a response to consideration f as normative reason **r**, if and only if
>
> (1) N comes to believe that p / to intend to ϕ in light of f (and thus on the basis of the belief that f),
> (2) thereby manifesting an implicit conception of normative reasons according to which f is normative reason **r,**
> (3) through a process that is under N's (virtual) agential control.

I have explained why responses to reasons are responses to considerations, and section **4.2.2** has motivated condition (1).[16] It seems fairly uncontroversial. Some believe that it is not only the beginning of but almost the whole story about responding to normative reasons. They

[16] Condition (1), as I explained it, could also be stated as 'the belief that p / intention to ϕ is *caused* by the belief that f *in the right kind of way*. Condition (2) and (3), one might argue now, just specify what this right kind of way is. If that is correct, then you act in light of f if and only if you respond to f as a normative reason. One might then argue that acting in light of f is acting for the reason that f. One might also argue that *RNR* as a whole gives an account of what it is to act for a reason. Again, acting for a reason would then be acting for a normative reason. Motivating reasons (or the contents of motivating reasons, if we accept psychologism about motivating reasons) are things we see as, treat or take to be normative reason. Alvarez (2016) takes this for granted, and also Smith (1994:131) appeals to this idea. Dancy gives an account of acting for the reason that p in terms of acting in light of p (or p's being a reason). Velleman (1992:4) argues that for something to be "one's reason for performing an action, one must be influenced by its bearing favourably on that action", that is, one is prompted by its "action-justifying character." All this is compatible with what I say.

hold that responding to normative reasons just is acting for reasons that are normative reasons for your action.[17] Lord (2018: 131) calls this view

> **Composite**: What it is for A to F for a normative reason r is for r to be A's motivating reason for F-ing and for r to be a normative reason to F.[18]

Given that **Composite** is supposed to apply to practical as well as theoretical reasons, 'F' stands either for an intention or a belief. Lord's view of motivating reasons seems compatible with the claim that whenever r is A's motivating reason for F-ing, then A F's in light of r. To avoid the language of motivating reasons, we could thus restate **Composite** as follows without changing its intended meaning:

> Composite*: A's ϕ-ing is a response to a normative reason r if and only if A ϕ's in light of r and r is a normative reason for A to ϕ.

Lord argues that **Composite** is false. He concludes that we cannot understand responding to normative reasons in terms of acting or believing for a reason, that is, in terms of acting or believing in light of a reason. With that conclusion, he denies the approach I have taken above, and condition (1) of *RNR*. I agree with Lord that **Composite** (and Composite*) is false. This, however, does not warrant his conclusion. We have to amend **Composite,** but this gives us no reason to doubt condition (1) of *RNR*.

---

[17] In the theoretical domain, some assume that you respond to reasons for the belief that p if and only if the basis of the belief that p are beliefs in reasons for believing p (e.g. Gibbons 2013; Evans 2013). Sylvan and Sosa (2014:15) propose an account of what it is to believe for a normative reason that builds on **Composite**. They require that the agent manifests "inferential competence", which seems roughly in agreement with my proposal. In the practical domain, Markovits (2010), Mantel (2017) Raz (2011:ch.2), Hornsby (2008) and arguably Dancy (2002:136-7) commit to **Composite**.

[18] Interestingly, Lord (2018:167) does not justify his move from the claim that 'when one acts for a reason, one act in light of what makes the action good, desirable or valuable' to **Composite**.

Lord gives the following case as an alleged counter-example to **Composite**, which he says is a version of one given by Swain (1981). I call it *Bad Case*.

Bad Case: "Sam is wondering whether Real Madrid lost yesterday. She believes that Barcelona won yesterday. This is, as it happens, a strong normative reason to believe that Real Madrid lost, as yesterday was the El Clasico game between Barcelona and Real Madrid. Sam infers that Real Madrid lost from her belief that Barcelona won [and the belief that Barcelona played against Real Madrid]. The inference rule she follows is Explosion: Infer <REAL MADRID LOST> from any proposition p."

Sam believes that Real Madrid lost in light of the consideration that Barcelona won. This is also a *reason for* Sam to believe that Real Madrid lost. Nevertheless, Lord claims, Sam's belief that Real Madid lost is not a response to a normative reason, because she fails to be sensitive to that fact *qua* reason. I agree with the first two, but object to the last judgment. Sam may indeed respond to the fact that Barcelona won *as* a normative reason, even though not *as the reason it is*. Put differently, Sam treats the fact that Barcelona won as a reason to believe Real Madrid lost, but she does not treat it *as the reason it is*. She responds to what I will shortly label a *supposed* normative reason.[19] In fact, Lord admits that in the *Bad Case*, Sam treats something as a reason which is not a reason (which contradicts everything else he says about the case). But what she treats as a reason is true, so she does not treat it as an

[19] The possibility that we can treat things as normative reasons that are not reasons in the same way in which we treat things as reasons that are reasons, has not been given enough attention. Objections similar to the one I raise against Lord can be raised against Mantel's (2013) competency-based account of acting for a normative reasons (where she assumes that acting for a normative reason is acting for a fact *that is a reason*), her later dispositional account of responding to normative reasons (Mantel 2017: 567), and Cunningham's (2019) analysis of what it is to believe for a normative reason. Dancy (2002: 3-4) only briefly mentions the possibility responding to supposed normative reasons.

apparent reason. The case shows, I shall argue, that we can treat a consideration as a normative reason, whether it is true or not, and whether it is (or would be, if it was true) a normative reason or not. To make sense of this and my objection to Lord, we first need a better understanding of normative reasons. We will then see the need for condition (2) of *RNR*, and the need for the notion of responding to a *supposed* normative reason. Only with all that at hand, will I be able to defend condition (1) of *RNR* against Lord's objection. Please bear with me.

4.3. Supposed Normative Reasons

Throughout this thesis, I have said, as we often do, that normative reasons are facts.[20] The fact that the child is drowning is a reason for you to jump in the pond to help. This way of speaking is not false but misleading. There is nothing normative about the fact itself; it is stated in fully non-normative terms, as Wedgwood (2018) reminds us. It might have a normative property or stand in a normative relation to other facts, agents or actions. Suppose that as you watch the child struggling in the water, everyone else around is leaving the site. The fact that you remain the only one there strengthens the reason you have to help. It changes the reason without changing the fact that the child is drowning. Thus, the reason is not identical to the fact. When we say that a fact *is* a reason, this 'is' expresses a predication, not identification. It would be more precise to say that the fact has the property of being a reason for you to help the child (Broome 2015b), or that the fact grounds the reason (Bader 2016), or that the fact stands in a certain relation to the agent's acting in a certain way. What I call the ground of a reason may be referred to as a reason-constituting fact (Wallace 2015). For the discussion here, we need to distinguish between the reason (or the property of being a reason) and the reason-constituting fact (the ground of the reason). Only then can we make

[20] For the account here, we can also think of facts as obtaining states of affairs or property instantiations.

sense of the idea that the same fact f can (partially) ground different reasons. For example, the fact that it is sunny may be a reason to go outside (because you would enjoy it), and also a reason to stay inside (because you will get a sunburn). As mentioned in the outset, we will see that *RNR* does not rely on the metaphysical claim that there are normative reasons, only on the metaethical claim that if there are normative reasons, this is how we must understand them.[21]

4.3.1. Same Fact, different Reasons

Because a fact can ground different reasons, a fact – or the corresponding consideration - can be treated as different reasons. I may treat f as a reason to ϕ, you may treat it as a reason to not-ϕ. One of us may make a mistake of some sort, but not necessarily. More complex yet, it may be that two agents treat the same consideration f as a reason to ϕ (or believe p), and yet treat it in different ways. They may treat it as different normative reasons. To give an example, let us contrast *Bad Case* with the following

Good Case: Smart comes to believe that Real Madrid lost by following the inference rule "From <BARCELONA WON> and

[21] The failure to recognize the distinction is one reason why also Mantel's (2017) dispositional account of responding to normative reasons is unsatisfactory. She recognises that "the action has to be connected not only to the belief in the right way, but to the normative reason itself" (575), but she thinks this calls for an epistemic requirement, namely: "The belief that p must be the manifestation of the agent's disposition to represent normative reasons correctly." (577). She explains that "the agent is not responsive to the reason if it is a mere coincident that her belief is true, e.g., if she merely guesses."(576). But because she does not distinguish between the reason and the reason-constituting fact p, what she requires is a disposition to represent the fact correctly – not a disposition to represent *reasons* correctly. She does not account for the problem the recognises. (Of course, I disagree with Mantel that the belief that p has to be true at all, let alone non-accidentally true, but that is another issue.) Cunningham (2019) similarly fails to distinguish between the reasons and reason-constituting fact, arguing that believing for a reason requires knowledge of the reason-constituting fact, and the ability to reason from that fact.

<BARCELONA PLAYED AGAINST REAL MADRID> infer

<REAL MADRID LOST>".

Both, Smart and Sam, believe that Real Madrid lost, they believe it o*n the basis of the same consideration*, namely that Barcelona won. They don't believe Real Madrid lost *for the same normative reason*. Sam and Smart both treat the consideration that Barcelona won as *a* normative reason to believe Real Madrid lost, but they do not treat it in the same way.[22] We need an account of what it is to *treat* a consideration as a normative reason that allows for the possibility that two agents may treat the same consideration as a reason to believe p (or as a reason to ϕ), while treating it as different normative reasons. We could also make this statement in terms of *responding* to a consideration as a normative reason, *seeing* a consideration as a normative reason or *taking* it to be a normative reason.

My proposal is the following: Responding to a consideration f as normative reason **r** for N to ϕ (or believe p) requires manifesting an implicit conception of normative reasons according to which f is normative reason **r** for N to ϕ (or believe p). The difference between Sam and Smart is that they manifest different conceptions of normative reasons. Because of this, they count as treating f as different reasons. They therefore both count as responding to reasons, but to different ones. Many suggest that responding to f as a normative reason requires (or just is) grasping f as favouring a certain response or having a conception of it as a normative reason.[23] The difference between my opponents and me is that I do not stop there. We need

[22] Note that we assume here that reasoning is a rule-following activity, and that in reasoning, the agent treats the premises as reasons. I have casted doubt on both. Nevertheless, we can make these assumptions here for the sake of argument and learn that two agents can respond to the same fact as different reasons.
[23] Wallace (2004:61-2), for example, argues that it is essential to rational agency that agents are guided by a "conception of their reasons for acting as they do" and that we act with "an implicit conception of the reasons that support the action we are performing." But he does not explain what this amounts to.

to explain what it is to grasp a consideration as a normative reason, i.e. to manifest an implicit conception of normative reasons.

In section **3.4**, I explained how we can understand an implicit conception of an object, namely a set of attitudes that an agent holds towards an object, so that the contents of these attitudes constitute a particular characterisation of the object. I here broadly followed Peacocke's (1998) account of implicit conceptions, according to which implicit conceptions influence judgments at the sub-personal level. The content need not be explicitly represented, the agent may not have explicit knowledge of the content and making it explicit may be difficult. Articulating the contents of our implicit conceptions gives us definitions of some sort. Mastering the observational concept 'cube', for example, requires having an implicit conception with a content that includes a description of the features that make an object a cube: "that cubes are closed figures formed from square sides joined at right angles along their edges" (Peacocke 1998: 52). Our judgments about cubes are based on these implicit conceptions of a cube. In a similar fashion, I suggest, our judgements about normative reasons require at least an implicit conception of what constitutes a normative reason, i.e. of that which explains why f is a reason for N to ϕ, and how it explains it. If fact f is a normative reason for N to ϕ (or believe that p), then we can explain why that is so. An (implicit) conception of normative reasons consists of a set of beliefs about that which makes a fact a normative reason. (This is analogous an implicit conception of what a cube is, the content of which describes that which makes something a cube.) There is one important difference between the two cases. When I judge that an object as a cube, my response is solely about the cube. The implicit conception at play is solely about the nature of cubes, it is about a monadic property. In contrast, when I respond to a consideration as a normative reason for me to ϕ, I respond to it as standing in a certain relation to my ϕ-ing. I treat it as something about my ϕ-ing. This is why I respond with an intention to ϕ, rather than another attitude towards f.

Therefore, not only explicit judgments about reasons, but also what we treat as reasons requires and manifests these implicit conceptions. They may be explicit; you may have explicit beliefs about what makes a fact a normative reason for one to do or believe something. But because this is rarely so, and an implicit conception is in a sense less demanding, I shall say that treating f as a normative reason for N to ϕ requires manifesting an implicit conception of normative reasons according to which f is a reason for N to ϕ.

With this, we can now say the following regarding the cases of Sam and Smart:

- Smart and Sam believe that Real Madrid (p) lost on the basis of the fact that Barcelona won (f).
- Smart treats f as reason **r** for her to believe p. She manifests an implicit conception of reasons according to which f is reason **r** for her to believe p because f entails, together with other facts, that p. In fact, f has this property **r**. Therefore, Smart responds to an actual normative reason.
- Sam, however, treats f to be reason **r'** for her to believe p. She manifests a conception of reasons according to which f is a reason for her to believe p because everything is a reason to believe p. She has a different conception of why f has the property of being a reason, and thus treat it as reason **r'**. She treats f as a reason it is not, and thus responds to a merely *supposed* normative reason.
- Sam and Smart manifest different conceptions of normative reasons, and therefore treat the same fact f as different reasons.

One could object by saying that both, Sam and Smart, treat f as a reason to believe p. They just have different opinions of why it is a reason. I suggest that we should avoid such talk. To see why, consider the following case. Suppose two persons are asked to pick a cube from a collection of objects. Both pick what we would call a cube, one a red one, the other a white

one. But the one who picks the red cube believes cubes are red objects. In this case, we would not say that both treat their chosen object as a cube, and that they just have different opinions of why it is a cube. They don't mean the same by 'cube'; one picks a cube, the other a "cube". They treat their chosen object as different things, even though they call them by the same name. In the same way, Sam and Smart treat f as different things. Sam treats f as having the property of being reason **r'**, Smart as having the property of being reason **r**. There is thus no such thing as responding to a fact as a normative reason, but only responding to a fact as the specific normative reason **r**. Smart and Sam both treat some consideration as a normative reason. They respond to the same facts, but to different reasons. This is so because they have different conceptions of what makes the facts in question the reason they respond to.

Further, we may say Smart responds to an actual normative reason, Sam responds to a merely *supposed* normative reason: the fact she treats as a reason is not a reason. Before I explain this concept in greater detail and give more examples, let us close the discussion of **Composite**. Recall:

> **Composite**: What it is for A to ϕ for a normative reason r is for r to be A's motivating reason for ϕ-ing and for r to be a normative reason to ϕ.

With the more fine-grained understanding of normative reasons introduced above, we see that **Composite** is ambiguous. What does "r" refer to? If we read the "r" in **Composite** as referring to the fact that Barcelona won, then *Bad Case* is not a counterexample to **Composite**. Sam believes that Madrid lost for that fact, and it is a reason for him to believe that Madrid lost. If we read the first appearance of "r" as referring to the *reason* Sam has to believe that Madrid lost, and the second and third appearances as referring to the reason-constituting fact that

Barcelona won, then *Bad Case* shows **Composite** to be false. To ϕ for the normative reason **r** (given by fact f), it is not sufficient that you ϕ in light of f and f is normative reason **r** for you to ϕ. Lord concludes that **Composite** must be rejected, and that we cannot understand the nature of responding to reasons as a special case of doing or believing something for a reason. But this is not what the case shows. It shows that **Composite** is incomplete, not that it is entirely false. We cannot understand responding to normative reasons only in terms of indenting or believing things for reasons, i.e. in light of certain consideration. Condition (1) of *RNR* remains uncontested but needs to be supplemented. Responding to a consideration f as normative reason **r** requires manifesting an implicit conception of normative reasons according to which f is normative reason **r**. In other words, condition (2) of *RNR* is necessary. Whenever N comes to belief that p / to intend to ϕ in light of f as a response to a normative reason, N is

(2) thereby manifesting a conception of normative reasons according to which f is normative reason **r.**

Bad Case is a counterexample to **Composite** only if we recognise that one can respond to the same facts as different reasons. Those reasons may be merely supposed normative reasons. We can then say that Sam does not respond to the normative reason **r** she has to believe p, but to a merely *supposed* normative reason **r'**. If this is correct, then to respond to a consideration f as a normative reason, f need not in fact be a normative reason; you may respond to a merely *supposed* normative reason. **Composite** states a condition that is not necessary. *RNR* therefore does not require 'f' to stand for a fact (or apparent fact) that is (or would be if it was true) a normative reason.

In general, a *supposed* normative reason is a consideration which would be a normative reason for you to ϕ or believe that p, if it was true and the normative facts were as you take them to be. We must not confuse *supposed* normative reasons with apparent normative reasons.[24] When you respond to supposed normative reasons, you make a mistake about normative facts, properties or concepts. In contrast, when you respond to an *apparent* reason, you make a mistake about the non-normative facts in the world. You can make both mistakes at the same time. A *supposed* normative reason is not an actual normative reason that is defeated or undercut in some way. It is not a reason at all. To this one might object that responding to supposed normative reasons is not responding to reasons at all for you cannot respond to a normative reason that is not there. But this is too quick. We can make sense of responding to *supposed* normative reasons just as we can make sense of responding to *apparent* normative reasons. From the agent's perspective, responding to an actual normative reason and responding to a merely *supposed* normative reason appear exactly the same. The agent manifests the same kind of capacity in treating considerations he takes to be true as normative reasons. In both cases, the agent manifests a conception of general normative facts. If this conception is incorrect, he responds to a merely *supposed* normative reason.

4.3.2. More Examples

In the remainder of this section I give three more examples to further illustrate the phenomenon of responding to merely supposed normative reasons. Here are another two from the theoretical domain. Turri (2010: 316) gives the case of Miss Improper, who treats a set of considerations as evidence, i.e. as reasons to believe a proposition, but not as the reasons they are.

[24] My supposed normative reasons may be Pettit's (2018: 271) "self-ascribed" reasons - reasons you wrongly think you have.

"Two jurors, Miss Proper and Miss Improper, are sitting in judgment of Mr. Mansour. Each paid close attention throughout the trial. As a result, each knows the following things:

(P1) Mansour had a motive to kill the victim.

(P2) Mansour had previously threatened to kill the victim.

(P3) Multiple eyewitnesses place Mansour at the crime scene.

(P4) Mansour's fingerprints were all over the murder weapon.

<Mansour is guilty> is propositionally justified for both jurors because each knows (P1–P4). As it happens, each comes to believe <Mansour is guilty> as the result of an episode of explicit, conscious reasoning that features (P1–P4) essentially.

Miss Proper reasons like so:

(Proper Reasoning) (P1–P4) make it overwhelmingly likely that Mansour is guilty. (P1–P4) are true. Therefore, Mansour is guilty.

Miss Improper, by contrast, reasons like this:

(Improper Reasoning) The tea leaves say that (P1–P4) make it overwhelmingly likely that Mansour is guilty. (P1–P4) are true. Therefore, Mansour is guilty."

Miss Improper believes that Mansour is guilty because of (P1) – (P4). Turri argues that she "does base her belief on (P1–P4) *as evidence for* <Mansour is guilty>. It is part of the content of her reasoning that (P1–P4) make <Mansour is guilty> overwhelmingly probable" (Turri 2010: 317). Miss Improper treats considerations which are reasons for her to believe that Mansour is guilty *as reasons* to believe that Mansour is guilty, but she does not treat them as the reasons they are. She makes the same kind of mistake as Sam and responds to merely supposed normative reasons. In contrast, Miss Proper is like Smart: she treats the facts as the reasons they are.

Now consider Lucy and Lilly. Both think: "I have never seen a unicorn" and conclude that they will never see a unicorn. Both treat the fact that they have never seen a unicorn as a reason to believe that they will never see one. But they have different conceptions of why the fact they respond to is a reason to believe that they will never see a unicorn. This is so because they have different conceptions of why the fact is evidence that they will never see a unicorn. That some fact f is evidence for a proposition p is to say that there is a correlation relation between the events described by f and p. Lucy and Lilly have different conceptions of what the evidence is because they have different views about the causal link underlying the presumed correlation relation between the appearance and the existence of unicorns. Lucy takes the fact that she has never seen a unicorn as evidence to believe that she won't see one because there are no unicorns. Lilly believes that the fact that she has never seen a unicorn is evidence to believe that she won't see one because unicorns are extremely good at hiding. Lucy and Lilly treat the same fact as a reason for belief, but they do not treat it as the same kind of reason.[25]

[25] I thank Sven Neth for the discussion of this example.

Now let's consider two cases from the practical domain. Analogously to the previous two examples, they show that responding to a fact as a normative practical reason is not necessarily the same as responding to that fact as the normative practical reason it is. First, there is Good Friend and Odd Friend. Both receive an invitation to a birthday party from a common friend. Good Friend thinks: "He invited me to his birthday party. So, I shall go!" She treats the fact that she was invited as a reason to go, thinking about her friend, and that he will be sad if she does not go. Odd Friend, upon receiving the invitation, also thinks: "He invited me to his birthday party. I shall go then." She also treats the fact that she is invited to the party as a reason to go but has nothing about friendship in mind. She believes it will make her feel appreciated and popular to go to the party. Both, Good Friend and Odd Friend respond to the same facts, and treat the same consideration as a normative reason, but they treat it as different reasons.

Here is a final example: Suppose Hedonistic Hanna treats the fact that a certain action will be pleasurable as a reason to do it. Ascetic Anna, who is in the same situation, treats the fact that a certain action will be pleasurable as a reason not to do it. Both respond to the fact that the action will be pleasurable, but they respond to different normative reasons. They treat the same fact as reasons of different valence. Who (if any) of them responds to an actual, and who to a merely supposed normative reason?

This concludes my discussion of supposed normative reasons. One responds to a supposed normative reason if one responds to a consideration f (which may be false) as a normative reason, where f is not a reason (and would not be one if it was true). That is to say that in responding to f, one manifests a false conception of normative reasons. For responding to reasons, I have argued, requires manifesting a conception of normative reasons. This is captured in condition (2) of *RNR*:

Responding to Normative Reasons (RNR):

Agent N's belief that p / intention to ϕ is a response to consideration f as normative reason **r**, if and only if

(1) N comes to believe that p / to intend to ϕ in light of f (and thus on the basis of the belief that f),

(2) thereby manifesting an implicit conception of normative reasons according to which f is normative reason **r,**

(3) through a process that is under N's (virtual) agential control.

An implicit conception of normative reasons consists of a set of implicit beliefs about what makes a fact a normative reason. In this section and the next, I elaborate on the contents of these beliefs. Section **4.4** concerns responding to reasons for action, section **4.5** responding to reasons for belief. I then turn to the question of what role these beliefs play in the process of responding to reasons. Once I have explained what having an implicit conception of normative reasons for action of belief is, I explain what it is to manifest it in responding to considerations as reasons in section **4.6.** This question will lead me to define condition (3) of *RNR*.

4.4. A Conception of Rightness

Responding to a fact as a normative reason for action requires an at least implicit conception of why the fact is the reason for action. If N responds to fact f as normative reason **r** to ϕ, N manifests an at least implicit conception of why f is normative reason **r** for N to ϕ. (This is not to deny that the conception can be explicit. Nevertheless, I take it that such conceptions are mostly implicit, which is all that is needed for responding to reasons. In what follows it will be convenient to speak only of implicit conceptions.) The content of this implicit conception is that which explains why a fact, say f, is a reason for N to ϕ. More precisely, it

explains why f is a *partial* ground of the reason for N to ϕ.[26] I say 'partial' ground, because f may not be the full ground of the reason. Additional facts, which play the role of conditions or modifiers, may be necessary to make it the case that there is a reason for N to ϕ. We commonly point out only the most prominent, unusual or salient of these facts, and call this fact (f) the reason. So, the correct question would be: How does a set of facts make it the case that f is a reason for N to ϕ? For the present purpose, however, we may leave the technicalities aside and simply ask: What makes it the case that f is a reason for N to ϕ?[27] In what follows, I explain how an answer to this question may be structured, without attempting to explain any particular reason for action that an agent may have. That would be to engage in first-order normative theorizing, which is not my project here.

To begin with, the notion of a normative reason for action is not an evaluative notion, but a deontic notion in the sense that it issues a directive and is related to 'ought'.[28] Reasons determine what you ought to do, that is, what is right to do. We can have different theories about how they do that. Let us assume, for the sake of argument, that what is right to do is determined by what is good. Rawls (1971: 22) calls such a theory a teleological theory. It says that "the right is defined in terms of the good", and "the good independently of the right." What is right to do is what promotes or honours what is good (Pettit 1993: 301-302).[29] On this view, normative reasons for action are to be understood as good-makers (and only in a

[26] As mentioned in chapter 3, note 15, I use the grounding relation or 'making it the case' - relation here as a placeholder for an explanatory relation, the investigation of which is a different project.
[27] I take for granted that a fact, as well as a collection or "fusion" of facts, can ground there being a reason for N to ϕ (Fine 2012: 72).
[28] A recent debate centres on how exactly they are related, see chapter 3, note 28.
[29] A conception of goodness involves a conception of value bearers. I assume here that states of affairs are the bearers of value or goodness (Broome 1991:3). It leaves open how exactly we are to interpret sentences like 'happiness is good', but they mean something like 'a state of affairs is better the more happiness it has'. An alternative would be to say that abstract objects like fairness, courage, kindness, trustworthiness, pleasure, etc are the primary bearers of value and that the world is better the more there is of those things. This is sometimes referred to as a "value-first" theory (Schroeder 2016).

secondary sense, as right-makers). If the fact f is a reason for N to ϕ, then the outcome of N's ϕ-ing is valuable or good in some way because of f. Reasons are facts that make it the case that certain action would lead to an outcome that is good in a certain way. Suppose, for example, that the fact that the girl would otherwise drown is a reason for you to jump in the pond and save her. The reason makes salient in what way the outcome of your action would be good: it would save the girl's life.

Now, if fact f is a reason for N to ϕ, and we ask why that is so, then one immediate explanation refers to (i) something that N's ϕ-ing brings about (because of f) and (ii) the claim that a state of affairs with that something is good to some extent (or, in short: that these things are good).[30] In other words, one way of explaining why f is a reason for N to ϕ is by reference to (i) certain properties and (ii) value claims about these properties. For example, if we ask why the fact that the child would otherwise drown is a reason for N to jump in, then one answer is: because (i) the action would bring it about that her life is saved and (ii) a state of affairs in which a life is saved is good (to a certain degree).

How do (i) and (ii) explain why f is a reason? They explain why f is a reason for N to ϕ by explaining why N's ϕ-ing would bring about something good that p makes salient.[31] The kind of explanation is a grounding explanation of the goodness of the outcome of N's ϕ-ing; it states what it is for this state of affairs to be good, what constitutes its goodness. The claims of (ii) can be seen as normative laws: general normative facts that link non-normative properties with normative ones. The normative property at stake here is the property of being

[30] Korsgaard (1996:226-228) argues the notion of a reason comes with idea that "normativity must be captured in regularity or law-like formulations." When I choose to act (on a desire), she says, I commit to acting in the same way in any similar situation. Agency, in particular, responding to normative reasons, cannot be explained "in absence of at least a claim to law or universality." I here agree with this and explain why that is so.

[31] The 'explain' in this sentence takes a wide scope: (i) and (ii) explain the explanatory relation between p and the goodness of the outcome of N's action.

good (to a certain degree). The role that general normative facts play in the explanation of particular normative judgements, like the goodness of the outcome of N's ϕ-ing, is up for debate. They may play the same role as the facts falling under (i). General normative facts together with claims about instantiations of non-normative properties explain why something is good, and thus why a fact is a reason (Rosen 2017). Alternatively, the general normative facts or laws may explain why facts of kind (i) can explain why a fact is a reason (Bader 2017).[32] Note also that in some cases, the considerations we respond to as reason for action already make explicit how the action would promote the good. For example, the fact that it would be unfair to do otherwise is a reason for the shopkeeper to give back the correct change. An explanation of why that fact is a reason would refer to the goodness of doing what is fair (ii).

The upshot is that our conception of normative reasons involves claims about general normative facts, which we may call value claims. An example of a value claim is the claim that fairness is good, or the claim that happiness is good. Understanding value claims requires an understanding the concepts they involve, such as happiness, pleasure, fairness, promise-keeping, courage, etc. If you believe that happiness is good you must have a conception of what happiness consists in. The set of value claims you believe, implicitly or explicitly, together with your conceptions of the things you believe to be good, form your conception of goodness.

With that, let us return to the example of the Good Friend and Odd Friend. They both treat the consideration that they received an invitation to a birthday party from a friend as a reason to go to the party. But they treat it as different reasons, because they thereby manifest different

[32] If the latter is correct, the metaphysics of normative properties seem to mirror the psychology of responding to normative reasons: That saving lives matters is not a fact we take to be a reason, but a fact that allows us to explain why things are reasons at all. We have to have a common ground of what matters to engage in the debate about what, and why facts are reasons.

conceptions of goodness. Good Friend manifests a conception of goodness according to which friendship is good, and a conception of friendship according to which keeping the friendship requires attending the birthday party. Good Friend sees the fact that she is invited to the party as a reason to go because this is what friendship requires, and friendship matters. This is what she has in mind when she treats the invitation as a reason to go to the party. Odd Friend, in contrast, as no such conception of goodness. As it so happens, going to the party, given that she was invited, will make her feel loved and appreciated. Odd Friend doesn't go to the party because she has friendship in mind, but because she has her well-being in mind.[33]

The assumption of teleology, that the right is determined by the good, is not essential to my argument. It only allowed me to illustrate what it is to have a conception of normative reasons for action. We can also state the argument on a non-teleological framework, according to which the right is not defined in terms of the good. In short, this would go as follows. Reasons are right-makers. If the fact that f is a reason for N to ϕ, then f makes N's ϕ-ing right in a certain way, to a certain degree. Reasons are right-makers. When we ask why a fact is a reason to ϕ, then one explanation refers to (i) a property about the action and (ii) the fact that actions

[33] One might object that Good Friend responds to a moral reason, and Odd Friend to a prudential reason. But responding to a prudential reason is also a response to a normative reason (which may be a merely a supposed normative reason). It involves treating considerations as reasons for one to ϕ just as responding to moral reasons does. It is different from, for example, acting out of a desire. When I mindlessly eat another gummy bear, my action can be explained by my desires and beliefs. But my action is not a response to a consideration as a normative reason. By contrast, when I treat the tastiness of the gummy bear as a prudential reason for me to eat it, I manifest a conception of goodness. Not moral goodness, but prudential goodness. In any case, the distinction between of moral and prudential reasons and different kinds of goodness (morality and prudence or happiness) raises the question of how the two different kinds of goodness can be combined to allow an assessment of overall goodness. It raises what Sidgwick(1874) called the problem of "dualism of practical reason." As Bader (2015:177) describes it, "the problem is that the radical heterogeneity of these two types of value implies that they cannot be ordered or weighed up against each other." It seems hard to deny that we recognise moral and prudential reasons. Our outlook to the practical reasons we have seems inevitably divided in the moral and the eudemonistic point of view (Wallace 2006:130). To get a sense of what we are overall required or permitted to do, we must integrate prudential and normative reasons, and thereby manifest a conception of the relation between prudence and morality. Sidgwick raises the question of whether there is coherent conception. Whether there is a coherent conception, and how we should respond to moral and prudential reasons is a substantial question I do not address here. The only claim I make is that we can respond to a prudential reason for action *as* such a normative reason.

with this property are right (to a certain degree). For example, when we ask why the fact that $5 is the correct change is a reason for the shopkeeper to return $5 change to the blind customer, one answer is: (i) because it the fair thing to do, and (ii) fair acts are right (to a certain degree). That something is fair is a reason to do it only on the presupposition that fairness matters. On this view of the right, a conception of normative reasons consists in beliefs about what kind of actions are right (e.g. promise-keeping, fair actions, saying the truth, etc.) Responding to normative reasons for action requires an implicit conception of rightness.

How we describe our implicit conception of normative reasons depends on our views about the relation between the right, the good and normative reasons. On a teleological view, reasons are good-makers because what is right is determined by what is good. A conception of normative reasons requires a conception of goodness because it is part of a conception of rightness. On a non-teleological view, reasons are right-makers and a conception of normative reasons requires a conception of rightness. In ordinary life, of course, we don't make these distinctions, and our implicit beliefs about goodness and rightness that guide our responses to the world may not be coherent. We assume that we ought to do what is best, and at the same time treat duties, obligations and rights we take ourselves to have as reasons for and constraints on our actions.[34] For example, we may say with Ross (1930: 19) that "normally promise-keeping, for example, should come before benevolence, but that when and only when the good to be produced by the benevolent act is very great and the promise comparatively trivial, the act of benevolence becomes our duty."

[34] Some argue that we can restate any non-teleological theory as a teleological one. If that is so, the distinctions I have drawn here can be explained away, see e.g. Portmore (2009) and Brown (2011) on 'consequentialising'.

The purpose of this excursion into the explanation of normative reasons was to elucidate condition (2) of *RNR*, according to which responding to a fact f as a normative reason requires manifesting an implicit conception of that which makes f a reason **r**. It is not sufficient to believe that the fact f is a reason for or favours the action. Given the analysis of this section, we can rephrase condition (2) for responding to normative reasons for action as follows: N manifests an implicit conception of rightness (or goodness).[35] I have laid out here the structure of the content of this conception. What is underlying our responses to facts as practical reasons is what we colloquially call our values. It is our set of (implicit and explicit) beliefs not just about what is right or good in a given context, but about what makes it so. Next, I explain the structure of the content of our implicit conception of normative reasons for belief.

4.5. A Conception of Epistemic Value

Responding to reasons for belief requires an implicit conception of normative reasons for belief, the content of which gives a characterisation of that which makes it the case that f is a normative reason for N to believe that p. If f is a normative reason for N to believe that p, what makes it so?[36]

One kind of answer appeals to the following: (i) because of f, the belief that p has a certain property, and (ii) beliefs with this property are epistemically good or right. To explain why a

[35] Velleman (1996: 723) says that the inclination towards the constitutive goal of action [i.e. the good] mediates the influence of your reasons for action, just as the inclination towards the truth mediates the influence of your reasons for belief." On my proposal, it is not an inclination, but more cognitive attitudes - those that constitute our conception of rightness or goodness- that mediate our responses to facts as practical reasons.

[36] I leave aside that f only partly grounds the epistemic reason for to N. There may be conditions and modifiers. As I mentioned in section **3.4**., a fact is an epistemic reason for N to believe p only if N has non-epistemic reasons to have any belief about p (Wallace (2006) makes a similar point). Epistemic reasons are conditional on non-epistemic reasons.

fact is an epistemic reason for belief, we have to appeal to claims about *epistemic value* and what instantiates it.

Regarding (i), one might appeal to the belief's being in accordance with the evidence provided by f, to the content of the belief being logically entailed by f or to the fact's being grounded in the fact that f. Persons may differ in what they take evidence, entailment and constitution to consist in. What is evidence for what? What are the laws of probabilification? Which deductively valid entailments are truth-indicating and justification transmitting? Which facts ground which? All these are substantive claims that we do not need to answer here, just as we did not have to answer any questions about which states of affairs are good, or which actions right.

Regarding (ii), that the belief that p has a certain property or stands in certain relation to (the belief that) f is not sufficient to explain why f is a reason for believing p. We need the claim that (ii) beliefs with this property epistemically ought to be held; the properties manifest a certain kind of epistemic value. The 'ought' is the epistemic ought, not the practical ought. It does not have to do with practical value, but epistemic value. Analogously to the practical case, we may disagree about the role of statements falling under (ii) in the respective explanations. Are they part of the explanans or do they explain why the other explanans are explanans?

The important point here is that there are some beliefs you epistemically ought to have, and some you ought not to have. If you epistemically ought or have sufficient epistemic reason to believe that p, then that fact is at least partly grounded in a non-normative fact. What exactly this non-normative fact is, is subject to debate. But most agree that it involves a property that the belief p has because of f, or a relation between the fact that f and the (content of the) belief p. Most also agree that you have to have some epistemic access to f. Maybe we can always state the relevant property in terms of truth. Maybe what matters is that f indicates the truth

of p. If f is evidence for p, or if f is entailed by p, for example, then f indicates the truth of p. On this view, what explains why f is an epistemic reason to believe that p is that (i) f indicates the truth of p, and (ii) that believing what the (supposed) facts indicate to be true is epistemically good.[37] The notion of 'truth indication', however, would have to be elucidated. In any case, we can assume that there is some way in which we can explain why a fact is a normative reason for belief. I shall not pursue this further here. Even though we mostly agree as to which facts are and which facts aren't reasons for belief, it proves hard to explain what exactly makes these facts reasons for belief. I leave this for another occasion, and now return to the example from section 5.

In the case of Sam and Smart, both treat the fact that Barcelona won as a reason to believe that Real Madrid lost, but they treat it as different reasons. Sam takes the fact that Barcelona won to entail that Real Madrid lost. He manifests a conception of entailment (or truth indication) according to which the fact that Barcelona won entails the truth of any proposition (or indicates the truth of any proposition). Smart, in contrast, manifests no such conception of entailment (or truth indication). They may not differ in their conception of what matters epistemically. They may have no normative dissonance: one ought to believe what is entailed by the facts (what the facts indicate to be true). Their disagreement, it seems, is only about the non-normative question of what entails what.

Finally, let us return to Turri's case of Miss Proper and Miss Improper. Both see (P1) – (P4) as reasons to believe p. But they have a different conception of why (P1) – (P4) are reasons for belief, and thus treat the facts as different reasons. Both, I want to say with Turri, treat the

[37] We could see this general version of (ii) as the general 'norm of belief'. Problems arise when we try to precisely formulate an acceptable version of the norm of belief in terms of truth, however. The that "you ought to believe p if only if p is true" is too strong. Formulations of the norm of belief in terms of evidence, which go back to Clifford (1876), seem much more plausible: you ought to believe that p if only if you have sufficient evidence for p (see e.g. Gibbons (2013: ch.1) for a comparison of the norm of belief in terms of truth and in terms of evidence). So, how exactly truth figures in the kind of explanation of reasons for belief I propose here, remains to be determined.

facts as evidence. But they have different conceptions of evidence consists of. While Miss Proper believes (P1) to (P4) are evidence for p because of some underlying correlation relation, Miss improper takes (P1) – (P2) to make p likely to be true because the tea leaves say they are. Again, we can explain the difference between the Miss Proper and Miss Improper without assuming a disagreement about normativity. They may both have an implicit conception of a reason for belief which includes that we ought to believe what the facts indicate is true, while they differ in their beliefs about what "truth indication" is.[38]

In summary, I have argued that responding to a consideration f as normative reason **r** to believe that p requires manifesting an implicit conception of normative reasons for belief according to which f is normative reason **r**. The content of this conception may be understood as including (i) that f indicates the truth of p, and (ii) that believing what the (apparent) facts indicate to be true is epistemically good. Let's say for short that this constitutes a conception of *epistemic value*: it's a conception of what kinds of beliefs are epistemically good to have.[39] One's conception of epistemic value plays the same role in responding to considerations as reasons for belief as one's implicit conception rightness (or goodness) in responding to reason for action. Recall the account of

> *Responding to Normative Reasons (RNR):*
>
> Agent N's belief that p / intention to ϕ is a response to consideration f as normative reason **r**, if and only if

[38] Worsnip (2016: 251) recognises that one can "fail to respond one's epistemic reasons" not because "one rejects the norm of believing what the evidence supports, but by having incorrect evidential standards." I explain here why, in those cases, one may still count as responding to reasons.

[39] Velleman (1996: 706, 713) argues an inclination to believe what is true is necessary for responding to normative reasons for belief. My account implies that an *inclination* is not sufficient for responding to epistemic reasons, just an *inclination* to do what you have reason to do is not sufficient for responding to practical reasons.

(1) N comes to believe that p / to intend to ϕ in light of f (and thus on the basis of the belief that f),

(2) thereby manifesting an implicit conception of normative reasons according to which f is normative reason **r,**

(3) through a process that is under N's (virtual) agential control.

The previous two sections aimed to elucidate condition (2) by explaining how we can understand the content of the implicit conception of normative reason for belief or action. I have not explained what it is to manifest this conception, however. What role do the attitudes constituting this conception play in the mental process leading to the response? Well, if N responds to f as a reason to believe p by forming the belief that p on the basis of believing f, an explanation of this basing must refer to N's implicit conception of epistemic value. Analogously, whenever N responds to f as a reason for ϕ-ing by forming the intention to ϕ on the basis of the belief that f, N's implicit conception of rightness (or goodness) partly explains why N comes to intend to ϕ on the basis of the belief that f. The next two sections say more on this. They are dedicated to explaining the nature of the process of responding to normative reasons and thus condition (3) of *RNR*.

4.6. The Capacity to Respond to Normative Reasons

You can respond to considerations as normative reasons for you to do or believe things without active reasoning. As mentioned in section 2, "not all ways of coming to R when taking F to be a reason for it are cases of reasoning" (Raz 2015:153). For some response to count as a response to reasons, you don't have to actively reason to the response. The process could be passive, an automatic manifestation of an acquired habit. However, the response could not just be the result of mere sub-personal causation. The process must be in some way

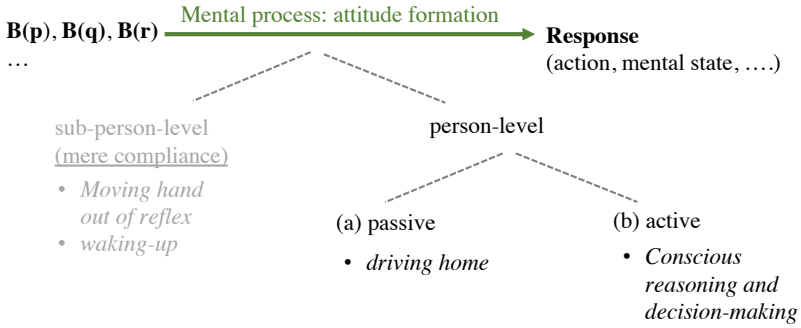
attributed to you as an agent. Even if conditions (1) and (2) are satisfied, and even if the consideration you treat as a reason for your response actually is that reason, you may comply with what that normative reason demands without responding to it.

Suppose, for example, you touch a hot surface. That it is hot is a reason to pull your hand away. You immediately withdraw your hand out of reflex. You cannot help but pull your hand away. You have a belief about a fact that gives you a reason, you respond with an action and the link between the belief and the response seems non-deviant in the right kind of way. But you do not respond to the fact as a normative reason - even if you comply with what it demands. You are lucky that your reflexes function well. But this is not an achievement of you as a rational agent. And would your reflex not function as it should, we would not criticise you for failing to respond to a reason (correctly). What distinguishes responding to reasons from mere compliance with reasons?

The difference is that only if the process is a person-level activity that is under your agential control, or at least under your virtual agential control, do you count as responding to, rather than merely complying with normative reasons. A reflex is a sub-person level mechanism that automatically causes you to act, so that the process steering the action is not under your agential control. That a process is under your agential control does not imply that it is the result of your exercising this control in that particular instance. But if a response is under your agential control, the process is deliberatively accessible to and re-constructible for you. For example, when you drive home from work while thinking about the climate crisis, you respond to many reasons regarding the driving without actively thinking about them. You respond to facts as reasons for you to stop, take a turn, accelerate, etc. You do so automatically, having internalised the procedures. But the process is still under your *virtual*

agential control, that is, control you do not actively exercise.[40] You could change your behaviour, if you chose to do so. If you were to focus on the driving, you could think about the situation and consciously and actively respond to the reason you have. In contrast, when you withdraw your hand out of a reflex, you could not at will initiate or deactivate your response to the hot sensation. (Of course, you could also respond the hot sensation as a reason to withdraw your hand, but then you would not act out of reflex. This would be a different scenario.) What indicates that the process is under your virtual control is that if something went wrong, you could and – if all goes well - would step in. If you accidentally took a wrong turn, your active reasoning would take over control and get you back on track.

The following figure summarises these remarks.



There two types of person-level processes which may underly responses to normative reasons: one is active and one is passive. Both are under your agential control, but only in the case of the former do you exercise this control. When the process is passive, you automatically respond to reasons, manifesting internalised patterns of behaviour. Sub-person level process, such as the process of a reflex or that of falling asleep, do not lead to responses

[40] I learned the concept of (virtual) agential control from Philip Pettit during his 'Locke Lectures' at Oxford in 2019. See also Pettit (2007a: 226-7).

to reasons. This is why condition (3) of *RNR* is necessary. This completes my presentation and defence of

> *Responding to Normative Reasons (RNR):*
>
> Agent N's belief that p / intention to ϕ is a response to consideration f as normative reason **r**, if and only if
>
> (1) N comes to believe that p / to intend to ϕ in light of f (and thus on the basis of the belief that f),
> (2) thereby manifesting an implicit conception of normative reasons according to which f is normative reason **r,**
> (3) through a process that is under N's (virtual) agential control.

But now we may ask: When you actively respond to a normative reason, exercising your agential control, what does this mental action consist in? What is the capacity you exercise? The next section argues that it is the activity of reasoning. Our responses to normative reasons are under the control that we have over our cognitive life in virtue of our capacity to reason.

4.7. Reasoning and Responding to Reasons

In the introduction of this chapter, I said it has two aims. The first was to develop an account of what it is to respond to normative reasons. I now argue that reasoning – on the *Approval Account* defended in chapter 3 - is a way of responding to normative reasons. Whenever you reason from the belief that f to the intention to ϕ or the belief that p, you count as responding to f as a normative reason. In other words: If a transition counts as an instance of reasoning according to the *Approval Account*, then it satisfies conditions (1) to (3) of *RNR,* so that the conclusion-attitude counts as a response to normative reasons.

The first condition of *RNR* states that your belief that p or intention to ϕ is a response to a normative reason only if you form it in light of f. In chapter 3, I have assumed that when you reason, you operate on the contents of your premise-beliefs as you perceive them and come to intend or believe something in light of them. In reasoning, the conclusion-attitude is based on the premise-attitudes. I take it that condition (i) is satisfied.

The second condition of *RNR* requires that you manifest an implicit conception of normative reasons according to which f is a normative reason. The beliefs that form this implicit conception play a role in the mental process leading to the response – they partly explain its being based on the belief that f. I argued in chapter 3 that reasoning essentially is a mental process you approve of. It is your approval of the transition from the first-person premises to the conclusion-attitude which makes the process count as an instance of active reasoning. The attitude of approval is evaluative in nature: you evaluate your behaviour against the standard given by your implicit conception of what makes what kind of response permissible. This is what your approval is based on. When you reason to the belief that p or the intention to ϕ on the basis of the belief that f, you therefore manifest this conception through your approval. Your conception of what makes what kind of response permissible is your conception of normative reasons and how they aggregate. So, when you reason from the belief that f to the intention to ϕ or the belief that p, you manifest this conception: it partly explains why and how the response is based on the belief that f. Therefore, when you reason from the belief that f to the intention to ϕ or the belief that p, condition (2) of *RNR* is satisfied: you manifest an implicit conception of normative reasons.

The third condition of *RNR,* which requires the process leading to the response to be under your (virtual) agential control, is obviously satisfied when you reason to the response. In section **3.8**, I have argued that on the *Approval Account*, reasoning is an intentional activity.

When you reason from the belief that f to a conclusion-attitude, the process is a paradigmatic example of a mental process that is under your agential control over your mental life. It is something you *do*.

In conclusion, *RNR* and the *Approval Account* imply that when you reason to the intention to ϕ or the belief that p on the basis of the belief that f, you respond to f as a normative reason. The capacity to reason is the capacity to respond to normative reasons. When you reason, you manifest not just an implicit conception of normative reasons, but also of how they aggregate. You need a conception of what constitutes sufficient reason for a response. In order to count as responding to a consideration as a normative reason, you do not need to manifest this more complex conception. Through reasoning, you can engage in what we call the weighing of normative reasons (see chapter **5**), which allows you to (correctly) respond to a consideration as a normative reason to ϕ with the intention to not-ϕ, namely when that intention is also based on other considerations you treat as reasons to not-ϕ. In any case, you can respond to a consideration as a normative reason without weighing it against other reasons. To put it in a slogan: to respond to reasons, you need to have values. To reason, you need to have a value system, a ranking of how different things that matter relate.

So, when you reason, you respond to reasons. You can respond to reasons without reasoning, but when you actively respond to normative reasons, do you then *reason* to the response? Plausibly yes, for both are mental activities through which you acquire attitudes on the basis of others, thereby manifesting your implicit conception of normative reasons. The question is whether this manifestation is of the same kind in both cases. I suspect it is.

In reasoning you manifest an implicit conception of normative reasons through the attitude of approval; you approve of the process, and the basis of this approval is your implicit conception of normative reasons. The attitudes constituting it are implicit premises of the reasoning (cf. **3.4**). The conclusion-attitude (or response) is not directly based on the set of

attitudes constituting the implicit conception of rightness. This is plausible, for you do not operate on and respond to the contents of these attitudes. Instead, the involvement of these attitudes determines the way in which you respond to the features of the situation you consider; they guide you.

Regarding responding to reasons, I only said that if an attitude is a response to consideration f as a normative reason, then how this attitude is based on your beliefs that f cannot be explained without reference to your implicit conception of normative reasons. But what role do the attitudes constituting this conception play in *how* the response is based on the belief that f? If we assume, as I suggest we do, that the way in which you manifest your implicit conception in actively responding to reasons is the same as in reasoning, we get an answer. You approve of the process, and your implicit beliefs about rightness (or goodness) and epistemic value are the basis of this approval. This approval – as I argued in chapter **3** – changes the quality of the process. It makes it something you do. So, the assumption does not only allow us to explain what role the attitudes constituting the implicit conception of normative reasons play when you respond to reasons (how exactly you manifest this conception). We can also explain what you do when you "actively" treat things as reasons: you reason on the basis of beliefs about them. This seems a plausible and coherent picture.[41]

This leads me to conclude that the capacity to reason is the capacity to respond to normative reasons. Reasoning is a way of responding to normative reasons, and sometimes we respond to reasons through reasoning. But responding to reasons does not require the active exercise of the capacity to reason. You can respond to normative reasons habitually, manifesting dispositions. The thought processes are still under the control given by your capacity to

[41] I criticised in chapter 2 that Kauppinen's account does not explain the claim that treating premises as reasons in reasoning is having "a positive attitude towards performing the act of forming the intention, given the premises, and negative attitude towards not doing so" (Kauppinen 2018: 419). I have now established this claim with the *Approval Account* and *RNR* and explained the nature and role of the attitude of endorsement.

reason, though you do not exercise it. You don't intentionally approve of the transitions you make, but 'virtually' you do. That is, you would approve of them, if you attended to them. As long as this is so, you do not waste any cognitive resources and let your internalised habits work for you. So even if you do not actively reason to a response, you manifest an implicit conception of normative reasons. In some sense, condition (2) is this entailed by condition (3) of *RNR*: you manifest an implicit conception of normative reasons in virtue of the process being under your agential control, that is, by your capacity to reason.

We need normative thoughts to see things as reasons, and to reason. I do not speculate about an order of priority or an explanatory relation between reasoning and normative thought.[42] All I say is that without the capacity to make sense of an 'ought' in contrast to an 'is', we would not be the kind of reasoners we are.

4.8. Responding Correctly to Normative Reasons

I was primarily concerned with what it is to respond to a consideration as a normative reason. But I have mentioned in passing that this responding may be correct or incorrect. For example, I said Sam and Miss Improper incorrectly respond to reasons; they treat considerations as normative reasons but fail to treat these as the reason they are. They respond to merely supposed normative reasons. I shall here just very briefly explain how we can make sense of this.

In general, we could say that agent N responds correctly to consideration f as the normative reason **r** if and only if N treat f as the reason it is (**r**). That is, N responds to reasons correctly

[42] Pettit's view, I believe, is that the capacity to reason (which "is put on stream by our capacity to speak") enables us to think in a normative way (Locke Lectures 2019, Oxford). He thinks there could be beings similar to humans without prescriptive concepts (Pettit 2018: 49). I doubt that. Such beings could not reason the way we do, for reasoning requires prescriptive concepts, more precisely, implicit conceptions of prescriptive properties.

whenever the conception of rightness or epistemic value which N manifests in responding to **r** is correct. What is it for this conception to be correct? My rough suggestion is that it is correct whenever it satisfies certain criteria or is 'good enough'. One agent's conception of rightness may be *better* qua conception of rightness than another's conception of rightness. It seems this is so when it, for example, gives a more accurate description of that which explains why certain facts are reasons and why certain responses are permissible, when it contains more true beliefs, is more fine-grained, more robustly held, or better understood.

This standard of correctness for responding to normative reasons proposed here is an objective, externalist standard. Put roughly, you respond correctly to considerations if and only if you treat them as the reasons they are. That is, you manifest a reasonably accurate conception of rightness or epistemic value. This assumes that there are facts about these matters, which we can correctly or incorrectly represent. It is not the notion of responding to normative reasons, but the notion of responding to normative reasons correctly (or incorrectly) which commits us to the existence of normative facts. With *RNR*, we can make sense of responding to reasons even if there are no normative facts. Anti-realist theories of normativity can still make sense of the phenomenon we describe as responding to reasons. These views don't have to deny that we respond to considerations as reasons, only that there is such a thing as responding correctly to normative reasons.[43] However, if we deny that there are facts about what is right to do and believe, if nothing matters, then there is nothing to guide our thought and action, and we effectively deny ourselves the capacity to reason.

One could, alternatively, also take an internalist or coherence–based approach.[44] You

[43] Broome (2013:101-102) argues that there is no such thing as responding correctly to *a* reason. For example, if you have a reason to ϕ and you intend to ϕ for that reason, you do not necessarily respond correctly. For example, he says, you do not respond correctly, if that reason is outweighed by a stronger reason not to ϕ. I agree with the first part. But we must distinguish between responding to a reason, responding to all your reasons, and doing what you ought to do. You can respond correctly to *a* reason to ϕ (in virtue of manifesting a correct conception of rightness), while being criticisable for being ignorant of facts which are reasons not to ϕ.

[44] Coherence could be understood, for example, as Worsnip (2018) proposes.

respond correctly to a fact as a reason if and only if your response is in line with your conception of rightness and epistemic value. It is a matter of meeting your own standards. But this approach would rule out the possibility of incorrectly responding to reasons. For, at least on the account I defend, when you respond to a consideration as a reason **r**, your response is by default in coherence with your conception of rightness, which you manifest in responding as you do. Responding to reasons would entail responding correctly. But the very notion of responding correctly to normative reasons is sensible only if there is such a thing as responding incorrectly. The coherentist approach is thus implausible. I conclude that you respond correctly to reasons if and only if you manifest a correct conception of rightness (goodness) and epistemic value.

4.9. Explaining the *Normative Leap*

Only those actions and beliefs which we can attribute to our implicit conceptions of rightness and epistemic value are responses to normative reasons. To treat a consideration as a normative reason, it is neither necessary, nor indeed sufficient, to explicitly believe that this consideration is a reason. According to my proposal, when we respond to facts as normative reasons for action or belief, we manifest implicit general normative beliefs about what it is for things to be good or right, or what kind of beliefs we ought to have. We count as responding to reasons, rather than as reacting to the respective descriptive reason-constituting facts, when we are guided by our implicit conception what is right to do or believe and why. They partly explain why we respond as we do.[45]

[45] Does the account of reasoning and responding to reasons I propose commit us to the doctrine of the 'Guise of the Good'? Does it commit us to the claim that intentional actions are actions in which the agent sees some good? In note 16 above, I suggested that one could understand conditions (2) and (3) of *RNR* as specifying condition (1). Acting in light of f then implies that you see f as a reason to ϕ, and thus that you treat f as making ϕ-ing valuable, good or right in some way. In other words, you act under the guise of the

Articulating our implicit, mostly unconsciously acquired conception of what makes what right to do or believe may be difficult, and few make the attempt. What we treat as a normative reason to believe something is influenced by what we (implicitly) believe epistemically matters. People largely agree that this is, in some sense, truth. We acquire a conception of what gets us to the truth as we learn about correlation relations and evidence, understand the nature of entailment and truth-preservation (without necessarily learning the words for it), learn whom to trust as testifiers, become proficient language speakers and learn the relations between concepts. This learning is largely automatic. It unfolds with our hard-wired cognitive capacities, independent of the particular cultural practices we grow up in. In contrast, our implicit conception of the rightness of actions– at least at first - largely consists of beliefs about cultural practices we observe in our societies and the desires of individuals around as. Our disagreement about what is and is not a reason for a certain action can be explained (even if not justified) by this factual, not normative, 'cultural relativism'.

Our implicit conceptions of rightness, goodness and epistemic value determine what we see as reasons. It is not that we count as having certain normative beliefs because of certain transitions of thought we make. One might argue that I count as holding a certain conception of rightness – that I value certain things – in virtue of treating certain facts as reasons for certain actions. But this would be false. While my actions and responses to certain situations *make explicit* my values, they do not *constitute* it. Instead, I only count seeing certain facts as (supposed) reasons because of my conception of rightness. The same transition of thought

---

good. One might further argue that acting in light of f is acting for the reason that f. Next, some argue that acting intentionally is acting for a reason or in light of a consideration (for example, Anscombe (1957:§5), Raz (2011: 13, 60, 98), and Davidson (2001a:6); Hyman (2015: 140-3) denies it). If we accept these claims, then my proposal commits us to the doctrine of the Guise of the Good about intention.[45] I take this conditional implication of my proposal to be plausible. If intentional actions are responses to (supposed and/or apparent) normative reasons, then when I act intentionally, I do so in the guise of the good. Note that my proposal does not commit us to claims about the relation between desires and intentional action, which Velleman (1992) takes to be central to the 'Guise of the Good' thesis (see also **Appendix I**), nor to the claim that the agent has a *belief* that 'there is some good in the action' they perform, which Raz (20122:62) defends.

(i.e. the same response to the same consideration) can manifest different conceptions of rightness. Two persons can treat the same consideration as a normative reason but not as the same (supposed) normative reason. Thus, if we said an agent counts as holding a certain conception of rightness in virtue of how they respond to (apparent) facts, then the content of these conceptions would be oddly underdetermined. Responding to reasons and reasoning are ways of *making explicit* our implicit beliefs about general normative facts. In practical reasoning, we make our implicit values somewhat explicit. But this does not imply that what it is to have those values just is to reason in certain ways.

With this conclusion, I am denying essential parts of, for example, Simon Blackburn's (1998:52) 'Expressivism', especially his claim that what is involved in an agent's thinking that X is good is that the agent values X (1998:50), where valuing X can be described in natural terms. However, it is worth briefly considering that in his later writings, Blackburn seems to commit to some parts of my view. For example, Blackburn (2010) argues that reasons-relations correspond to good movements of the mind (8). "We say that p is a reason for [believing] q or that the fact of x is a reason for doing y when we think it is good to infer q from p" (9). So far, I agree. But Blackburn holds that the fact p (or x) is a normative reason because we endorse or comment the corresponding movement of the mind. He says: "By invoking the relation we commend or endorse the kind of guidance of the mind that it indicates" (9). Our attitude of commendation, he says, imputes normativity to the fact. My claim is that to endorse the movement of the mind is to *see* the consideration as a reason, but that does not make it one. Blackburn and I agree again that the endorsements of our transitions are not baseless, and that whether a movement of the mind is good or bad - in my terms: correct or incorrect - depends on the relation between the attitudes (see chapter **5**). He argues that "in the case of theoretical reason, our sense of how truths relate to each other give us our

standards for good or bad movements" (11). I agree (though I think there is more to say). When it comes to the practical domain, Blackburn says that when "I commend or endorse or express approval" of a movement of thought that takes one from a belief about the world to an action, he "is entering an ethical judgment" (12). The basis of this approval are his beliefs about goodness ("I will have my own reason for [the approval]: I hold that things go better if people are guided like this"). Similarly, he says that "the addiction to turning silent radios on, is pointless, and potentially costly and irritating. So, of course, we are not inclined to endorse the movement of the mind from awareness of a silent radio to the motivation to turn it on, that manifests the compulsion." The judgment that something is pointless or costly is obviously evaluative. So again, if such evaluative beliefs determine what movements of the mind we approve of, Blackburn's view seems not far from mine. There is a standard of correctness for reasoning, which is our basis for approving movements of the mind. Unless Blackburn shows that this can be stated in fully non-normative terms, his claim that our approval of movements of the mind imputes, but does not itself involve normativity, is not established.[46]

4.10. Conclusion

In this chapter, I have developed and defended an account of what it is to respond to normative reasons. In short, I argued that responding to reasons is a mental process that leads an agent from beliefs about the world to a response (such as a belief or an intention) where this process manifests the agent's implicit conception of normative reasons and is under the agent's (virtual) control.

[46] Wallace (2015: 253) makes a similar point, suggesting that Blackburn does not seem to see the implications of acknowledging that our endorsements of movements of the mind are not baseless.

Even if we have no false beliefs, we sometimes treat things as reasons which are not reasons. Some of the reasons we respond to are merely *supposed* normative reasons. A *supposed* normative reason is a consideration which would be a normative reason for you to ϕ or believe that p, if it was true and the normative facts were as you take them to be. We must not confuse *supposed* normative reasons with *apparent* normative reasons, which are false considerations that would be reasons if they were true. What is essential to count as responding to practical or theoretical reasons is that you manifest an implicit conception of normative reasons, which consists of a set of beliefs about that which makes a fact normative reason r. When you respond to reasons for action, you manifest an implicit conception of rightness (or goodness). When you respond to reasons for belief, you manifest an implicit conception of epistemic value.

When you respond to considerations as normative reasons, the process is under your agential control. You do not necessarily exercise this control, you may respond to reasons automatically and habitually, having internalised patterns of behaviour you would approve of. When you do exercise your control, you engage in reasoning. Our responses to normative reasons are under the control that we have over our cognitive life in virtue of our capacity to reason.

Furthermore, I have argued that reasoning – according to the Approval Account proposed in chapter **3** – is a way of responding to reasons (what I have called the *Reasons Claim*). Whenever you reason from the belief that f to the intention to ϕ or the belief that p, you count as responding to f as a normative reason. The capacity to reason is the capacity to respond to normative reasons.

Lastly, I have suggested that you respond to reasons correctly whenever you manifest a correct conception of rightness or epistemic value. If reasoning is responding to reasons,

then our reasons for action and belief also set a standard of evaluation for reasoning. In the next chapter, I argue that they set the standard of correctness.

Appendix I

Three Objections to the Question of the Normative Leap

In this appendix, I respond to three objections one might have to my approach of investigating the nature of responding to normative reasons for action. Suppose Sara decides to go to the climate protest because she believes her presence will make a significant contribution to the effectiveness of the protest. She treats this fact as a normative reason to go. But in virtue of what? One might argue that this question arises only if we assume a particular view about normative reasons or motivation. In particular, one might argue that it arises (i) only if we deny Internalism about normative reasons, (ii) only if we deny a Humean Theory of Motivation, or (iii) only if we deny that all responses to practical reasons are responses to instrumental reasons. However, I believe none of the claims is true.

The first objection is that if one accepts the doctrine of Internalism about normative reasons, then the question of the *Normative Leap* does not arise.[47] There are many versions of this view, but the general idea is that which normative reasons you have depends in some way on some fact about your motivation. My overall reply is that either version of Internalism is a view about which normative reasons you have, or can have, not how you respond to what you take to be a reason, whether it is a reason or not.

So, one type of Internalist may say that a fact is a reason for you to ϕ only if you have, or in some circumstance would have, a certain motivational state which could play a role in motivation. According to Smith (1995), for example, you have a reason to ϕ in situation S

[47] Williams (1981) famously argues that a fact is a reason for you to ϕ only if ϕ-ing furthers or serves some motive of yours.

only if, if you were fully rational, you would then desire that you ϕ in situation S. But we can accept this and still ask what it is for you to treat a consideration as a normative reason. The question of the *Normative Leap* is not a question of which normative reasons you actually have, but about which considerations you to treat as normative reasons and why. As argued above, one can treat facts as normative reasons which are not reasons.

Another kind of Internalist may argue that whenever you respond to a consideration as a normative reason, you are or can be motivated to ϕ. Of course, when you respond to a consideration as a reason for action by reasoning on the basis of it to the intention to ϕ, you then evidently are motivated to ϕ. Now, either this motivation to ϕ derives from another motivational state you have, or you motivate yourself to ϕ through the reasoning to the intention to do it.[48] But an Internalist might deny the second possibility. You cannot motivate yourself. When you respond to a fact as a normative reason, you must already be motivated to ϕ, or must already have a certain motivational state, say a desire, that can explain your action of ϕ-ing, or your decision to ϕ. For example, to respond to consideration f as a reason to ϕ you must have a desire to ϕ that could motivate you to ϕ. But notice that this is a matter of motivation, not of responding to normative reasons! It is a matter of which consideration you can treat as reason for action, given your motivational states. It is independent of the question of what it is for you to respond to f as a reason, if you do. The objection assumes a particular theory of motivation, rather than one of normative reasons, namely a broadly Humean theory of motivation, according to which motivation has its source in a desire and a means-end belief (e.g. Smith 1994).

[48] My *Approval Account* of reasoning may thus be seen as an account of how you can motivate yourself to do what you take yourself to have some or sufficient reason to do. I here take the thought that you can motivate yourself to do things through practical reasoning from Broome (2013: ch.13), though he defends a radically different idea of what this practical reasoning consists in.

However, even if one accepts this Human theory of motivation, the *Normative Leap* does still arise. Suppose Sara decides to go to the climate protest because she believes her presence will make a significant contribution to the outcome of the protest. The objection now is that Sara must have a certain desire that motivates her to go to the protest, and if she has this, there is no question of the *Normative Lea*p. But even if Sara has a strong desire to go to the protest, she can still treat considerations as normative reasons, and she can do so without treating her desire as a reason. As Pettit and Smith (1990) explain, a desire can have a motivational presence in decision making without appearing in deliberation. The desire may be in the background, so that it explains the agent's choice or action, without the agent realising that he has the desire (p. 271). In those cases, the agent responds to the reasons she takes there to be, not counting her desire as one of them. Of course, an agent can treat a desire as a reason, especially as a prudential reason for action. Sara can treat her desire to go to the protest as an additional reason to go. She then responds to the consideration that she desires to go as a normative reason. There is again a *Normative Leap* to be explained (even if the nature of the normativity involved may then not be morality but prudence).

In sum, even if one believes in a Humean theory of motivation, according to which one can respond to a consideration as a normative reason to do something only if one already has a certain motivational state to do that thing, the question of the *Normative Leap* arises.

There is another version of this objection. One might argue that the presence of a desire in any decision or action entails that the agent regards what he decides to do as good. According to the doctrine of the *Guise of the Good* about desire, to desire x is to believe that x is good, or to represent x as being good or for x to appear good to the agent (Schroeder 2017). Because the desire is in some sense normative in nature, one might think, there is no *Normative Leap* to be explained. But even if whenever you decide to ϕ on the basis of some considerations,

you have a desire to ϕ and because of that regard ϕ-ing *as good* (*Guise of the Good* doctrine), this does not explain why you count as treating the considerations on the basis of which you act as normative reasons. Your desire to ϕ and your seeing ϕ-ing as good may have nothing to do with the consideration you respond to. If this is so, then the desire cannot explain why you respond to these in a way that seems to inherently involve normativity. Another possibility is that you treat your desire to ϕ as a reason to ϕ, which I discussed above.

To this, the defender of the *Guise of the Good* doctrine might still reply as follows: responses to reasons are instrumental in character. Suppose you have the desire to have pizza for lunch, and because of that, decide to go to the café next door that serves pizza. You then acquire the desire to go to the café. You treat the fact that the café next door serves pizza as a reason to go to the café. You treat it as a reason, because you believe it is something that allows you to satisfy a desire. That is, you believe it is something that allows you to do something good. And this just is treating it as a reason. There is no *Normative Leap* left to be explained. But this is so only because it is explained! And the explanation of the *Normative Leap* given here is very much like the one I give. You count as treating a consideration as a reason for an action because you have a belief about what is good, and a belief about why the action in question is good. The details will depend on what version of the doctrine of the *Guise of the Good* we assume here, but I conjecture that the explanations of why you count as treating the consideration as a reason will be similar. Two differences between the approaches, however, are the following. First, the defender of the *Guise of the Good* thesis seems to commit to the claim that if I treat a consideration as a reason to do something, then I see the action as good *because of a desire* of mine. My view is less restrictive than that. It may be that you see the action as good in virtue of something else, say a fundamental conviction about what is good. The second difference is that my view does not postulate that

whenever you see a consideration as a reason to do something, you see the favoured action as allowing you to satisfy a desire. I think this is another advantage of my view. Even if our actions can be explained in terms of beliefs and desires, the considerations I treat as reasons for action are not necessarily seen in the light of my desires. If I see a person has fallen on the street and go over to help, then I do not respond to the fact that the person is in need of help because I see doing so as satisfying a desire of mine.

In sum, the doctrine of the *Guise of the Good* seems to explain of the *Normative Leap* in some cases of responding to reasons, where these explanations seem ultimately to be similar to the ones I offer. But the doctrine seems to be unnecessarily restrictive and cannot plausibly explain all cases. In any case, I don't believe the doctrine of the *Guise of the Good* is plausible, so I do not pursue this issue further here. I believe I can find myself desiring x without in any way representing x as good. And I can also do what is necessary to satisfy this desire, without representing doing that as good. Further, it seems some animals have desires without having the capacity to represent things as good in the genuine normative sense that would be required to explain the *Normative Leap*.

A third objection to the question of the *Normative Leap* may be given by appeal to intentions. Sometimes we do things partly because they allow us to do what we intend to do already. Suppose Tom has the intention to have pizza for lunch. He realised that the café next door serves pizza today and on the basis of this belief, decides to go there. Tom treats the fact that the café serves pizza as a reason to go there. We explain his doing so partly by reference to his intention to have pizza. We may wonder whether this fact is a normative reason for Tom to go to the café and if it is, why. But it seems hard to deny that Tom *treats* the fact as a reason for him to go to the café. Treating considerations as so-called instrumental reasons is still treating them as reasons. So even if all our responses to normative reasons were responses to

instrumental reasons – which I take to be implausible - we would still have to explain the *Normative Leap*.

In conclusion, the question of the *Normative Leap* arises also for those who believe in a Humean theory of motivation, those who advocate the *Guise of the Good* for desires, and even for those who deny that there are any normative reasons for action, as arguably Hume did (Millgram 1995; Hume 2003).

Appendix II

A Sensitivity to Normative Properties

Many have argued or stated that responding to normative reasons essentially requires a sensitivity to normative properties. When you respond to normative reasons you track normative features in the world. Lord (2018: 133) calls this claim the *Sensitivity Platitude*: you do what is demanded of you not just coincidentally, but "as a result of some type of sensitivity to the right, good, or required" (Lord 2018: 127). Blackburn (1998:1) says that "to develop an ethical personality is to become sensitive to different aspects of things and to be disposed to use them to influence or determine attitudes, emotions and choices." He describes the ethical agent as one who takes *inputs*, representations of features of the world around him, and responds with certain *outputs*, attitudes and behaviours, where the function from input to output is an "ethical sensitivity" (p. 5). Dancy and Skorupski appeal to a sensitivity for normativity in the context of reasoning. Dancy (2018: 143) says, for example, that "of course reasoning requires a normative sensitivity. One needs to be able to distinguish what is a reason from what is not, and to respond accordingly." Skorupski (2010:139) argues that practical

reasoning as well as reason-responsiveness "essentially involve[s] a sensitivity to evaluate practical as well as epistemic norms." The list goes on.[49]

A claim about the sensitivity to normative features can at most be metaphorical. It misleads us to believing that we have a special sense with which we perceive normative properties. If responding to reasons was a matter of a sense perception, we could hardly be criticised for failing to use it correctly. We could at most be blamed for not training it correctly, but it would be mysterious how we are expected to improve it substantially and immediately through our rational faculties.

Some suggest the sensitivity can be understood in terms of dispositions. You manifest the sensitivity to normativity if you have a disposition to do what the contents of your beliefs, if they were true, would require you to do. But we can be disposed to act in line with what we have reason to do, given our beliefs, without recognizing the normative import of what we believe. In virtue of what would we count as responding to reasons, rather than complying with what the reasons demand? In virtue of what would we not merely respond to the facts as we perceive them? A dispositional approach would not explain the *Normative Leap*.[50] For example, Mantel (2017) argues that responding to normative reasons (in her terms: acting for a normative reason) requires "the disposition to be motivated to do what is favoured by normative reasons" (568). But for his disposition to do the work it is supposed to do, namely account for the role of the "normative force of the reason" (581) in the agent's action, it must be a disposition to do what is favoured by the normative reason one responds to *because it is*

[49] Raz (2011), Scanlon (1998:57) and Kauppinen (2018) argue that responding to normative reasons is manifesting a sensitivity to normative features in world. Pettit (2018: 280, 323) also adopts this language when he argues that persons are "sensitive to the relevant desiderata" for actions and "sensitive to facts about desirability."

[50] This objection also applies to a similar, but different interpretation of the sensitivity to normative properties, which relies on the epistemological concept of 'safety'. One could argue that a belief that p or intention to ϕ counts as a response to normative reason f if and only if you respond to f with the belief that p or intention to ϕ with modal robustness: in close-by worlds, you respond to f in the same way.

favoured by that reason. That is, it must be a disposition to do what one recognises to be favoured by the reason. And with that, we are back to the beginning, standing before the question of what it is to recognise or see something as a reason.

Another, related problem with the dispositional approach is this. It seems to offer a kind of causal explanation of the *Normative Leap*. A dispositional explanation of a response to reasons seems to say that the response is "caused by its being in the relevant stimulus condition on that occasion" (Wedgwood 2007: 184). If the stimulus condition is simply a non-normative fact about the world, then we could again not distinguish between responses to normative reasons and responses to the reason-constituting facts as such. We could not explain the *Normative Leap*. If, alternatively, the stimulus condition is the instantiation of a normative property, then the argument relies on the idea that normative properties are causally efficacious, which is at the very best controversial (see e.g. Harman 1984). A purely dispositional interpretation of the *Sensitivity Platitude* does not seem sufficient to explain the *Normative Leap*.

The point is that the *Sensitivity Platitude* by itself does not explain what it is to respond to normative reasons. It requires interpretation, that is, an account of what manifesting the sensitivity to normativity consists in. The account I proposed here does not rely on the notion of a sensitivity to normativity. We are not sensitive to normative properties or relations. Instead, responding to reasons is a way of *expressing* our conceptions of normative properties or general normative facts. They determine what we treat as reason and how.

5. Correctness, Goodness, and the Aim of Reasoning

5.1. Evaluations of Reasoning

Our reasoning is subject to evaluation and demands. We may say that

Our reasoning is not always correct.

Our reasoning is not always good.

Our reasoning does not always reach its aim.

Our reasoning is not always how it ought to be.

Each of these statements has intuitive appeal and is more or less part of ordinary language. But neither their meaning, nor their truth is evident. Do they all express the same or do they make different claims? If they make different claims, then there is more than one standard of evaluation for reasoning. We must then ask what it is for reasoning to be correct, what is it for reasoning to be good, and what the aim of reasoning is.

This chapter addresses these questions. **Section 5.2** explains what it is for reasoning to be correct. The correctness of reasoning is a matter of the relation between the premise-attitudes and the conclusion-attitude. The *Approval Account* implies that you reason correctly if and only if you manifest a correct conception of normative reasons and how they make it the case a certain response is permissible. Whether an instance of reasoning is correct depends thus on what we have sufficient reason to believe or do in a given context. Without entering into substantive normative debates, I suggest that this is a plausible result. In **section 5.3,** I discuss deductive and non-deductive theoretical reasoning. **Section 5.4** is dedicated to practical reasoning. I suggest that the *Approval Account* allows us to understand the activity of 'weighing reasons', and how we deal with cases of equipollence and indeterminacy. **Section 5.5** argues that there is no such property as 'being good *qua* reasoning', but that reasoning

may only be good *for* something or good *in some respect*. The attributive goodness of an instance of reasoning depends on how well the reasoning serves the purpose for the sake of which you reason. In **section 5.6** I propose to understand the claim that reasoning has an aim as a metaphor. I discuss the currently popular claim that reasoning aims at answering questions, as well as the classical claim that reasoning aims at the truth and the good.

5.2. Correctness

5.2.1. Attributive Correctness

It is not merely a contingent fact about reasoning, but essential to reasoning that it has a standard of correctness. The activity of reasoning is, in the following sense, a correctness-fixing. What it is to reason sets a standard for reasoning to be correct reasoning (Thomson 2008). The activity 'reasoning' is correctness-fixing analogously to the way in which 'reciting a poem' is correctness-fixing. What it is to recite Shakespeare's Sonnet 18 sets a standard for doing so correctly. Only mental processes that count as reasoning can be measured against this standard of correctness. Only reasoning can be correct or incorrect *qua* reasoning.

Following Geach (1956) and Thomson (2008), we must note that the 'correct', in 'correct reasoning' is an *attributive* adjective, not a predicative one that describes a property. We are not dealing with the property of correctness, but the property of correct reasoning. As Thomson explains,

> "An adjective '*adj*' is a predicative one just in case
>
> A is an *adj* K1
>
> and
>
> A is a K2
>
> entails the proposition
>
> A is an *adj* K2" (Thomson 2008: 4).

An example for a predicative adjective is ‘red’:

‘A is a *red* car’

and

‘A is a Mercedes’

entails

‘A is a *red* Mercedes’.

An example for attributive adjective is ‘big’:

‘A is a big mouse’

and

‘A is an animal’

does not entail

‘A is a big animal’.

Now consider the reasoning “I promised to skip the meeting. So, I shall skip the meeting” and call it *R*. Forming an intention is making a decision to act, so the reasoning expresses a decision-making process. ‘*R* is a correct instance of reasoning’ and ‘*R* is a decision-making process’ does not entail ‘*R* is a correct decision-making process’. Correct reasoning leading to an intention is not necessarily a correct decision-making process, if there is such a thing at all. In the case at hand, the decision-making process might be incorrect because you forgot to take into consideration that the meeting is extremely important. Therefore, ‘being a correct instance of reasoning’ is an attributive adjective, like ‘big’. ‘Correct’ when applied to reasoning is an attributive adjective, and what we are concerned with is the property of being a correct instance of reasoning.

5.2.2. A Content and a Psychological Dimension

One might hold that theoretical reasoning is correct only if it yields a true belief (or knowledge), and practical reasoning is correct only if it concludes in the intention to do what is permissible (or obligatory). This view is implausible, however. You can reason correctly from some false belief to another false belief, for example. Similarly, you can reason correctly to do something that you ought not to do. For example, you can reason from the unjustified and false belief that the building is on fire to the intention to run out as fast as possible. On the other hand, you can reason incorrectly with only true beliefs. Suppose you truly believe q, and if p, then q. You reason "If p, then q. q. So, p." You commit the fallacy of affirming the consequent, and your reasoning is not correct even though it concludes in a true belief (Broome 2015a). This shows that the correctness of reasoning is not a matter of the properties of the conclusion-attitude (that it is true or permissible, for example), but a matter of the relation between the premise- and conclusion-attitudes. But this content-related condition is not sufficient for correct reasoning. There is a psychological condition, too. An instance of reasoning is correct only if the right kind of relation holds between the attitudes constituting the reasoning *because of the reasoner's action*, that is: in virtue of the reasoner's recognition of it.

You do not reason correctly if, by sheer luck, you form an attitude on the basis of others so that it stands in the right kind of relation to those. You must perform the right kind of mental action. On this view, the correctness of reasoning is a *local* assessment: it is a matter of the attitudes constituting the reasoning. Your reasoning may be correct, for example, even if you forget to take crucial considerations into account. We may then criticise you on other grounds, for being ignorant, negligent or closed-minded, for example.

5.2.3. The *Approval Account* of Correct Reasoning

To say that reasoning is a correctness-fixing kind is to say that the nature of reasoning sets the standard of correctness for reasoning.[1] In chapter **3**, I proposed the *Approval Account* of reasoning, according to which reasoning is essentially a transition from some attitudes to another which you approve of. To approve of your transition is to manifest a conception of normative reasons and how they amalgamate to make a certain response permissible. This sets a standard of correctness, namely:

Correct Reasoning: Agent N reasons correctly from attitudes $A_1(a_1)$ … $A_n(a_n)$ to the intention to ϕ or the belief that p if and only if, and if so because,

(i) if the premises of the reasoning as N perceives them were all that was relevant to whether or not N ought to ϕ or believe that p, then they would (at least partly) make it the case that N has sufficient reason to ϕ or believe that p.

(ii) N's approval of the transition manifests a correct conception of normative reasons and how they aggregate, and

(iii) each premise plays a role in (partly) making it the case that N has sufficient reason to ϕ or believe that p.

When you approve of a transition, I argued in chapter **3**, you rationally commit to the belief that, if the premises were all that was relevant to whether or not you ought to ϕ or believe that p, and if things were as you experience them, then they would make it the case that you have sufficient reason to ϕ or believe that p. It would be implausible if correct reasoning rationally

[1] I am here still following Thomson. Similarly, Velleman defines a formal standard of correctness as one that is "solely stated in terms of, or in terms that depend on, the very concept of being the object of that enterprise" (2000:176), like the formal object of a game is winning.

committed you to false beliefs. So, the *Approval Account* implies that if you reason correctly, this normative belief is true. Condition (i) is therefore necessary for correct reasoning. It specifies what kind of relation must hold between the premise-attitudes and the conclusion-attitude for the reasoning to be correct. It is implied by the nature of reasoning and the claim that it would be implausible if correct reasoning committed you to false beliefs.

But this condition is not sufficient for correct reasoning. You may do some reasoning which, as a matter of luck, satisfies condition (i). In approving of your transition, you may manifest a conception of normative reasons which implies (i), while this conception of normative reasons is, in fact, mistaken. It just happens to yield the correct result in this particular instance. If you were asked to explain your reasoning, your explanation would be misguided. In this case, you respond to the premises as normative reasons, but not as the reasons they are. It is a pure coincidence that what you treat as reasons really are reasons. You are not reasoning correctly. For reasoning to be correct, we therefore need condition (ii), which requires that you manifest a correct conception of normative reasons (and their aggregation). If you satisfy condition (ii), you also satisfy condition (i). (It is therefore somewhat superfluous, but I believe it makes things clearer.)

Condition (iii) requires that every premise of an instance of reasoning plays a role in making the response permissible. It ensures that reasoning from irrelevant premises is incorrect.[2]

Note that in *Correct Reasoning*, the psychological condition of correct reasoning is implicit: When you reason and thereby manifest a correct conception of normative reasons (ii), then (i) cannot be a matter of luck. You reason in the way you do partly because the transition allows you to approve of it, that is: because it satisfies condition (i). Your implicit conception of normative reasons guides your reasoning, as I have argued in chapter **3**. If condition (i)

[2] Strictly speaking, also condition (iii) is implied by condition (ii), but no harm is done in stating it explicitly.

was not satisfied according to your implicit conception of normative reasons, you would not have approved of it, and you would not have reasoned in this way. And if your approval was not based on this conception in the right kind of way, you would not count as reasoning.

Condition (i) implies that correct reasoning is non-monotonic. If reasoning from some set of attitudes to another is correct, this does not imply that reasoning from a set that includes all of the former but also additional premises is also correct. Suppose you reason from the belief that you promised to go to lunch to the intention to go, and that this is correct reasoning. This does not imply that the reasoning from the belief that you promised to go to lunch, and the belief that if you do go to lunch, you will leave someone to die, to the intention to go, is also correct. (I here agree with McHugh and Way (2016b).) The same holds for theoretical reasoning. A body of evidence may give you sufficient reason to believe that p, while a larger body of evidence may give you sufficient reason to believe that not-p. I posit that even deductive reasoning is non-monotonic. Reasoning from believing p and if p then q to the belief that q is correct. But this does not imply that reasoning from believing p, if p then q, r and if r then not-q to the belief that q is also correct. The explanation for this derives from the account of correct reasoning: What makes an instance of theoretical reasoning correct is not that the conclusion is entailed by the premises, but that the first-person premises, i.e. the contents of the premise-beliefs if they were true, would make it the case that you have sufficient reason to believe the conclusion.[3]

In sum, the view is that reasoning essentially involves tracking what makes certain beliefs and intentions permissible to hold. You are guided by your conception of normative reasons

[3] Logical validity, it has been argued, is insufficient for assessing practical reasoning (Piller 2001:200; Broome 2002; Audi 2004: 131). I believe it is also insufficient for assessing theoretical reasoning, even deductive reasoning.

and how they aggregate. You reason correctly whenever this conception is correct, that is, whenever what you take to give you sufficient reason to believe or do something indeed gives you sufficient reason to believe or do it – given things are as you believe them to be. The standard of correctness of reasoning is thus external in the sense that it does not only depend on the reasoner's state of mind. It is not just a matter of internal coherence. Correct reasoning is responding correctly to normative reasons.

5.2.4. Is correctness normative?

Is 'being a correct instance of reasoning' a *normative property?* Let's say normative statements are divided into evaluative ones and deontic or directive ones (Ross 1930). A statement about the property of 'being correct reasoning' does not entail a normative directive. You may have no reason to reason to a certain conclusion, even if it was correct. A correct instance of reasoning is, in this respect, like a correct move in chess.

> "It is entirely consistent to say that a certain move would be correct though you ought not make it, and that a certain move would be incorrect though you ought to make it" (Thomson 2008: 90).

That an instance of reasoning is a correct instance of reasoning does by itself not license any normative directives.[4] You can reason correctly and have no reason to reason in this way. Suppose you reason "This is a cockatoo. If it's a cockatoo, it belongs to the family of Cacatuoidea. Individuals from this species are not native here. So, this bird must have escaped from captivity." Though this reasoning is correct, you may have no reason to reason in this way. It may be that you ought not to think about the cockatoo at all, but about the hurricane

[4] Chang (2013:7) holds that 'it is correct' can be seen a placeholder for a normative operator. That it is correct for you to φ might be taken to mean, for example, that you ought to φ or that you have a reason to φ. Here, an instance of reasoning may be correct but not something that is correct for you to do.

that is about to hit the area you are hiking in. You ought to get to a safe place as fast as possible, and not waste time wondering about the origin of birds around you.[5]

Even if the correctness or incorrectness of reasoning does not issue directives (i.e. imply any claims about how you ought to reason), these properties do have normative *evaluative* significance.[6] For when you reason incorrectly, you manifest an incorrect conception of normative reasons. Whether incorrect reasoning is an epistemic, moral, or prudential failure depends on the specific instance of reasoning. It depends on what kinds of normative reasons you track. An instance of incorrect reasoning may be a moral failure, for example when you reason "I promised to go to lunch. But I don't feel like it anymore. So, I won't go." An evaluation of correctness of reasoning has evaluative normative significance because it tells us something about how you track and responds to what the normative reasons require.

5.2.5. Rationality

One might object to this account that it makes no mention of rationality. Is not incorrect reasoning incorrect because it is irrational? To assess this claim we need an account of what it is for an instance of reasoning to be irrational. One view is that rationality consists in responding correctly to normative reasons. If we understand rationality in this way and responding to normative reasons as a process, for example as I propose in chapter **4**, then reasoning is incorrect whenever, and if so because, it is irrational.[7]

---

[5] Cullity (2018) discusses reasons for deliberation.

[6] Not according to Thomson's (2008: 73) idea of correctness properties that are evaluative, because she thinks there is only attributive goodness. She argues that 'being a correct K' is an evaluative property F only if it is a virtue in a K. A property is a virtue in a kind K only if (i) K is a goodness-fixing kind and (ii) a K is as good a K as a K can be [qua K] only if it has F. On my view, reasoning is not goodness-fixing kind.

[7] Lord (2018) defends the view that rationality is responding to reasons and defines the letter as a process (I object to his argument in chapter **4**), but he does not set up his theory as one of reasoning. Kiesewetter's (2017) account of rationality as responding to normative reasons, by contrast, is not fit to govern reasoning. He formulates the requirements of rationality as state requirements. Responding correctly to reasons, for Kiesewetter, is not a process. It is not about forming attitudes on the basis of others.

However, it is highly controversial whether rationality consists in correctly responding to normative reasons. Rationality, many believe, is about coherence of your attitudes. You are rational in so far as your attitudes stand in a certain relation to each other. Rationality, thus understood, makes structural, not substantive demands on what to believe and intend (T. M. Scanlon 2007). For example, rationality requires of you that you do not, at the same time, believe that p and believe that not-p. And, one might say, rationality requires of you that you intend to ϕ whenever you believe you ought to ϕ. These requirements of rationality govern mental *states*, not processes. They require you to hold or not hold certain (combinations) of attitudes, but they say nothing about the process of forming attitudes on the basis of others.[8] They do therefore not govern reasoning.

John Broome has argued that we can derive so-called basing permissions from the state-requirements of rationality. These basing permissions, Broome argues, tell us when it is rational to base one attitude on a set of others. Basing permissions give rise to the rules of reasoning. Correct reasoning is correctly following these rules. Thus, we get a standard of correctness from the state-requirements of rationality, or so Broome argues.[9] I have objected to the view that reasoning consists in following such rules in chapter 1. However, I have not objected to the general view that rationality is a system of coherence requirements. If we understand rationality in this way, incorrect reasoning, on my view, is not necessarily irrational. But it may be. For example, we may say that an instance of reasoning is irrational whenever it brings you in a state of mind in which you violate a requirement of rationality. Through reasoning, you can satisfy the requirements of rationality, but you can also violate

---

[8] Kolodny (2005: 521) argues that the requirements of rationality govern processes, but what he calls process requirement do not describe what it is to rationally reason to one attitude on the basis of others. For objections to his view of process-requirements, see Fink 2011 and Kieswetter.

[9] McHugh and Way point out problems with this approach. The major problem is that Broome does not say how exactly the basing permissions and rules of reasoning are to be derived from the requirements of rationality. Broome (2015a:202) admits that something else is needed but does not say what.

requirements of rationality through reasoning. The process can make you less coherent. An instance of reasoning may also be irrational because you hold certain normative beliefs. For example, it is irrational to reason from believing p to believing q while believing that p is no reason to believe q (see chapter 3). However, it can be perfectly coherent to have an incorrect conception of normative reasons.

5.3. Correct Theoretical Reasoning

According to the account of *Correct Reasoning*, you reason correctly only if you thereby manifest a correct conception of normative reasons and how they aggregate to make a response permissible. You then treat as normative reasons what really are normative reasons or what would be reasons if things were as you take them to be. Therefore, whether an instance of reasoning is correct depends on what normative reasons there are. I now illustrate that this central claim of *Correct Reasoning* is plausible. Without entering substantive debates about what reasons we have, I suggest that intuitively correct instances of reasoning correspond to plausible conceptions of normative reasons. In other words: judgements about the correctness of reasoning correspond to judgements about reasons.

I first consider modus ponens reasoning as an example of deductive reasoning then non-deductive or evidential reasoning, and finally what we may call reasoning about constitution.

5.3.3. Deductive Reasoning

Suppose you reason from the belief that the barometer falls rapidly (p), and the belief that if the barometer falls rapidly, then a storm is approaching (if p then q) to the belief that a storm is approaching (q). In deductive theoretical reasoning from the belief that p and if p then q to

the belief that q, you manifest a conception of normative reasons for belief according to which the fact that p together with the fact that if p then q in this context make it the case that you have sufficient epistemic reason to believe that q.

I wish to suggest that the conception of normative reasons for a belief that you manifest in presumably correct deductive reasoning is plausibly correct, too. To begin with, I suggested in chapter 3 that a set of facts make it the case that you have sufficient epistemic reason to believe that p only if they have certain epistemic properties. It seems plausible that this condition is satisfied if some facts entail that p but not that not-p. Furthermore, the account of correct reasoning does not commit us to a conception of normative reasons for belief according to which you have some (or even sufficient) epistemic reason to believe whatever follows from what you believe. This would be implausible, for we should avoid cluttering our mind with trivialities.

The reason why we are not committed to this claim is because we can hold that some facts are epistemic reasons for one to believe that p only if one has sufficient value-driven reason to have a true belief about p.[10] But, for modus ponens reasoning like the one at hand here to be correct, this condition would have to be satisfied whenever one reasons to the belief that p. I think this is so (not at least because the standard for having sufficient value-driven reason to have a true belief about something seems rather low). We are usually prompted to reason about something because we are curious about it, because we wondered about it in the past, because the subject matter has pragmatic relevance for our actions, or because we believe we ought to make our mind up about the issue (as voters in a democracy or as students in a degree program, for example). In any of these situations, we presumably have sufficient value-driven

[10] See Steglich-Petersen (2011) for a discussion of reasons to have beliefs about p.

reason to have a true belief about p. However, even if you reason entirely aimlessly to the belief that q, you have sufficient reason to have a true belief about q simply in virtue of your forming the belief that q. When you approve of the transition and implicitly judge that the premises make it the case that you have sufficient reason to believe that q, then this belief is partly self-verifying. If the premises entail that q, they are reasons for the conclusion-attitude partly because your beliefs in those facts cause the conclusion-attitude. If you have a belief that q, or are about to form the belief, facts pertaining to the truth of q become reasons for you to believe or not believe that q. We don't bootstrap reasons for belief into existence by thinking about certain things, but 'activate' them by satisfying the condition on which we are subject to them. (Similarly, if I go to France, I have a reason to comply with the French law.)

In sum, I believe a plausible conception of epistemic reasons implies that the fact that p, and the fact that if p then q, given that you reason on their basis, make it the case that you have sufficient reason to believe that q. This is so partly because p and if p then q entail that q. Maybe not all instances of reasoning that are formally deductively valid are correct as instances of reasoning. But neither do all formal entailment relations ground epistemic reasons. The conjecture is that some facts are sufficient epistemic reason for an agent to believe that p whenever the corresponding reasoning is correct.

5.3.4. Non-Deductive Reasoning

Suppose you entertain some considerations which are evidence that p, and on their basis reason to the belief that p. For example, you might reason from the belief that there are muddy footsteps on the kitchen floor to the belief that the housemate is at home.

It seems that whenever the premises give you sufficient epistemic reason to believe that p in virtue of being evidence for p, reasoning on their basis to the belief that p is correct.[11] Again, I cannot give a substantive theory of evidence and reasons grounded in evidential relations here but shall make a few remarks in support of the proposal. I have already suggested in chapter 3 that some facts which are evidence that p give you sufficient evidential reason to believe that p only if they are sufficient evidence that p, but not sufficient evidence that not-p. Evidential reasons for the belief that p are reasons grounded in evidential properties. A set of facts gives you sufficient evidential reason only if the evidence they provide for p surpasses a certain threshold *and* outweighs the evidence they provide for not-p. The first condition implies that having some evidential reason to believe that p, i.e. some evidence that p, and no evidence that not-p does not make it the case that you have sufficient reason to believe that p.[12] Further, the second condition implies that if the evidence that p surpasses a certain threshold, and you have equally strong evidence that not-p, then it is not the case that you have sufficient reason to believe p and sufficient reason to believe that not-p. In this case, you ought to not believe p, and not believe not-p, but suspend judgement about p. If you reasoned to the belief that p or the belief that not-p on the basis of beliefs in this set of conflicting evidence, your reasoning would not be correct. If however, your reasoning leads you to suspend judgment about p, where suspending judgment is a genuine attitude not reducible to the absence of belief, then you are reasoning correctly.

So far, I have only considered theoretical reasoning to beliefs, but the account can be extended to theoretical reasoning that leads to credences or degrees of beliefs. To have a credence in p is to hold a certain level of confidence in p. Suppose I reason: "His fingerprints were at the

[11] Dancy (2018: 67) makes the same point in arguing that evidential reasoning is correct whenever the premises favour believing the conclusion, whereby that which explains the favouring is a probabilifying relation.
[12] I here follow Reisner (2008:21).

crime scene. He has a motive. He is the only suspect. So, he is the murderer", whereby the first sentences express beliefs, and the last a credence. My approving of the transition commits me to believing that the first-person contents of the premise-attitudes make it the case that I have sufficient reason to hold the conclusion-attitude, i.e. certain level of confidence in the proposition that he is the murderer. It would be irrational for me to believe that the premises do not give me sufficient reason to have the level of confidence in the conclusion that I do. Again, it seems plausible that this belief is true whenever the reasoning is correct. Even if my premises, if they were true, would not give me sufficient reason to believe that p, they may give me sufficient reason to have a certain credence that p. The reasoning is correct if the conclusion is not a belief, but a credence, and the degree of the credence corresponds to the level of evidence for p provided. For example, one view is that the fact that p gives me sufficient reason for the credence of degree *d* that q only if the probability of p given e is 0.*d*.

The central point is, again, that our intuitions about having sufficient evidential reason for belief seem to be aligned with our judgments about correct evidential reasoning, so *Correct Reasoning* appears plausible.

5.3.5. Constitutive reasoning

Lastly, my argument also applies to a third kind of theoretical reasoning which is rarely discussed, but which I believe cannot be subsumed under either deductive or evidential reasoning. Suppose you reason as follows: "This chair is crimson. So, this chair is red." This reasoning is correct because if it is true that it is crimson, this fact would give you sufficient reason to believe that it is red. This, in turn, is so because of an underlying relation of constitution: the fact that something is crimson makes it the case that it is red. What guides your reasoning here is an understanding of the concept 'crimson' and the concept 'red', or an

understanding of the relation between the properties that these predicates pick out. We could model such reasoning as enthymematic deductive reasoning, where the implicit premise is 'if the chair is crimson, then the chair is red'. But I don't think one must have this belief in order to reason in this way. In reasoning, we may be guided by our implicit conceptions of what being crimson involves, without having accessible knowledge that could guide our reasoning. Through reasoning, we make our implicit conceptions explicit; we don't need them to be explicit in order to reason. Furthermore, a fundamental difference between what I call constitutive reasoning and deductively valid reasoning is that the former, but not necessarily the latter, tracks an explanatory relation. To see this, consider the reasoning: "The barometer falls rapidly. If the barometer falls rapidly, then a storm is approaching. So, a storm is approaching." The reasoning corresponds to a valid inference and is, partly because of that, correct. The premises do not explain why a storm is approaching, but they do explain why the reasoner has sufficient reason to believe that a storm is approaching. In contrast, when you reason "This chair is crimson. So, it is red", the premises explain not only why you have sufficient reason to believe the conclusion, but they explain why the conclusion is true. Constitutive reasoning is explanatory in a way deductive reasoning need not be. My suggestion is that some facts give us reasons to believe things because they make them true, not because they – together with other facts, entail their truths. If this is so, then there is a corresponding type of reasoning, distinct from deductive reasons. It tracks reasons for belief grounded in constitutive relations, not grounded in entailment relations.

5.4. Correct Practical Reasoning

According to the account of *Correct Reasoning,* practical reasoning on the basis of some beliefs to the intention to ϕ is correct only if the premises make it the case that the agent has

sufficient reason to respond with the intention to ϕ, that is, only if the agent responds correctly to practical reasons.

5.4.1. Simple Cases of Responding to Reasons

In chapter 4 I differentiated between simple and complex cases of responding to normative reasons. In simple cases, the considerations you respond to have the same polarity: they all favour ϕ-ing. In complex cases, some of the considerations favour some course of action (ϕ-ing), other consideration favour another course of action (Ψ-ing). Some simple cases of responding to reasons are instances of reasoning. And in some of those cases, you respond to *pro tanto* reasons for action. A pro toto reason to ϕ is a reason to ϕ that does not make it the case that you ought to ϕ. Suppose you reason: "I am cold. So, I shall go to the kitchen and get a cup of tea", where the first sentence expresses a belief and the second an intention. The reasoning is correct only if, if the fact that you are cold was all that was relevant to whether you ought to get a cup of tea, it would make it the case that you have sufficient reason to get a cup of tea. In chapter 3 I followed Parfit in that you have sufficient reason to ϕ if and only if your reasons to ϕ are not weaker than, or outweighed by, your reasons to act in any of the other possible ways. The sufficiency of practical reasons is thus a comparative matter. Does this imply that to reason correctly to any intention to ϕ, you have to compare the reasons for ϕ-ing with those for any alternative action? No. You can reason correctly to an intention to ϕ, when the premises you operate on are only *pro tanto* reasons to ϕ. This is so because an implicit premise of any instance of practical reasoning is that the premises you operate on include all considerations that are relevant to whether or not you ought to ϕ. On the assumption that the fact that I am cold is all that is relevant to whether or not I ought to get a

cup of tea, I have sufficient reason to get a cup of tea if I am cold. So, for the reasoning to be correct, I don't have to explicitly consider the reasons not to get a cup of tea.

5.4.2. Weighing Reasons

Of course, few situations and few instances of reasoning are that simple. Often some of the things we consider favour one course of action (ϕ-ing), others another (Ψ-ing). To make a decision - to reason to an intention on the basis of beliefs about these facts - we have to *weigh* the reasons for the one against the reasons for the other. To say that assumes that normative reasons have different strengths or *weights*. The weights of normative reasons determine what you ought to do, ought not to do and what is permissible to do. For example, "whether or not you ought to eat cabbage is determined by the weight of reasons for you to eat cabbage against the weight of reasons for you not to eat cabbage (Broome 2008: 108). The weight of reasons for action is a substantive normative issue I shall discuss here. Nevertheless, it seems plausible again that our intuitions about the weight of reasons correlates with our intuitions about the correctness of reasoning. On the assumption that the reasons to eat cabbage are weightier than the reasons not to eat cabbage, reasoning on the basis of the reasons for and against eating cabbage seems correct only if it leads you to the intention to eat cabbage. If this is so, the central claim of the account *Correct Reasoning* is plausible with regards to this type of practical reasoning, too.

But what do we do, when we weight the reasons? The weighing of reasons is the psychological process that corresponds to how the reasons make it the case that you have sufficient reason (or ought) to do something. Few have taken seriously the question of what this psychological process consists in. Cullity (2018:428-9) argues that "summing the weights of reasons cannot be thought of as the addition of numerical values." He comes to this

conclusion upon recognizing that when alternative courses of action are incommensurable, we cannot determine which of two options is more supported by normative reasons: "we are unable even to make an ordinal comparison of the relative strengths of the reasons", let alone put them on a cardinal scale. Cullity concludes that rational deliberation requires the assignment of weights to and the weighing of (apparent) normative reasons, but he leaves us with no account of what we do when we engage in this process.[13]

I suggest that the *Approval Account* of reasoning can help to decipher the metaphor of *weighing reasons*. For the sake of argument, let us suppose a teleological framework in which reasons are good-makers in the sense explained in chapter 4. Now suppose you consider that p, which makes ϕ-ing good in a certain way and to a certain extent. You then consider q, which makes Ψ-ing good in a certain way to a certain extend. On the one hand, it may be that the way in which p makes ϕ-ing good is the same as the way in which q makes Ψ-ing good. Both actions are good, say, because they help the poor. Weighing the reasons given by p and q is then a matter of determining which action scores higher on some non-normative property that is at issue. It is here a matter of determining which action helps the poor more. On the other hand, the way in which p and q make the respective actions good may differ. For example, the fact that flying is a much faster way of getting to Paris is a reason to take a flight. The fact that it is environmentally friendly is a reason to take the train. What is it to weigh the reasons and make a deliberate decision? According to the *Approval Account*, when you reason on the basis of these considerations to an intention, you respond to them as normative reasons, thereby manifesting an implicit conception of what makes a consideration a

[13] Kearns and Star suggest the activity of weighing reasons is the activity of weighing evidence that one ought to ϕ. However, without an account of what it is to weigh the evidence for or against the proposition that one ought to ϕ - where this is thought of as a psychological process - we are no further in our investigation, see also Hawthorne and Magidor (2018).

normative reason, and what it is to have sufficient reason to do something. To respond to the consideration that flying gets you to Paris very fast as a reason to take a plane is to manifest conception of goodness according to which travelling fast is good to some extent. To respond to the fact that it is environmentally friendly as a reason to take the train is to manifest a conception of goodness according to which living an environmentally friendly life is good. The question of which reason is weightier is the question of how these reasons aggregate to make some response permissible. In the teleological framework we are assuming, this is the question of which course of action has the better outcome. Weighing reasons is weighing goods.[14]

When we "weigh" reasons, we thus compare the goodness of the outcomes of the alternative actions. The hesitation that we associate with the process is a state of uncertainty: we are uncertain as to which action would lead to the better outcome. The uncertainty may be about the non-normative, namely when we are uncertain about the relevant features of each possible course of action (for instance, we may be uncertain how long the train will really take). But the uncertainty may also be about the normative, namely when we have no beliefs about the relative ordering of the values at stake. How much does climate protection matter compared to personal comfort and travel time saved? As said in chapter 4, reasoning requires not just having values (beliefs about what is good, which are part of our implicit conceptions of normative reasons) but having a value system which establishes an order among the things that we believe matter. This order may be incomplete. Maybe the system is incomplete because we were never prompted to think about how the values at stake relate. Maybe we never had to make a similar choice, and never observed anyone whose judgment we trust

[14] The weighing of reasons thus understood raises questions of separability. What if we weigh reasons of different sources against each other? Do we first get the moral ought, the rational ought, the prudential ought and then from these, the overall ought? Or do we get straight from the individual reasons of different sources to what we overall ought to do?

make such a choice.[15] If our value system is incomplete, there is nothing to guide our reasoning. We falter. The more fine-grained our guiding conception of goodness is, and the more confident we are about it (justifiably or not), the less we hesitate. Familiar choices appear easier, because we have firm value rankings that have proven to reliably guide us and we have done the work of subsuming the relevant features of the situation before.

Suppose you reason to the intention to take the train, thereby manifesting a conception of goodness according to which taking the train leads to a better outcome than taking the plane. We can say that the reason to take the train is weightier than the reason to take the flight. It is weightier *because* of the relative goodness of the outcomes of the actions. How much weightier it is depends on how much we could increase the relative speed of travel (that which makes taking the plane good) without changing the fact that taking the train is overall better. If only a little increase in relative travel time would tip the balance and make flying better than taking the train, then climate protection matters little compared to time spent travelling and the reason to take the train is only marginally weightier than the reason to take the plane. Whatever the ordering of the option is in one situation, it may be different in another. Because of that, the strength of the reasons may vary across contexts. (Just suppose the life of thousands of people depends on your timely arrival in Paris).

[15] The weighing of reasons can also be guided by our ends. Suppose you intend to get a nice dinner but also intend not to spend much money. Now you try to figure out what is the best restaurant to, given these two ends. Again, you might be uncertain about the facts: which option will get you what kind of food at what price? Even if you know the relevant facts, you may still hesitate. You may be uncertain about the relative weighting of each goal: if the cheapest restaurant is not also the one with the best food, you need to make a judgment about the relative importance of each goal. What you then do when you weigh the reasons for each option is really making your mind up about how much the price and food quality matter to you (in this context). The mental activity you engage in is the same as when goodness of each option is independent of your self-set goals. In both cases, you commit to a ranking of what matters. One difference is that in the case above, what matters how much is independent of your goals. A second difference is that in this case, when you resolve this normative uncertainty and commit to a view about what matters more, this is less a matter of figuring out what matters how much, but a matter of committing to certain goals, defining what matters to you.

I cannot develop this account of the weighing of practical reasons further here. But what I have said indicates that the *Approval Account* allows us to give a plausible theory of how we weigh reason in practical reasoning. The process of the weighing of reasons does not involve assigning numerical values to reasons that represent their "weight." It is not a calculation. We may give up the idea that there is a cardinal ordering of practical reasons according to their 'weights'. According to the *Approval Account*, the weighing of reasons – as any other reasoning process – is guided by our implicit conception of goodness (or rightness), which allows an ordinal ordering of the options we consider.[16]

5.4.3. When there is no guidance

I have so far only considered cases in which the considerations the agent reasons from make it the case that exactly one of the alternative courses of action is permissible. But what if this is not the case? What if we allow for the possibility of equipollence, for example? If the reasons one considers for ϕ-ing and Ψ-ing are on balance, that is, if ϕ-ing and Ψ-ing are equally right then, according to the account of *Correct Reasoning*, one can reason correctly to the intention to ϕ, and one can reason correctly to the intention to Ψ. In which direction one does reason may depend on one's motivations, desires and commitments, for example. The arbitrary character of the choice one makes stems from the arbitrariness of how one chooses to reason. Whether one reasons to the intention to ϕ or to the intention to Ψ, one manifests the same conception of normative reasons.

[16] As already seen in chapter **3**, my proposal is not compatible with the buck-passing account of value. Which reasons we have and what weight they have depends on what is good (or, on a non-teleological framework, on what kind of actions are right). It is not the other way around, which is what I believe the buck-passer would argue. If we take reasons to be primitive, it seems nothing can guide us when we must weigh reasons.

Another way in which correct reasoning may fail to be determinate is when, for two comparable courses of action, neither is one of them better than the other, nor are they equally good. It may be indeterminate which option is better, and thus indeterminate whether reasoning to the intention to ϕ or reasoning to the intention to Ψ is correct. The options, we may say, are incommensurable.[17] Or we may say that the two options are "on a par", as Chang (2002; 2013) proposes. Either way, the *Approval Account* implies that in cases like this, you can choose to reason in one way or the other. Whether you reason to the intention to ϕ or the intention to Ψ may depend on your desires, disposition, interests, past experiences, talents, etc. Other people may reason differently than you do, but none of you is reasoning incorrectly. In cases of incommensurability or parity, "reasonable people may differ about what we have most [or sufficient] reason to do" (Chang 2013:171). However, if you reason in one way rather than the other, say to the intention to ϕ, you thereby manifest a specific implicit conception of normative reasons, and rationally commit to some specific normative beliefs. Unless you would equally approve of the transition to Ψ, and thus manifest the same conception of normative reasons no matter how you reason (in which case we would have equipollence, not indeterminacy), your reasoning lacks the arbitrary character that it has in cases of equipollence. How you reason matters, because it rationally commits you to different normative beliefs about what is good and what is right. Your choice commits you to a certain value system. How you reason thus constraints how you can reason in the future without becoming incoherent. If, in the near future, you reason in a way that is incoherent with the normative beliefs your reasoning to the intention ϕ commits you to, then you are, it seems, irrational. With the *Approval Account*, we can thus make sense of the intuition that your choices in cases of incommensurability or parity have a bearing on future reasoning, while

[17] We may understand this indeterminacy of what is better as vagueness of the "better than"-relation (Broome 1997).

your choices in cases of equipollence do not. Importantly, we can do so without relying on Chang's claim that your reasoning "creates" a reason for you. Neither the approval, nor the intention to act in one way or the other ground a normative reason. We don't have to postulate any special normative powers or will-based normative reasons to make sense of reasoning when there are no reasons to guide us.

This completes my discussion of correct reasoning. The correctness of reasoning is set by what it is to reason. According to the *Approval Account*, reasoning is responding to normative reasons. You reason correctly whenever you respond correctly to normative reasons, that is, whenever you manifest a correct conception of what is good, right to do and epistemically right to believe.

As a final note on the correctness of reasoning, I should mention so-called enkratic reasoning. Put roughly, enkratic reasoning takes you from the belief that you ought to do something to the intention to do it, as in, for example, "I ought to pay the bill today. So, I shall pay the bill today" (Broome 2013: ch. 6). I have not included enkratic reasoning in my discussion of correct reasoning because, according to the *Approval Account*, it is not a form of reasoning at all.[18] When you reason, you are guided by your conception of what makes what kind of response permissible. But the fact that you ought to do something does not make it permissible to do it. That you ought to do something is itself not a reason for doing it. I do not deny that one can choose to do what one believes one ought to do. The control we have over our minds is not restricted to the control we have in virtue of the capacity to reason. The

[18] Others who take this line of reasoning are Southwood (2016a) and Raz (2011:88; 2015). Wallace (2006:150) argues we have the "ability to form volitional motivations", which allows us to motivate ourselves to do what we believe we ought to do.

power to imagine things is another case in point. How exactly we capture the ability to decide to do what we believe we ought to do is another issue I cannot discuss here.

5.5. Goodness

In ordinary language, talk about good reasoning comes even more natural than talk about correct reasoning. I argue that there are two standards of evaluation for reasoning: correctness and goodness, and that the respective properties are not identical. Correct reasoning may be bad reasoning and incorrect reasoning may be good reasoning.

5.5.1. Good for something vs. good qua reasoning

If we say that an instance of reasoning is good, do we use 'good' as an attributive adjective (like 'correct' above)? If yes, we say that reasoning is a goodness-fixing kind like it is a correctness-fixing kind, and that there is the property of 'being good reasoning qua reasoning'. Some kinds are correctness-fixing and goodness-fixing. For example, my reciting of a poem may be correct, but not very good. But reasoning is not such a kind, or so I shall argue now.

In general, goodness-fixing kinds are functional kinds.[19] A functional kind is something that intrinsically has a function. That is, it is something that has an instrumental purpose and having that purpose is part of its nature. For example, a knife intrinsically has the function of cutting. A good knife is one that cuts well because cutting is what knives are means for. If

[19] Thomson thinks that there are goodness-fixing kinds that are not functional kinds. She argues that a tiger can be good qua tiger, and a tomato can be good qua tomato. I agree with McHugh and Way (2016a) that this is odd. A tiger can be strong or healthy. A tomato can have qualities we look for but having qualities we happen to look for does not make it good qua tomato. If not all goodness-fixing kinds are functional kinds, then I assume here that if reasoning is a goodness-fixing kind, it must be a functional kind.

reasoning is a functional kind, then reasoning is a means to an end, and good reasoning is reasoning that serves this end well.

Does reasoning itself have an intrinsic function, so that good reasoning is reasoning that fulfils this function, or reaches this aim? There is no limit to purposes we may reason for. You might want to prove a theorem, find the shortest way to the train station or decide which career to pursue. The more useful the conclusion you reach, the faster the process, etc., the better is your reasoning. How good your reasoning is may depend on processes that are not reasoning, like remembering, the search for relevant considerations, or the weights you attach to the considerations you adduce. However, the better the means to your ends you find and the more cognitively expedient your reasoning is, the better is your reasoning - not *qua* reasoning, but *for figuring something out*. The fact that we employ reasoning as a means to ends does not imply that reasoning is a goodness-fixing kind. The aim is not part of the nature of reasoning, but a contingent fact about it, depending on our changing interests. In general, our contingent interests in a kind K do not fix that there is such a property as being a good K qua K (Thomson 2008: 24). Here is an analogy: I collect pink pebbles. I see a pink pebble and say 'That's a good pebble!'. This does not make 'pebble' a goodness-fixing kind. There is no such property as being a good pebble *qua* pebble, let alone is a good pebble one that is pink. Similar to my interest in pink pebbles, the aim of reasoning sets a standard of evaluation. Reasoning can be better or worse for achieving the aim for the sake of which I reason. But this aim is not intrinsic to the activity of reasoning. Therefore, the goodness property involved here is not the property of being good *qua* reasoning but being good *for* achieving an end.

One might argue that the intrinsic aim of reasoning is to get us what we want, to allow us to pursue the ends for the sake of which we reason, whatever they may be. We always reason as a means to an end, we always reason for a purpose, for example, to figure something out to

learn about something we are curious about or to make a decision about what to next weekend. The purpose for the sake of which we reason may be encoded in an intention to do something, a belief or a desire.[20] It may be subconscious. How good an instance of reasoning is, the argument goes, depends on how well that aim is achieved. The property of being *good qua* reasoning just is the property of being *good for* achieving the end for the sake of which the reasoning is done. However, to say that every instance of reasoning has a function, or an aim is not to say that the activity of reasoning has an intrinsic function or aim. It is trivial to say that an activity has the aim of fulfilling the purpose for which it is done. Suppose everyone runs only to achieve some end. Nobody runs for no reason. This does not entail that running has the intrinsic aim of allowing you to achieve whatever you want to achieve when you run. Furthermore, it is hard to deny the possibility of aimless reasoning. Even though we mostly reason for a specific purpose, it may be that sometimes we reason aimlessly, without asking ourselves a specific question, for example. When I notice that my dad's keys are lying on my kitchen counter, I may just infer that he is in my flat, for no purpose.

I conclude that reasoning is not a goodness-fixing kind. Reasoning can be good *for* achieving some end, but not good *qua* reasoning. In contrast to the local scope of correctness assessments, the goodness of reasoning may depend on attitudes and capacities that constitute the reasoning process. The purpose for which you reason may be represented by some belief, intention, hope, or desire, for example. How good an instance of reasoning is also depends on the choice of the premises, whether you include all and only the relevant ones, which in turn may depend on your ability to recall things from memory.

[20] Alvarez (2010:7) suggests, along these lines, arguably following Aristotle and Anscombe, that the aim of practical reasoning is that "for the sake of which the agent engages in practical reasoning."

5.5.2. Good in some respect

Reasoning can not only be good for achieving a certain aim, but also *in some respect*. There are such properties as being morally good reasoning or being prudentially good reasoning, for example. An instance of correct reasoning, may, for example, be morally bad if it happens to lead to dramatic consequences. Or it may manifest a virtue such as courage or patience. The property of being good in some respect is not the property of being good qua reasoning. We simply assess an activity, your reasoning, according to some standard of evaluation, which may have nothing to do with the nature of reasoning.

5.5.3. Goodness vs. Correctness

And out of two correct instances of reasoning, one can be better in some respect than the other.[21] We should not conflate the two standards of evaluation. McHugh and Way, I argue, make this mistake. They ask "What is good reasoning?" but give an account of correct reasoning. On their account, reasoning is a goodness-fixing kind, according to which good reasoning is reasoning that reaches its intrinsic aim. The aim of reasoning, they argue, is getting fitting attitudes. Fitting attitudes are taken to be, for example, beliefs that are true and intentions that are permissible. The claim about the aim of reasoning neither sets a standard of goodness, nor a standard of correctness. As argued above, good reasoning as well as correct reasoning can lead to false beliefs and intentions to do impermissible things.[22] McHugh and Way (2018:13, 16) recognize this problem and thus argue that a pattern of reasoning is good if and only if, if the premises are fitting, then the conclusion is fitting, too, and that this is so

[21] Here is an analogy: you make a map in order to set your enemy on the wrong track. The incorrect map will be better, i.e. it serves your purpose of misleading your enemy better. Importantly, it will be better for misleading your enemy, not better qua map. A misleading map is a rather bad map as a map.
[22] In response to McHugh and Way, Broome (2015a: 201-202) gives an example that "shows that the point of reasoning is to acquire attitudes that stand in an appropriate relationship to each other, not merely attitudes that are correct."

in virtue of a relationship between the premise-attitudes and the conclusion-attitude. Their account then seems to be one of correct reasoning, rather than good reasoning. First, their account specifies a relation between the premise-attitudes and the conclusion-attitude, which, I argued, a standard of correctness does. Second, if it was a necessary and sufficient condition for reasoning to be good that the conclusion is fitting *in virtue of a relation to the premises, if they were fitting,* then the goodness of reasoning would be a function of a binary condition. Reasoning would then be either good, if the condition is satisfied, or bad, if it is not. We would thus have to give up the plausible idea that the goodness of reasoning is gradable. There would then be no sense in which one instance of reasoning can be better or worse than another, where neither is entirely good or bad. Contrary to what they claim, the usage of the locution 'good reasoning' does not fit ordinary language.

To conclude, an instance of reasoning may be good *for* a purpose, or good in some respect. The goodness of reasoning is not the correctness of reasoning, and there is no property of being good *qua* reasoning because reasoning is not a goodness-fixing kind.

5.6. The Aim of Reasoning

The claim that reasoning has an aim appears in several ethical and epistemological debates. Some hold that the aim of practical reasoning is to choose what is good, while the aim of theoretical reasoning is to believe the truth (e.g. Kolodny 2007; Velleman 2000: 181; McHugh and Way 2018:12). Wedgwood (2002a) argues that the aim of practical reasoning is to figure out what is best, and Williamson (forthcoming) that it is to get one to "act on what one knows." Several authors suggest that the aim of reasoning is to answer questions, say,

whether to φ or whether p is true.[23] Each of these locutions has some intuitive appeal. However, neither their exact meaning, nor their truth is evident. What does it (semantically) mean to say that reasoning has an aim? Does reasoning have an aim, and if so, what is it and why? Without answers to these questions, talk of the aim of reasoning bears the risk of misleading the debates in which it appears. This section answers these questions. I focus on the most prominent versions of the claim that reasoning has an aim, namely

Reasoning aims at answering questions.

Reasoning aims at the truth and the good.

To begin with, the phrase that reasoning has an aim must be understood as a metaphor.[24] Reasoning is a process or, more precisely, an activity. Activities do not literally have aims. An agent can literally aim at a target or aim to achieve something in performing an activity or a series of activities. An archer aims at the target. I aim at having the tent set up before sunset. You might aim to finish reading this chapter fast. We might say that burglary aims at stealing, and that travel aims at getting somewhere. But these are metaphors for saying that burglary is an act performed with the intention of stealing; and that to travel is to do things with the intention of reaching a certain place. So, as metaphors, claims about the aim of reasoning need interpretation.

5.6.1. Answering Questions

Let us first consider the claim that the aim of theoretical reasoning is to settle the question of whether p, for some specific p, and the aim of practical reasoning is to settle the question of

[23] See especiall Hieronymi (2005; 2009; 2013); Southwood (2016b; 2016a) and Shah (2003), Arpaly and Schroeder (2012: 212, 215, 235), also Raz (2015).
[24] I am here inspired by Wedgwood (2002b: 267), who reminds us that "[b]eliefs are not little archers armed with little bows and arrows: they do not literally 'aim' at anything. The claim must be interpreted as a metaphor."

whether to ϕ, for some specific ϕ. On one interpretation, the claim makes a statement about the nature of reasoning. For example, we may take it to mean that practical reasoning is reasoning that leads to the intention to ϕ, which settles the question of whether to ϕ. Theoretical reasoning is reasoning to the belief that p, which settles the question of whether to believe that p. We can, however, define the two types of reasoning this way without appealing to the metaphor of the 'aim' and without appealing to the answering of questions. The latter is especially misleading, for we can come to answer for ourselves the questions of whether to ϕ or whether p without reasoning. Not every thought process leading to the belief that p or intention to ϕ is reasoning. Even if I explicitly ask myself whether to ϕ and then come to intend to ϕ, no reasoning may have been involved.

Southwood (2016b) argues that the aim of practical reasoning is to answer the question of *what do to*. But, understood as a claim about the nature of reasoning, this cannot be right. Suppose I reason "It will be snowing all day. I shall not go running today." I respond to the consideration that it will be snowing, but before I considered it, I did not think about what to do today, or what to do next. Maybe I had already settled that question, I intended to go running. In fact, the reasoning *raises* rather than answers the question of what to do.

Maybe we should interpret the claim that reasoning aims at answering questions as saying that in reasoning, the agent aims at answering a question. But, again, that seems false. Of course, sometimes we ask ourselves questions and reason in order to answer them. But this is a contingent fact about our psychology, and not an essential characteristic of the activity of reasoning itself. Often we reason to the intention to do something without having asked ourselves whether to do it. Here is an example: "I should eat something healthy for dinner. A salad is healthy. So, I shall have a salad for dinner." I did not aim, let alone intend to settle the question of whether to have a salad for dinner. An analogous point can be made for theoretical reasoning, too. In modus ponens reasoning, for example, I do not necessarily

ask myself whether q before starting to think about p and if p then q. In no sense do I aim to settle the question of whether q, though of course, through the reasoning I do settle it. Therefore, having looked at three ways of interpreting the claim that reasoning aims at answering questions as a claim about the nature of reasoning, I conclude that it is either misleading or false.

When we do ask ourselves questions and reason in pursuit of them, our reasoning may be better or worse for answering them. We might say, for example, that the more original, simple, comprehensive or interesting our reasoning is, the better it is for answering a certain question in a certain context. The claim that reasoning aims at answering questions could thus be interpreted as setting a standard of goodness for reasoning. However, few authors seem to have this in mind when they talk about the aim of reasoning as answering questions.

A third interpretation of the claim that reasoning aims at answering questions is as setting a standard of correctness. You reason correctly whenever you answer the question of whether p or whether to ϕ correctly. But what is it to answer the question of whether p correctly? To answer the question of whether p correctly is, it seems, to come to believe that p if and only if p is true. Answering the question of whether to φ correctly is to decide to ϕ if and only if one has sufficient reason to ϕ. But this says nothing about reasoning, let alone correct reasoning. I can reason incorrectly to a correct answer to the question of whether p, for example. As argued in section 2, the correctness of reasoning is a matter of the relation between the attitudes constituting the reasoning, not a matter of independent properties of the conclusion-attitude. This claim that one reasons correctly whenever one answers the questions of whether to ϕ or whether p correctly, however, makes no mention of this relation.

One could respond to this by saying that my reasoning is correct whenever my answer to the question of whether to ϕ or whether p is correct *given* what I believe to be true.[25] This approach, it seems, collapses into the claim that reasoning to the intention to ϕ or belief that p is correct only if the premises of my reasoning give me sufficient reason to ϕ or believe that p. The asking and answering of questions would then play no role in setting the standard of correctness for reasoning.

Therefore, I conclude that the metaphor that reasoning aims at answering questions is misleading on some interpretations and false on others. Aiming to answer a question is neither necessary nor sufficient for reasoning and does not set a standard of correctness.

5.6.2. Truth and the Good

We have seen that reasoning that is good for something or in some respect, as well as reasoning that is correct may lead us to false beliefs or to intentions to do impermissible things. So reasoning does not aim at the truth and the good in the sense that good or correct reasoning leads us to true beliefs and good intentions. But we can accept the claim that reasoning aims at the truth and the good in the following sense. In reasoning, we are guided by our conception of what constitutes having sufficient reason to believe or do something. What we have reason to do is in some way related to what is good, and what we have sufficient reason to believe is related to what is true. The claim about the aim of reasoning is thus a claim about the nature of reasoning and the correctness of reasoning. However, it is not to be taken too literally. What exactly is the relation between goodness and our reasons for action? What exactly is the relation between truth and our reasons for belief? These are substantive questions, which are likely to have quite different answers. Nevertheless, we can accept the

[25] Hieronymi (2009: 209) argues along these lines.

claim that reasoning aims at the truth and the good, according to the *Approval Account*, if it is understood as a claim about the nature and correctness of reasoning. The aim of reasoning does not set an independent standard of evaluation.

5.7. Summary

Reasoning is a correctness fixing kind. The standard of correctness for reasoning is set by what it is to reason. According to the *Approval Account* of reasoning, you reason correctly if and only if you manifest a correct conception of normative reasons and how they aggregate, that is, of what is right to do and epistemically right to believe. You reason correctly only if you respond correctly to normative reasons. You then treat as normative reasons what really are normative reasons or what would be reasons if things were as you take them to be. Therefore, whether an instance of reasoning is correct depends on what normative reasons there are. I illustrated that this central claim of *Correct Reasoning* is plausible for paradigmatic instances of theoretical and practical reasoning. I suggested how the *Approval Account* of reasoning can be developed to deal with cases of incommensurability, and how it can explain the psychological process of 'weighing reasons.' I argued that reasoning is not a goodness-fixing kind, but that reasoning can be good *for* some end or *in some respect*. Independent of the goodness or correctness of reasoning, we may have normative reason to reason in one way rather than another.

Conclusion

Reasoning is a mental activity through which we form, revise or discard attitudes on the basis of others. In contrast to (often unconscious) thought processes that we have no direct control over and that we are merely subject to, reasoning is something we do. I have argued that the most prominent accounts of the nature of reasoning, which define reasoning in terms of rule-following, are unsatisfying. Furthermore, the view that reasoning can be explained in terms of responding to normative reasons, which many philosophers assume but few defend, is surprisingly underdeveloped. To be explanatorily satisfying, the view would require an account of what it is to respond to normative reasons.

I proposed a new account of reasoning. Accordingly, the mark of reasoning is the attitude of approval. An instance of reasoning is a transition from some attitudes to another that the agents intentionally approves of. The attitude of approval that is essential to reasoning is normative in nature. When we reason, we evaluate our transitions of thought according to our implicit conception of the normative reasons we have for belief and action, and how they make it the case that certain beliefs or intentions are permissible for us to adopt. These implicit conceptions consist of beliefs about what is good, what is right to do, and what is epistemically right to believe. In practical reasoning, we are guided by our beliefs about goodness and rightness – what we colloquially call our values. In theoretical reasoning, we are guided by our implicit conception of what kind of beliefs are epistemic good to have. Reasoning is therefore an inherently normative activity. Without the ability to envisage not only what is the case but also what ought to be, we would not be the kind of reasoners that we are. Those who deny the reality of normativity, deny themselves the ability to reason.

The discussion has raised questions about the grounding of normative reasons which needs further investigation. What role do truth and goodness play in that which makes it the

case that some facts are reasons for us to do or believe things? Are the buck-passing account of value and the view of metaethical expressivism compatible with a plausible theory of reasoning? Further, we may ask: can the account of the mental activity of reasoning generalised to an account of *agency*?

Reasoning, I have argued, is one way of responding to normative reasons. But we don't have to reason to respond to normative reasons. We can also do so through habits, that is, internalised patterns of thought and behaviour that we perform without consciously thinking about them, but that are nevertheless under our control in the sense that we would approve of the respective thought processes, if we focused our attention on them. In general, what it is to respond to reasons, or to treat considerations as reasons, is to form beliefs or intentions through a process that manifests one's conception of rightness and is at least under one's virtual agential control. But we do not necessarily have to actively exercise this control. We can respond to reasons through acquired habits. If, however, we do exercise the control we have over our responses, then we engage in reasoning.

Given the prevalence of the notion of "responding to reasons" in philosophical debates, the account calls for application. Is our capacity to respond to normative reasons our capacity to act intentionally? Or is it that which gives the capacity for reflective self-control, a free will?

We reason correctly when we respond correctly to normative reasons, that is, when we form beliefs or intentions and thereby manifest a correct conception of what we have sufficient reason to believe or do. Sometimes we have false beliefs and respond to merely apparent reasons - considerations that would be reasons if they were true. If we reason on the basis of false beliefs about the world, our reasoning may be correct, but not very good for achieving

our ends. Correct reasoning may also be bad in some respect. For example, it may lead to morally bad consequences. Our reasoning is incorrect when it is guided by an incorrect conception of what we have sufficient reason to believe or do given what believe to be the case. We then treat considerations as reasons which are not reasons and would not even be reasons if they were true; we respond to merely supposed reasons. If nihilism is true, and there are no normative reasons for us to do or believe anything, then there is no correct reasoning. Does this conclusion imply that nihilism is self-defeating?

Through our practical reasoning and the actions based on it, we make explicit to ourselves and others our implicitly held, often unarticulated conceptions of normative reasons for action which constitute what we colloquially call our values. The values we have determine what we see as reasons for us to do certain things. To see a fact as a reason to do something is to see it as making a certain action good or right in some respect. It is to apply a conception of goodness or rightness to the state of the world as we perceive it. But our seeing facts as reasons does not make them reasons. There might even be no reasons for us to believe or do anything, and the thought that there is an 'ought' in addition to what is an illusion. But if we deny that there are facts about what is good, what is right to do and what is epistemically right to believe, then we deny ourselves the capacity to reason. If we deny that some things matter, there is nothing to guide our reasoning.

References

Alvarez, Maria. 2009. “How Many Kinds of Reasons?” *Philosophical Explorations* 12 (2): 181–93.

———. 2010. “Reasons for Action and Practical Reasoning.” *Ratio* 23 (4): 355–73.

———. 2018. “Reasons for Action, Acting for Reasons, and Rationality.” *Synthese* 195 (8): 3293–3310.

Anscombe, G. E. M. 1957. *Intention*. 2nd ed. Library of Philosophy and Logic. Oxford: Basil Blackwell.

———. 2005. “Practical Inference.” In *Human Life, Action and Ethics : Essays by G.E.M. Anscombe*, edited by Mary Geach and Luke Gormally. Exeter: Andrews UK Ltd.

Asarnow, Samuel. 2017. “The Reasoning View and Defeasible Practical Reasoning.” *Philosophy and Phenomenological Research* 95 (3): 614–36.

Audi, Robert. 2004. “Reasons, Practical Reason, and Practical Reasoning.” *Ratio* 17 (2): 119–149.

Bader, Ralf. 2015. *Kantian Axiology and the Dualism of Practical Reason*. Edited by Iwao Hirose and Jonas Olson. Oxford University Press.

———. 2016. “Conditions, Modifiers, and Holism.” In *Weighing Reasons*, edited by Errol Lord and Barry Maguire, 27–55. Oxford University Press.

———. 2017. “The Grounding Argument against Non-Reductive Moral Realism.” In *Oxford Studies in Metaethics 12*, edited by Russ Shafer-Landau. Oxford University Press.

Bengson, John. 2018. “The Unity of Understanding.” In *Making Sense of the World: New Essays on the Philosophy of Understanding*, edited by Stephen R. Grimm. N. Oxford University Press.

Bittner, Rüdiger. 2001. *Doing Things for Reasons—The Idea*. Oxford University Press.

Blackburn, Simon. 1998. *Ruling Passions*. Oxford: Oxford University Press.

———. 2010. “The Majesty of Reason.” *Philosophy* 85 (1): 5–27.

Boghossian, Paul. 2014. “What Is Inference?” *Philosophical Studies* 169 (1): 1–18.

———. 2016. “Rationality, Reasoning and Rules: Reflections on Broome’s Rationality through Reasoning.” *Philosophical Studies* 173 (12): 3385–97.

Boghossian, Paul A. 2008. “Epistemic Rules.” *The Journal of Philosophy* 105 (9): 472–500.

Broadie, Alexander. 1968. “The Practical Syllogism.” *Analysis* 29 (1): 26.

Broome, John. 1991. *Weighing Goods: Equality, Uncertainty and Time*. John Wiley & Sons.

———. 2001. “Are Intentions Reasons?” In *Practical Rationality and Preference: Essays for David Gauthier, Edited by Christopher Morris and Arthur Ripstein*, 98–120. Cambridge University Press.

———. 2002. “Practical Reasoning.” *Reason and Nature: Essays in the Theory of Rationality*, 85–111.

———. 2004. “Reasons.” In *Reason and Value: Themes From the Moral Philosophy of Joseph Raz*, edited by R. Jay Wallace, 2004–28. Oxford University Press.

———. 2007. “Does Rationality Consist in Responding Correctly to Reasons?” *Journal of Moral Philosophy* 4 (3): 349–374.

———. 2008. “Reply to Southwood, Kearns and Star, and Cullity.” *Ethics* 119 (1): 96–108.

———. 2013. *Rationality through Reasoning*. John Wiley & Sons.

———. 2014a. “Comments on Boghossian.” *Philosophical Studies* 169 (1): 19–25.

———. 2014b. “Normativity in Reasoning: Normativity in Reasoning.” *Pacific Philosophical Quarterly* 95 (4): 622–33.

———. 2015a. “Replies.” *Teorema* 34 (2): 191–209.

———. 2015b. “Reason versus Ought: Reason Versus Ought.” *Philosophical Issues* 25 (1): 80–97.

———. 2016. “Responses.” *Philosophical Studies* 173 (12): 3431–48.

———. 2019. “A Linking Belief Is Not Essential for Reasoning.” In *Reasoning*, edited by Brandon Balcerak Jackson and Magdalena Balcerak Jackson. Oxford University Press.

———. 2013. “Enkrasia.” *Organon F. Medzinárodný Časopis Pre Analytickú Filozofiu* 4 (20): 425–36.

Broome, John, and Christian Piller. 2001. “Normative Practical Reasoning.” *Aristotelian Society Supplementary Volume* 75 (1).

Brown, Campbell. 2011. “Consequentialize This.” *Ethics* 121 (4): 749–71.

Brunero, John. 2010. “The Scope of Rational Requirements.” *The Philosophical Quarterly* 60 (238): 28–49.

Carroll, Lewis. 1895. “What the Tortoise Said to Achilles.” *Mind* 4 (14): 278–280.

Chan, Timothy, ed. 2013. *The Aim of Belief*. Oxford University Press.

Chang, Ruth. 2002. “The Possibility of Parity.” *Ethics* 112 (4): 659–688.

———. 2013. “Grounding Practical Normativity: Going Hybrid.” *Philosophical Studies* 164 (1): 163–87.

Clifford, William. 1876. “The Ethics of Belief”.

Cullity, Garrett. 2018. “Weighing Reasons.” In *The Oxford Handbook of Reasons and Normativity*, edited by Daniel Star. Vol. 1. Oxford University Press.

Cunningham, J. J. 2019. “Is Believing for a Normative Reason a Composite Condition?” *Synthese* 196 (9): 3889–3910.

Dancy, Jonathan. 2002. *Practical Reality*. Oxford University Press.

———. 2004. *Ethics without Principles*. Oxford University Press.

———. 2008. “On How to Act—Disjunctively.” In *Disjunctivism: Perception, Action, Knowledge*, edited by Adrian Haddock and Fiona Macpherson. Oxford University Press.

———. 2014. “From Thought To Action.” In *Oxford Studies in Metaethics, Volume 9*, edited by Russ Shafer-Landau. Oxford University Press.

———. 2018. *Practical Shape*. Oxford University Press.

Darwall, Stephen. 2003. “Desires, Reasons, and Causes.” *Philosophy and Phenomenological Research* 67 (2): 436–43.

Davidson, Donald. 1963. “Actions, Reasons, and Causes.” *The Journal of Philosophy* 60 (23): 685–700.

———. 2001a. *Essays on Actions and Events: Philosophical Essays*. Vol. 1. Oxford University Press.

———. 2001b. “A Coherence Theory of Truth and Knowledge.” In *Subjective, Intersubjective, Objective*. Oxford University Press.

Evans, Ian. 2013. “The Problem of the Basing Relation.” *Synthese* 190 (14): 2943–57.

Fernandez, Patricio A. 2016. “Practical Reasoning: Where the Action Is.” *Ethics* 126 (4): 869–900.

Fine, Kit. 2012. “Guide to Ground.” In *Metaphysical Grounding: Understanding the Struktur of Reality*, edited by Fabrice Correia and Benjamin Schnieder, 37–80. Cambridge University Press.

Fink, Julian. 2016. “The Ladder of Rationality: John Broome: Rationality Through Reasoning” *Ethical Theory and Moral Practice* 19 (3): 787–91.

Frankfurt, Harry G. 1971. “Freedom of the Will and the Concept of a Person.” *The Journal of Philosophy* 68 (1).

Frege, Gottlob. 1979. *Posthumous Writings*. Oxford: Blackwell.

Friedman, Jane. 2013. “Suspended Judgment.” *Philosophical Studies* 162 (2): 165–81.

Geach, P. T. 1956. “Good and Evil.” *Analysis* 17 (2): 33.

Gibbard, Allan. 2002. “Normative and Recognitional Concepts.” *Philosophy and Phenomenological Research* 64 (1): 151–67.

Gibbons, John. 2010. “Things That Make Things Reasonable.” *Philosophy and Phenomenological Research* 81 (2): 335–361.

———. 2013. *The Norm of Belief*. Oxford University Press.

Ginsborg, Hannah. 2006. “Reasons for Belief.” *Philosophy and Phenomenological Research* 72 (2): 286–318.

Gjelsvik, Olav. 2015. “Rationality, Capacity and Inference.” *Teorema* 34 (2): 105–16.

Green, Mitchell S. 2000. “The Status of Supposition.” *Noûs* 34 (3): 376–99.

Gregory, Alex. 2016. “Normative Reasons as Good Bases.” *Philosophical Studies* 173 (9): 2291–2310.

Grice, H. P. 2001. *Aspects of Reason*. Oxford University Press.

Hacker, P M S. 2007. *Human Nature – The Categorial Framework*. Oxford: Blackwell.

Harman, Gilbert. 1984. “Logic and Reasoning.” *Synthese* 60 (1): 107–27.

———. 1986. *Change in View: Principles of Reasoning*. Cambridge, Mass: MIT Press.

Hawthorne, John, and Ofra Magidor. 2018. *Reflections on the Ideology of Reasons*. Edited by Daniel Star. Vol. 1. Oxford University Press.

Hieronymi, Pamela. 2005. “The Wrong Kind of Reason.” *The Journal of Philosophy* 102 (9): 437–457.

———. 2009. “The Will as Reason.” *Philosophical Perspectives* 23 (1): 201–220.

———. 2013. “The Use of Reasons in Thought (and the Use of Earmarks in Arguments).” *Ethics* 124 (1): 114–127.

Hills, Alison. 2007. “Practical Reason, Value and Action.” *Journal of Moral Philosophy* 4 (3): 375–92.

Hlobil, Ulf. 2014. “Against Boghossian, Wright and Broome on Inference.” *Philosophical Studies* 167 (2): 419–29.

———. 2019. “Inferring by Attaching Force.” *Australasian Journal of Philosophy* 97 (4): 701–14.

Hornsby, Jennifer. 2008. “A Disjunctive Conception of Acting for Reasons.” In *Disjunctivism: Perception, Action, Knowledge*, edited by Adrian Haddock and Fiona Macpherson. Oxford University Press.

———. 2016. “Intending, Knowing How, Infinitives.” *Canadian Journal of Philosophy* 46 (1): 1–17.

Hume, David. 2003. *A Treatise of Human Nature*. Courier Corporation.

———. 1904. *Enquiry Concerning Human Understanding*.

Hyman, John. 1999. “How Knowledge Works.” *The Philosophical Quarterly* 49 (197): 433–451.

———. 2011. “Acting for Reasons: Reply to Dancy.” *Frontiers of Philosophy in China* 6 (3): 358–68.

———. 2015. *Action, Knowledge, and Will*. Oxford University Press.

Jackson, Frank, and Robert Pargetter. 1986. “Oughts, Options, and Actualism.” *The Philosophical Review* 95 (2): 233.

Jackson, Frank, Philip Pettit, and Michael Smith. 2004. *Mind, Morality, and Explanation: Selected Collaborations*. Oxford University Press.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow [Electronic Resource]*. New York: Farrar, Straus and Giroux.

Kant, Immanuel. 2002. *Groundwork for the Metaphysics of Morals [Electronic Resource]*. Rethinking the Western Tradition. New Haven: Yale University Press.

Kauppinen, Antti. 2018. "Practical Reasoning" in *The Oxford Handbook of Reasons and Normativity*, edited by Daniel Star. Vol. 1. Oxford University Press.

Kearns, Stephen, and Daniel Star. 2009. “Reasons as Evidence.” In *Oxford Studies in Metaethics*. Oxford University Press.

———. 2013. “Weighing Reasons.” *Journal of Moral Philosophy* 10 (1): 70–86.

Kelly, Thomas. 2002. “The Rationality of Belief and Some Other Propositional Attitudes.” *Philosophical Studies* 110 (2): 163–196.

———. 2003. “Epistemic Rationality as Instrumental Rationality: A Critique.” *Philosophy and Phenomenological Research* 66 (3): 612–40.

———. 2007. “Evidence and Normativity: Reply to Leite.” *Philosophy and Phenomenological Research* 75 (2): 465–74.

Kiesewetter, Benjamin. 2015. “Instrumental Normativity: In Defense of the Transmission Principle.” *Ethics* 125 (4): 921–946.

———. 2017. *The Normativity of Rationality*. First edition. Oxford, United Kingdom: Oxford University Press.

Kietzmann, Christian. 2018. “Inference and the Taking Condition.” *Ratio* 31 (3): 294–302.

Kim, Jaegwon. 1988. “What Is ‘Naturalized Epistemology?’” *Philosophical Perspectives* 2: 381–405.

Kolodny, N. 2007. “State or Process Requirements?” *Mind* 116 (462): 371–85.

Kolodny, Niko. 2005. “Why Be Rational?” *Mind* 114 (455): 509–63.

———. 2018. “Instrumental Reasons.” In *The Oxford Handbook of Reasons and Normativity*, edited by Daniel Star. Oxford University Press.

Korsgaard, Christine. 1996. *The Sources of Normativity*. Cambridge University Press.

———. 1997. “The Normativity of Instrumental Reason.” In *Ethics and Practial Reason*, edited by Garrett Cullity and Berys Nigel Gaut. Oxford University Press.

———. 2010. “Reflections on the Evolution of Morality.” Amherst Lecture in Philosophy.

Korsgaard, Christine M. 2009. “The Activity of Reason.” In *Proceedings and Addresses of the American Philosophical Association*, 83:23–43.

Leite, Adam. 2007. “Epistemic Instrumentalism and Reasons for Belief: A Reply to Tom Kelly’s ‘Epistemic Rationality as Instrumental Rationality: A Critique.’” *Philosophy and Phenomenological Research* 75 (2): 456–64.

Libet, Benjamin. 1999. “Do We Have Free Will?” *Journal of Consciousness Studies* 6 (8–9): 47–57.

Littlejohn, Clayton. 2012. *Justification and the Truth-Connection*. Cambridge University

———. 2018. “Reasons and Theoretical Rationality.” In *The Oxford Handbook of Reasons and Normativity*, edited by Daniel Star. Oxford: Oxford University Press.

Lord, Errol. 2010. “Having Reasons and the Factoring Account.” *Philosophical Studies* 149 (3): 283–96.

———. 2017. “What You’re Rationally Required to Do and What You Ought to Do (Are the Same Thing!).” *Mind*.

———. 2018. *The Importance of Being Rational*. Oxford University Press.

Mantel, Susanne. 2013. “Acting for Reasons, Apt Action, and Knowledge.” *Synthese* 190 (17): 3865–88.

———. 2017. “Three Cheers for Dispositions: A Dispositional Approach to Acting for a Normative Reason.” *Erkenntnis* 82 (3): 561–82.

Markovits, J. 2010. “Acting for the Right Reasons.” *Philosophical Review* 119 (2): 201–42.

McDowell, John. 1979. “Virtue and Reason.” *The Monist* 62 (3): 331–50.

———. 2013. “Acting in the Light of a Fact.” In *Thinking About Reasons*, edited by David Bakhurst, Brad Hooker, and Margaret Olivia Little, 13–28. Oxford University Press.

McHugh, Conor, and Jonathan Way. 2016a. "What Is Good Reasoning?" *Philosophy and Phenomenological Research* 96 (1): 153–74.

———. 2016b. "Against the Taking Condition: Against the Taking Condition." *Philosophical Issues* 26 (1): 314–31.

———. 2018. "What Is Reasoning?" *Mind* 127 (505): 167–96.

Millgram, Elijah. 1995. "Was Hume a Humean?" *Hume Studies* 21 (1): 75–93.

Parfit, Derek. 2011a. *On What Matters Vol I*. Oxford University Press.

———. 2011b. *On What Matters Vol. II*. Oxford University Press.

Peacocke, Christopher. 1998. "Implicit Conceptions, Understanding and Rationality." *Philosophical Issues* 9: 43.

Pettit, Philip. 1993. "Consequentialism." In *A Companion to Ethics*. John Wiley & Sons.

———. 1996. "The Common Mind: An Essay on Psychology, Society and Politics, Ch.2." *Ethics* 107 (4): 752–54.

———. 2007a. "Joining the Dots." *Common Minds: Themes from the Philosophy of Philip Pettit*, 215–344.

———. 2007b. "Rationality, Reasoning and Group Agency." *Dialectica* 61 (4): 495–519.

———. 2016. "Broome on Reasoning and Rule-Following." *Philosophical Studies* 173 (12): 3373–84.

———. 2018. *The Birth of Ethics*. Oxford University Press.

Pettit, Philip, and Michael Smith. 1990. "Backgrounding Desire." *The Philosophical Review* 99 (4): 565.

Piller, Christian. 2001. "Normative Practical Reasoning: II - Christian Piller." *Proceedings of the Aristotelian Society, Supplementary Volumes* 75: 195–216.

———. 2016. "Evidentialism, Transparency, and Commitments: Evidentialism, Transparency, and Commitments." *Philosophical Issues* 26 (1): 332–50.

Portmore, Douglas W. 2009. "Consequentializing." *Philosophy Compass* 4 (2): 329–47.

Rawls, John. 1971. *A Theory of Justice*. Harvard University Press.

Raz, Joseph, ed. 1978. *Practical Reasoning*. Oxford University Press.

———. 1999. "Explaining Normativity: On Rationality and the Justification of Reason." *Ratio* 12 (4): 354–379.

———. 2002. *Engaging Reason: On the Theory of Value and Action*. Oxford University

———. 2011. *From Normativity to Responsibility*. Oxford University Press.

———. 2015. "Normativity: The Place of Reasoning: Normativity: The Place of Reasoning." *Philosophical Issues* 25 (1): 144–64.

Reisner, Andrew. 2008. “Weighing Pragmatic and Evidential Reasons for Belief.” *Philosophical Studies* 138 (1): 17–27.

Rippon, Simon. 2010. “In Defense of the Wide-Scope Instrumental Principle.” *J. Ethics & Soc. Phil.* 5: i.

Rosen, Gideon. 2010. In *Metaphysical Dependence: Grounding and Reduction*, edited by Bob Hale and Aviv Hoffmann. Oxford University Press.

———. 2017. “Ground by Law: Ground by Law.” *Philosophical Issues* 27 (1): 279–301.

Ross, W. D. 1930. *The Right and the Good*. Oxford: Clarendon Press.

Scanlon, Thomas. 1998. *What We Owe to Each Other*. Cambridge, Mass: Belknap Press of Harvard University Press.

Scanlon, Thomas M. 2007. “Structural Irrationality.” *Common Minds: Themes from the Philosophy of Philip Pettit*, 84–103.

Schroeder, Mark. 2004. “The Scope of Instrumental Reason.” *Philosophical Perspectives* 18 (1): 337–364.

———. 2007. *Slaves of the Passions*. Oxford University Press.

———. 2008. “Having Reasons.” *Philosophical Studies* 139 (1): 57–71.

———. 2016. “Value Theory.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2016. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/fall2016/entries/value-theory/.

Schroeder, Tim. 2017. “Desire.” In *The Stanford Encyclopedia of Philosophy*, Summer 2017. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/sum2017/entries/desire/.

Searle, John R. 2010. *Intentionality*. Oxford University Press.

Setiya, Kieran. 2014. “What Is a Reason to Act?” *Philosophical Studies* 167 (2): 221–35.

Shah, Nishi. 2003. “How Truth Governs Belief.” *The Philosophical Review* 112 (4): 447–482.

———. 2006. “A New Argument for Evidentialism.” *The Philosophical Quarterly* 56 (225): 481–498.

Sidgwick, Henry. 1874. *The Methods of Ethics*. London: Macmillan.

Siegel, Susanna. 2017. *The Rationality of Perception*. Oxford University Press.

Silva, Paul. 2015. “On Doxastic Justification and Properly Basing One’s Beliefs.” *Erkenntnis* 80 (5): 945–55.

Skorupski, John. 2002. “The Ontology of Reasons.” *Topoi 21* 21: 113–124.

———. 2010. *The Domain of Reasons*. Oxford University Press.

Sluga, Hans. 2011. *Wittgenstein*. Wiley-Blackwell.

Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell.

———. 1995. "Internal Reasons." *Philosophy and Phenomenological Research* 55 (1): 109.

Southwood, Nicholas. 2016a. "*'The Thing To Do' Implies 'Can'*: 'The Thing To Do' Implies 'Can.'" *Noûs* 50 (1): 61–72.

———. 2016b. "The Motivation Question." *Philosophical Studies* 173 (12): 3413–30.

Stanovich, Keith E, and Richard West. 2000. "Individual Differences in Reasoning: Implications for the Rationality Debate?" *Behavioral And Brain Sciences* 23: 495–519.

Steglich-Petersen, Asbjørn. 2011. "How to Be a Teleologist About Epistemic Reasons." In *Reasons for Belief*, edited by Asbjorn Steglich-Petersen and Andrew Reisner, 13–33. Cambridge University Press.

———. 2017. "Epistemic Instrumentalism, Permissibility, and Reasons for Belief."

Swain, Marshall. 1981. *Reasons and Knowledge*. Ithaca, N.Y.: Cornell University Press.

Sylvan, Kurt. 2015. "What Apparent Reasons Appear to Be." *Philosophical Studies* 172 (3): 587–606.

———. 2016a. "Epistemic Reasons I: Normativity: Epistemic Reasons I: Normativity." *Philosophy Compass* 11 (7): 364–76. https://doi.org/10.1111/phc3.12327.

———. 2016b. "Epistemic Reasons II: Basing: Epistemic Reasons II: Basing." *Philosophy Compass* 11 (7): 377–89.

Sylvan, Kurt, and Errol Lord. forthcoming. "Prime Time (for the Basing Relation)." In *Well-Founded Belief: New Essays on the Basing Relation*, edited by J. Adam Carter and Patrick Bondy.

Sylvan, Kurt, and Ernest Sosa. 2014. "The Place of Reasons in Epistemology."

Thomson, Judith Jarvis. 2008. *Normativity*. Chicago, Ill: Open Court.

Turri, John. 2009. "The Ontology of Epistemic Reasons." *Noûs* 43 (3): 490–512.

———. 2010. "On the Relationship between Propositional and Doxastic Justification." *Philosophy and Phenomenological Research* 80 (2): 312–326.

———. 2011. "Believing For a Reason." *Erkenntnis* 74 (3): 383–97.

Valaris, Markos. 2017. "What Reasoning Might Be." *Synthese* 194 (6): 2007–24.

Velleman, J. David. 1992. "The Guise of the Good." *Noûs* 26 (1): 3.

———. 1996. "The Possibility of Practical Reason." *Ethics* 106 (4): 694–726.

———. 2000. *The Possibility of Practical Reason*. Oxford: Clarendon Press.

Wallace, R. Jay, ed. 2004. *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford University Press.

———. 2006. *Normativity and the Will: Selected Papers on Moral Psychology and Practical Reason*. Oxford: Oxford University Press.

———. 2015. "The Fugitive Thought." In *Passions and Projections*, edited by Robert N. Johnson and Michael Smith, 246–66. Oxford University Press.

Way, Jonathan. 2013. "Reasons and Rationality."

———. 2016. "Two Arguments for Evidentialism." *The Philosophical Quarterly* 66 (265): 805–818.

———. 2017. "Reasons as Premises of Good Reasoning." *Pacific Philosophical Quarterly* 98 (2): 251–70.

Wedgwood, Ralph. 2002. "The Aim of Belief." *Noûs* 36 (s16): 267–297.

———. 2006. "The Normative Force of Reasoning." *Noûs* 40 (4): 660–686.

———. 2018. "The Unity of Normativity." In *The Oxford Handbook of Reasons and Normativity*, edited by Daniel Star, 23–45. Oxford University Press.

Williams, Bernard. 1981. "Internal and External Reasons." In *Moral Luck*. Cambridge University Press.

Williamson, Timothy. n.d. "Acting on Knowledge."

Wittgenstein, Ludwig. 2009. *Philosophische Untersuchungen: Philosophical Investigations*. Edited by P. M. S. Hacker and Joachim Schulte. Translated by G. E. M. Anscombe. Rev. 4th ed. Chichester, West Sussex, U.K. ; Malden, MA: Wiley-Blackwell.

Worsnip, Alex. 2016. "Moral Reasons, Epistemic Reasons and Rationality." *The Philosophical Quarterly* 66 (263): 341–61.

———. 2018. "What Is (In)Coherence?" In *Oxford Studies in Metaethics 13*, edited by Russ Shafer-Landau. Vol. 1. Oxford University Press.

Wright, Crispin. 2014. "Comment on Paul Boghossian, 'What Is Inference.'" *Philosophical Studies* 169 (1): 27–37.