

This electronic thesis or dissertation has been downloaded from the King's Research Portal at <https://kclpure.kcl.ac.uk/portal/>

The Will as a Capacity for Knowledge

Dietz, Simon

Awarding institution:
King's College London

The copyright of this thesis rests with the author and no quotation from it or information derived from it may be published without proper acknowledgement.

END USER LICENCE AGREEMENT



Unless another licence is stated on the immediately following page this work is licensed

under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International

licence. <https://creativecommons.org/licenses/by-nc-nd/4.0/>

You are free to copy, distribute and transmit the work

Under the following conditions:

- Attribution: You must attribute the work in the manner specified by the author (but not in any way that suggests that they endorse you or your use of the work).
- Non Commercial: You may not use this work for commercial purposes.
- No Derivative Works - You may not alter, transform, or build upon this work.

Any of these conditions can be waived if you receive permission from the author. Your fair dealings and other rights are in no way affected by the above.

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

The Will as a Capacity for Knowledge

Simon Dietz

Master of Philosophical Studies in Philosophy

King's College, London

Abstract

The topic of this thesis is a debate between proponents of two competing interpretations of the claim that our will is a capacity for practical knowledge. On one side of the debate are those who argue that practical reason is a capacity for knowledge of the good, which can be understood as analogous to theoretical reason, our capacity for knowledge of the true (e.g. Tenenbaum, 2009a). The rival conception construes practical knowledge as having non-evaluative content concerning one's actions and their explanation (Setiya, 2012). On this view, the will is a capacity to non-observationally know what one is doing and why by forming intentions.

The thesis consists of three chapters. The first focuses on examining Kieran Setiya's (2010) argument for his claim that the will is a capacity for non-evaluative knowledge of action and against the rival, evaluative conception of the will. This argument aims to show that evaluative representations are superfluous to an account of acting for a reason, which in his view requires no more than beliefs about the appropriate explanation rather than justification of one's action. I will provide a preliminary case for taking his argument against the opposition to be inconclusive as well as questioning a key element in Setiya's non-evaluative account of the will, namely his extremely thin account of practical knowledge, know-how, and their relation.

The second chapter examines the relationship between theoretical and practical reasoning. Discussion of this issue is motivated by the observation that a theoretical analogue of Setiya's argument, designed to establish that we can believe a proposition for a reason without taking that reason to be justifying, does not seem compelling. Hence, Setiya recognises the need to reinforce his arguments by undermining any temptation we might feel to treat theoretical and practical reasoning in a parallel fashion. I will introduce two lines of response to Setiya on this point. According to the first, pursued by Eric Marcus, there are, contrary to Setiya, reasons for thinking that actions relate to reasons in the way beliefs do. However, an alternative to this kind of approach seems to me to be to question Setiya's claim that treating theoretical and practical reasoning as parallel necessarily commits one to thinking of action as the practical analogue of belief.

In the third and final chapter of my thesis, I take up a potential problem for Marcus' view which has been identified by Matthias Haase, namely that the success of his account depends

on its being able to provide an explanation of how knowledge of what we have done (rather than are doing) can be practical. I will suggest that Marcus' reply to that objection is problematic, which provides some initial motivation for considering the alternative line suggested in chapter two.

My tentative conclusion will be that proponents of a view of the will as a capacity for knowledge of the good might benefit from considering an account of the relationship between practical and theoretical reasoning that differs slightly from that which they sometimes favour.

Acknowledgements

I would firstly like to thank Matthew Soteriou. I could not imagine a more generous, supportive, and committed supervisor. It is impossible for me to do justice to all the ways in which he helped me here. He has gone far beyond anything a student could ever reasonably hope for. I also owe a major debt of gratitude to Eliot Michaelson and Mark Textor for helping me out with some difficult situations. Throughout the degree, I have also been supervised by James Stazicker, Thomas Pink, Maria Alvarez, and John Callanan, and I would like to thank them for what they have taught me. Finally, Willa Saadat, Fintan Mallory, Calum McNamara, and Colette Olive have taken the time to read some of this work, saving me from committing at least some offenses against the English language.

Table of Contents

Introduction	5
Chapter 1: Practical Knowledge and ‘Sympathy for the Devil’	8
1.1. Setiya’s Argument	8
1.2. Intention-Cognitivism and Know-how	12
1.3. Conclusion	18
Chapter 2: How to Compare Theoretical and Practical Reasoning	20
2.1. Practical-theoretical parallels	21
2.2. Dynamic inference and the ‘taking’-condition	27
2.3. Anscombe on practical inference	34
2.4. Buridan’s Ass	41
2.5. A concessive response	46
2.6. Comparing the proposals	49
2.7. Tenenbaum’s argument for Aristotelianism	52
2.8. Revisiting the Guise of the Good	54
2.9. Conclusion	59
Chapter 3: The Reach of Practical Knowledge	61
3.1. From Normativism to Presentism	61
3.2. Normativism without Presentism?	63
3.3. A problem for Presentism	68
3.4. Practical inference revisited	69
3.5. Conclusion	71
Conclusion	72
Bibliography	73

Introduction

In his ‘Knowing How’, Kieran Setiya argues for what he calls an ‘Anscombean conception’ of the will as a capacity for practical knowledge (2012, p.285). Contemporary authors usually take ‘practical knowledge’ to refer to the distinctively non-observational way in which agents may normally know their own intentional actions. However, as Sergio Tenenbaum (2009a, p.95) “boldly asserts”, there is historical precedent for taking practical knowledge to be “knowledge of the good in the same way that theoretical knowledge is knowledge of the true” (*ibid*, p.96). According to Tenenbaum, this view underwrites what he calls the ‘Scholastic View’, i.e., the thesis that intentional action essentially involves judging what one does to be good. On the resulting picture, just as successful theoretical reasoning concludes in accurate representations of what is the case, successful practical reasoning concludes in acting (or at least intending to act) well. Thus, the agent’s intention or action constitutes their “final stance” on what is good to do, just as their beliefs constitutes their stance on what is true (*ibid*).

This thesis engages with a debate about the proper way of understanding the Anscombean conception of the will which arises from these two different ways of approaching practical knowledge. On the one side of the debate, there is Setiya who argues that once we get clear on the way practical knowledge in the contemporary sense (henceforth just: practical knowledge) figures in agency, we will see that alternative accounts like Tenenbaum’s Scholastic View are mistaken. Setiya’s thought is roughly that since practical knowledge on its own is enough for intentional action, no further evaluative element can be essential. On the other side, there are those who argue that we cannot properly understand our distinctively non-observational agential knowledge without giving a proper place to the Good in our theory of agency.

In what follows, I examine Setiya’s argument for his non-evaluative version of the Anscombean conception of the will with a view to contrasting and exploring two ways in which friends of views like Tenenbaum’s might respond. As I explain, Setiya’s argument relies on the claim that there are deep disanalogies between theoretical and practical reasoning, and both responses I examine dispute that contention. The difference between the responses lies in how theorising about reasoning in general is approached, and consequently, where parallels between the theoretical and practical are sought.

According to the first approach, theorising about reasoning examines premise-conclusion relations between theoretical and practical attitudes (and possibly actions). Proponents of this approach are usually less interested in describing the way some psychological process unfolds. Within this framework, the response to Setiya hangs on establishing that actions or at least intentions can be conclusions of reasoning thus understood. As explained above, these practical conclusions will be stances on what is good, and practical knowledge will be awareness of taking such a stance, which normally results from self-consciously making up one's mind on what is good to do.

The second view, by contrast, embraces a picture of reasoning as a kind of process. Correspondingly, it seeks parallels between occurrences that happen, as it were, 'on the way' towards theoretical and practical conclusions. I will suggest that there may be a plausible way of developing this second approach in such a way that the practical analogue of a theoretical conclusion is not taken to be the act itself, but the state in which action terminates. Consequently, practical knowledge extends beyond what we are intending or doing, towards (at least some of) what we actually make happen (this idea is taken from Matthias Haase). Moreover, the second view also gets evaluative knowledge into the picture in a different way. Taking inspiration from Gilbert Ryle, I characterise the processes of both theoretical and practical reasoning as depending on our exercising intellectual or practical know-how. This know-how may be seen as constituting knowledge of good rules of reasoning and exercising it may thus put one in a position to cite a justification for the upshot of this exercise.

While I argue that the second account of practical reasoning and practical knowledge has certain advantages over the first, I do not aim to make anything like a decisive case for it. Furthermore, I do not provide any direct arguments for favouring one of these proposals to a completely non-evaluative account of intentional agency. That said, I suggest that the specific proposal Setiya develops has significant drawbacks which my preferred account might help overcome. Thus, I hope to provide motivation for my proposal by illustrating its attractions.

One note: I follow most authors in speaking of the 'Guise of the Good' (GG) rather than the 'Scholastic View' and will use the term to cover a broad range of views, not just Tenenbaum's. For my purposes, all that matters is distinguishing views according to which intentional action essentially involves taking *some* kind of normative or evaluative stance, from views like Setiya's, which deny that any such element is essential. I will not be concerned with disputes about whether we should adopt a relatively weak reading of the

Guise of the Good, according to which action involves only taking what one does to be good in some respect, or a stronger reading like Tenenbaum's, that requires 'all-out' judgments of goodness (). Moreover, I do not believe it will matter whether one favours speaking of a 'Guise of Normative Reasons' rather than a 'Guise of the Good' (cf. Gregory, 2013).

Chapter 1: Practical Knowledge and ‘Sympathy for the Devil’

The chapter critically examines the more recent version of Kieran Setiya’s argument against the Guise of the Good. 1.1. briefly outlines the argument while noting some apparent limitations. Thereafter, 1.2. aims to provide a detailed discussion of the account of practical knowledge that underwrites Setiya’s reasoning, suggesting that it has significant defects. Finally, 1.3. concludes with a brief sketch of how the version of the Guise of the Good developed in the rest of the thesis might help in overcoming the difficulties Setiya appears to face.

1.1. Setiya’s Argument

In ‘Sympathy for the Devil’, Setiya targets a version of the Guise of the Good according to which acting for a reason requires believing that reason to be a good reason, i.e., to justify what one does. While he claims that his reasoning can be extended to weaker versions of his target thesis (e.g., a requirement that one represents one’s reason as good in a kind of perceptual state rather than belief), I do not have the space to consider these alternatives here and focus on the argument about belief.

Setiya’s starting point is Anscombe’s thesis that intentional actions are those to which the question ‘why?’ is given application; where ‘why?’ is read as having the ‘special sense’ referring to an agent’s reason (2010, p.88). Now, Setiya’s first premise is:

- (1) It is sufficient to answer the question “Why?” that one has a belief of the form, “I am doing ϕ because p ,” in the sense of “because” that gives an agent’s reason (*ibid*, p.90)

The thought is that the content of such a belief would not appear to entail that p is actually a good reason for action. After all, people often act on considerations that do not *in fact* show

that there is anything good about what they are doing, even if it could somehow be established that acting thus somehow requires taking them to be good reasons anyway. Setiya consequently adds a second premise:

(2) That I am doing ϕ because p , in this sense, is consistent with the fact that p not being a reason for me to ϕ (*ibid*)

Finally, we get from (1) and (2) to the falsity of GG by adding

(3) If one proposition is consistent with the negation of another, it is possible to believe the first without believing the second (*ibid*)

Hence, given (1)-(3) it appears that we should be able to act for reasons without believing them to be good reasons.

The argument seems to have limitations. As Joseph Raz (2010, p.127) emphasises, we still need a way of distinguishing between the kinds of explanations that show one to be acting intentionally (e.g., ‘I did it to inherit’) and those that do not (e.g., ‘I did it because I was angry’). Only the former kind will make (1) true. By contrast, the belief that one’s anger caused one to lose self-control and thus made one do it obviously would not involve the right sense of ‘because’. But then it seems that until we have unpacked the relevant sense of ‘because’ we cannot be completely sure about what is or is not entailed by the propositions mentioned in (1).

Of course, part of the point of the Guise of the Good is to give a theory of the ‘because’ distinctive of intentional explanation. If such a theory turned out to be attractive, that might give us reasons for denying (2) or at least some variant of it. Setiya is sensitive to this concern. In a footnote, he quotes Raz (2002, p.23) as claiming that “intentional action is action done for a reason; and [...] reasons are facts in virtue of which those actions are good in some respect and to some degree” while noting that this would apparently amount to

denying (2) (Setiya, 2010, p.88, note 18). In response, he complains that such a claim would be “surely too strong”, while citing examples to provide additional intuitive support for (2).

These cases involve acting in ways that are known to be valuable only under certain conditions, which are conspicuously violated in the relevant context. The first example is saving one’s own life at the expense of one’s wife and children, the second is the pursuit of wicked pleasure (*ibid*). In each case, we are supposedly inclined to affirm that the agent’s reason was not in fact good (which would support (2)), and it even seems plausible that the agent himself might reach that verdict despite what he does (which would support (3)).

However, these considerations do not strike me as decisive. To begin with, it does not seem impossible for the GG theorist to simply bite the bullet and insist that saving your life or obtaining some pleasure are good reasons for action, only ones that are massively outweighed in the given context. After all, we do often seem to say there was ‘no reason’ for doing something to mean only that there was an extremely weak reason (cf. Schroeder, 2007, pp.92-7). Likewise, it might be acceptable to say that there is some value in saving your life or experiencing some bad pleasure, even though this value is negligible in the relevant context.

Alternatively, it would be possible to weaken the GG-theorist’s claim further, while still preserving the core idea. One might retreat from the proposal that reasons are actually good-making to the claim that they need only be of a kind that is normally good-making. Pleasant actions, for instance, might be good in general, though particular, defective instances of this kind are not (Clark, 2010, p.236). The point of this move is to illustrate the possibility of agreeing with Setiya’s intuitions about the cases, while holding onto the idea that being moved as the agents of the examples are, is to be moved by something that is, as it were, ‘close enough’ and in salient respects like a good reason. If that were so, then it might still be plausible that the best way of making sense of the ‘because’ of intentional explanation is in terms of the agent’s exercising a capacity for recognising and responding to value. Imperfect exercises of such a capacity may result in our responding to properties that merely appear to be values (which might, for example, be understood as defective instances of good kinds, as Clark suggests).

Setiya recognises the possibility of this type of view. Indeed, he proposes that his argument might be read ‘ecumenically’ as suggesting that a successful GG-theory would be of this sort (2010, p.93). As he expresses the idea, GG-theorists should argue that to accept the kind of explanation figuring in (1) is to accept a “weak proposition about the justification of action” which he unpacks as indicating that the agent is somehow “approximating rationality” in acting for her reason (2010, p.94). He even acknowledges that this is a plausible thing to say in some cases, arguably including his earlier examples:

Some defects of character are recognizable distortions of virtue, and their reasons mimic those of a decent person [...] Here, despite its obscurity, the claim of approximation gets some grip. In giving such reasons, one shows oneself to be in touch with the sorts of considerations that do provide reasons, if not in just this case. One’s motivation can be seen as the flawed or imperfect exercise of a capacity to get things right. (*ibid*, p.99)

However, he contends that there are more extreme examples of bad action, where this sort of account is no longer plausible. He mentions actions done from “bitterness and spite”, “pessimism and despair” and “bigotry and prejudice” (*ibid*). For illustration, he asks us to imagine a community of xenophobes. While it may be that some members of such a community come up with justifications for their xenophobic behaviour (they have a story about the value of ethnic purity etc.), most of its members may be entirely unreflective. They would deny help to a person purely because she ‘one of them’, without having any further thoughts on where the value in this behaviour might lie. What is more, they might explicitly deny that such questions would be relevant:

“Who cares whether it is right or wrong?” they ask. “This is what we do.” (*ibid*)

I am not sure how much weight cases like this one can pull in an argument against the Guise of the Good. When people speak as in Setiya’s quotation above, they can plausibly be taken to deny caring about goodness in some much narrower sense than that figuring in GG (such as specifically moral value). A workable version of GG will need to accommodate conceptions of the good in which, say, ethnic purity takes precedence over morality (to account for the motivations of Nazis). When read in this encompassing fashion, there seems

to be less of an issue in maintaining that racists take what they are doing to be good, even though they would not put things that way (cf. Tenenbaum, 2009b, p.417). Moreover, as Raz argues, there are some positive reasons for taking our racists to think of their xenophobic reasons evaluatively, for they are presumably disposed to try to make them cohere with other considerations they do recognise as values (2010, p.131-2). They will reflect on how much weight one should give to someone's ethnicity when discriminating against them would seriously undermine one's own interests etc.

Finally, Setiya overlooks one additional resource. Particularly unreflective agents might be connected to values or perversions of them in virtue of their deference to others (phrases like 'this is just how we do things' appear to suggest that much). The unreflective racists might not quite comprehend what should be disloyal about helping foreigners, but they might simply defer to those who do have a suitably warped conception of loyalty to genuinely see matters that way. Our capacity to be attuned, at least to some extent, to values we do not quite grasp by 'going along' with others seems to play an important role for moral development and education. It would not be surprising if defective instances of the same phenomenon also played a role in bringing about common patterns of vicious behaviour.

1.2. Intention-Cognitivism and Know-how

The previous remarks were mainly defensive. In what follows, I raise some doubts about Setiya's account of practical knowledge, which underwrites the argument against the Guise of the Good.

As Setiya acknowledges, those beliefs about what we are doing and why that are relevant to his argument are epistemologically distinctive (2016, p.258). As Anscombe (1963, p.51) argues, we seem to know what we are intentionally doing and why non-observationally. We do not need to gather evidence about our motives to know what we are up to. Of course, we might form these beliefs in non-standard ways. By deferring to an analyst, one might come to believe that one is only doing a PhD to upset one's parents. But such beliefs are not plausibly sufficient for us to be acting for the relevant reasons (Setiya would deny that they stand in the

‘right relation’ to one’s actions and thus do not provide successful reason-explanations) (cf. Marcus, 2013, p.521).

At this stage, one might propose that there is a place for the Guise of the Good in an explanation of the non-observational character of our beliefs about the explanations of our actions. A comparison with our knowledge of what we believe and why might encourage one to take this hypothesis seriously. Again, we can form beliefs about our beliefs and our reasons for them through self-observation or deference to the analyst. But under normal circumstances, we seem to be able to answer the question whether we believe that p by answering the question whether we should believe that p, i.e., by weighing up the relevant evidence and making up our minds (see e.g., Moran, 2001, pp.60-5; Boyle, 2011, pp.7-8). Moreover, when arrived at in this latter way, knowledge of our beliefs appears to share the non-observational character of our knowledge of our intentional actions. Thus, we might hypothesise that our non-observational knowledge of action is similarly linked to our capacity to make up our minds what to do.

It is not my aim here to argue for this hypothesis. Rather, the point is just to illustrate that the success of Setiya’s argument partially depends on his being able to offer a convincing account of the distinctive epistemology of the beliefs about the explanation of action that figure in the premises of his argument. Setiya recognises this explanatory burden, but takes himself to have a satisfactory, non-normative account of practical knowledge. His proposal is that acting intentionally requires the presence of an intention, and that intentions are themselves at least partly constituted by beliefs about what one is doing (call this: Intention-Cognitivism). Thus, forming an intention or taking some consideration as one’s reason just is forming the belief that would, if all goes well, constitute one’s knowledge of what one is doing and why (Setiya, 2016, p.259). In the rest of this chapter, I critically examine this view. If it turned out to be unsatisfactory, that would significantly weaken Setiya’s argument.

Any account of intention as constitutively involving belief will need to explain how it can be epistemically permissible for us to form the relevant belief. Ordinarily, formation of a belief is usually licensed by the possession of sufficient evidence. However, if it is genuinely up to me what I will do, the evidence available when I form an intention will not settle this

question. Thus, we need to explain what, if not evidential support, rationalises the belief that (perhaps only partly) constitutes my intention (Setiya, 2008, p.397).

One solution to the problem is to suggest that it is enough for my belief to be epistemically justified that I should have sufficient evidence for it *once it is formed*. This is David Velleman's (1989, pp.56-7) approach to the problem. His idea is that since we are generally inclined to do what we intend, our awareness of intending to do something is sufficient evidence for the belief that constitutes that intention. This supposedly also makes it acceptable to form the intention to begin with, even though at the time of formation one does not yet have the evidence needed to support it.

Setiya finds Velleman's solution unacceptable. He thinks it intuitively implausible that being justified in believing that a given belief would be self-fulfilling once formed (so that having the belief would give one evidence of its truth) is sufficient to make the formation of that belief epistemically permissible. To bring this out, he asks us to consider a character who has good grounds for thinking that a benevolent spirit will make it true that whichever horse he believes will win the race will do so. Setiya argues that it would be implausible to suggest that our character is epistemically permitted to form the belief that some particular horse will win solely based on his belief about the spirit. He concludes that intention-cognitivists need to find some "epistemically relevant difference" between forming an intention and the horse-racing case (call this the 'wishful-thinking-problem') (Setiya, 2008, p.401).

The wishful-thinking-problem is, Setiya suspects, related to a second challenge. Intuitively, it is a necessary condition on acting intentionally that the agent should know how to do what she intends.¹ If knowing how to A and believing oneself to be A-ing are both necessary conditions on intentionally A-ing, we are faced with the problem of showing how one of these necessary conditions is explicable in terms of the other, lest we end up postulating a brute necessary connection between the know-how- and belief-conditions (*ibid*, p.404).

Velleman is arguably unable to explain the necessity of know-how. On his account, it suffices for having knowledge in intention that I will A or am A-ing, respectively, that I am aware of my intention and have knowledge of "ability in the simple conditional sense", i.e. I know that I will succeed if I form the intention (Setiya, 2008, p.404-5). However, this knowledge of ability is unnecessary for intentional action itself. I can irrationally form the intention to do

¹ There are complications here which are irrelevant to my purposes (cf. Setiya, 2008, p.404)

something I do not know I know how to do (Setiya gives the example of clenching one's fist while thinking one is paralysed). Thus, the necessity of knowledge of ability for practical knowledge does not help explain the necessity of know-how for intentional action. Moreover, knowledge of intention is insufficient for know-how, since one can form these kinds of irrational intentions (*ibid*). Velleman's best bet is, according to Setiya, to identify knowledge how to A with the capacity to A (in the sense of 'capacity' on which A-ing intentionally entails the capacity to A). However, and this is crucial to Setiya's thought, Velleman is not entitled to regard this capacity as a form of knowledge. He writes:

On his account, the capacity to Φ intentionally is epistemically inert: all the work is done by knowledge of intention and ability (*ibid*)

Now, Setiya proposes that we can solve both the wishful-thinking-problem and explain the necessary connection at hand if we suppose that know-how licenses our formation of the belief that we will do what we intend. Without it we would not be epistemically permitted to form intentions.

Sarah Paul (2009) is not convinced that Setiya's proposal is much of an improvement over Velleman's. She argues that on this picture, forming an intention is still a matter of illicitly jumping to a conclusion. The trouble is that even if I know how to A before I form the intention-belief to A, this belief does not seem more strongly supported than the belief that I will do something else that I know how to do (*ibid*, p.556).

However, I take it that Setiya did not set out to explain how it can be that the intentions we do in fact form are antecedently more strongly supported epistemically than other intentions we might form. That would appear to be in tension with the thought that it is genuinely up to us what we will do (cf. Setiya, 2008, p.396). All that he aims to do is to identify some relevant difference in the epistemic standing we have with respect to intention-beliefs that it would be permissible to form, and beliefs like that of the character in the horse-racing story.

Nevertheless, it seems like the success of Setiya's case depends on his being able to flesh out his account of know-how in such a way as to make plausible that we should indeed consider it epistemically significant. This task is closely bound up with that of explaining why know-how is more than knowledge in name only.

On the account that Setiya recommends, know-how is a disposition to reliably execute an intention. While I cannot rehearse the details of his argument here (see Setiya, 2012), it is worth pointing out that he is driven to such a thin conception of know-how by the need to account for its necessary presence in intentional action (this allegedly rules out intellectualist conceptions that identify know-how with propositional knowledge, such as Stanley and Williamson, 2001). As Setiya himself acknowledges, it may not be immediately obvious why a mere disposition to reliably execute an intention should be thought of as a genuine form of knowledge.

In response, Setiya emphasises that this disposition ensures that our intention-beliefs will be non-accidentally true. Non-accidental truth is arguably a necessary condition of knowledge, so know-how at least explains why our intentions, when formed correctly, satisfy this necessary condition. It is worth emphasising that non-accidental truth is plausibly not sufficient for knowledge and Setiya makes it clear that he does not want his remarks to be understood as expressing a commitment to such a view. All he aims to show is that know-how does “epistemic work” (2012, p.304), and to that end, he thinks it sufficient to illustrate how know-how enables our intentions to meet at least one necessary condition of knowledge.

However, it is not obvious that this is enough. That some state explains why some belief meets a necessary condition on knowledge does not entail that that state is itself an instance of knowledge. In one passage, Setiya takes up this concern. He writes:

The explanation is not that every enabling condition of knowledge is itself an instance of knowledge, but that, on the Anscombean view, the disposition involved in knowing how to Φ is a capacity to know. (*ibid*)

The proposal is supposed to be that knowledge how to do something is a determination of the general capacity for knowledge that is the will (Setiya, 2012, p.204). If I understand it correctly, the argument for this view proceeds, very roughly, in two steps. First, it is established that Setiya’s dispositional account of know-how is required to account for its necessary involvement in intentional action. Then, he argues that the best explanation of how such a disposition could be a form of knowledge is the Anscombean conception of the will, paired with the claim that know-how “[...] is a determination of the capacity for practical knowledge” (p.304).

There are at least two misgivings one might have at this stage. First, Setiya set out to explain why know-how is a form of knowledge, but he ends up arguing that it is a capacity to know

something. These are obviously two different claims. However, this might not be too much of a problem in the context of the wishful-thinking objection. There the issue was identifying an epistemically relevant difference between the horse-racing case and forming an intention. If know-how is a capacity to know something, we could say that the difference is that forming an intention involves exercising a capacity for knowledge, whereas forming the belief that some horse will win without evidence does not manifest such a capacity. Indeed, it is worth noting that this story actually seems significantly more attractive than simply saying that when I correctly form an intention, it will be true of me that I have some knowledge that the wishful thinker lacks (i.e., knowledge how to execute my intention). After all, the wishful thinker may also know that the benevolent spirit will make his belief true. It looks like we each have some knowledge that is relevant to the truth of the belief we will form, so why would my knowledge make me epistemically better off than the wishful thinker? If the explanation is that I do, whereas he does not, exercise a capacity for knowledge in forming my belief, the significance of our difference is much clearer.

Nevertheless, the need to explain how know-how gets to count as knowledge is plausibly a desideratum for any adequate conception of know-how. Thus, it still poses a challenge to Setiya's overall picture that he might not have properly addressed.

The second misgiving is that it is not clear that the argument for thinking of know-how as an epistemic capacity, a determination of the will understood as the capacity for knowledge of action, works. Note that one exercises the will by forming intentions, which, according to Setiya, constitute one's knowledge of action. Consequently, if knowledge of how to A were a determination of that capacity, it would be natural to assume that to exercise it would be to acquire an intention to A. By analogy: colour vision is a determination of vision. Exercising either capacity will result in my acquiring knowledge of visible properties, it is just that the more specific capacity yields knowledge of a narrower range of properties. This intuitive picture seems to be operative in 'Practical Knowledge', where he explicitly says that "decision is an exercise of knowledge how" (2008, p.407).

However, it is not clear that the account of know-how as a disposition for the reliable *execution* of intentions entitles Setiya to the claim that the *formation* of intentions is itself an exercise of it. In response to being pressed on this issue by Paul (2009, p.555), he concedes this point, retracts the earlier formulation and makes it clear that know-how only "mediates between intention and action" (2009, p.136). Moreover, he adds that the possibility of

irrationally forming intentions despite lacking the relevant know-how indicates that the formation of the intention is not an exercise of know-how.

It seems to me that this concession considerably diminishes the attractiveness of Setiya's overall picture. For it is now clear that the will and know-how are at best 'capacities for knowledge' in rather different senses. The first is a capacity to acquire beliefs that (ideally) constitute knowledge. The second, by contrast, is a capacity to ensure that some beliefs which have already been acquired (by exercising the will) will reliably turn out to be true. Once we are explicit about this, however, it no longer seems appropriate to say that know-how is an epistemic capacity which is a determination of a more general capacity for knowledge of action. Know-how rather seems to operate as an enabling condition of the will's functioning properly, by ensuring that its exercises do not fail the reliability condition on knowledge.

Finally, it is worth repeating the earlier point that the idea that forming an intention is *exercising* know-how seemed congenial to Setiya's response to the wishful thinking objection. The wishful thinker might know that his belief will reliably be made true by the spirit. In my case truth is ensured by my disposition to execute my intention. It is not immediately obvious why I should be epistemically better off. This changes once we allow that I *acquired* my belief by exercising an epistemic capacity while the wishful thinker did not. By conceding that this is not how know-how works, Setiya thus deprives himself of a valuable resource for fending off the objection.

1.3. Conclusion

My purpose has been to draw attention to a few potential gaps in Setiya's account of know-how, knowledge of action, and the interrelations between both. These gaps are significant for two reasons. First, as explained, Setiya needs a non-normative account of knowledge of action to make his argument against the Guise of the Good work. That his success in offering such an account appears unclear threatens to undermine his argument. Second, it might turn out that we can see how the gaps which the examination of Setiya's view uncovered might naturally be filled by thinking about know-how and practical knowledge in terms that are more congenial to the Guise of the Good.

In the following two chapters, I will make proposals about practical reasoning (chapter 2) and practical knowledge (chapter 3) that might constitute parts of an alternative way of understanding the idea that the will is a capacity for knowledge. My aim will not be to give

decisive arguments for these proposals, but just to illustrate their attractiveness. Before moving on, however, let me briefly anticipate some of the claims to show how they relate to the problems I identified with Setiya's view.

The main gaps were vindicating the claim that know-how is knowledge, that it is an epistemic capacity, and that decision and intentional action are ways of exercising know-how. In chapter two, I propose a Rylean conception of both theoretical and practical inference according to which drawing an inference involves exercising theoretical or practical know-how, which is at the same time to be understood as a kind of normative knowledge (namely knowledge of good rules of reasoning). Consequently, know-how, thus understood, is both itself a kind of knowledge and a capacity to know (by performing inferences). Moreover, I will suggest that both acting and deciding involve drawing practical inferences (albeit in different ways), and thus are ways of exercising know-how. Finally, the proposed view is evidently congenial to the Guise of the Good.

In chapter 3, I draw on some recent work of Matthias Haase's on practical knowledge. I follow him in arguing that acting is a distinctive way of knowing the world by changing it and suggest that my preferred account of practical inference fits nicely into his framework. While Haase himself takes his results to threaten a prominent version of the Guise of the Good (Eric Marcus'), I take it as a cause for revising our understanding of it. In a nutshell, according to the resulting proposal, the will is a capacity to know what one has done based on understanding the good of having done it.

Chapter 2: How to Compare Theoretical and Practical Reasoning

The focus of this chapter is on where, if at all, we should expect to find parallels between theoretical and practical reasoning. As I explain, this issue is crucial for the prospects for Setiya's argument against the Guise of the Good. However, it makes a significant difference to how we approach the question about parallels whether we consider theorising about reasoning to be aimed primarily at characterising dynamic or state-like phenomena (mental occurrences or relations between attitudes). Call these dynamic and structural approaches to reasoning. As I will illustrate, both Setiya's arguments against theoretical-practical parallelism and an important response to them share a tendency towards the structural approach. Therefore, I develop a dynamic alternative, hoping to illustrate how proponents of the Guise of the Good might benefit from taking it seriously.

Section 2.1. introduces Setiya's argument against theoretical-practical parallelism and Eric Marcus' response. I show how the question of dynamic versus structural approaches to reasoning fits into the debate and argue that exploring the dynamic approach might reveal some potentially fruitful possibilities that are not often considered.

Sections 2.2. and 2.3. discuss leading arguments for the structural approach to reasoning in the theoretical and practical domains. In 2.2. I take up regress-arguments aimed at showing that theoretical reasoning in the sense subject to the 'taking-condition' (i.e., the requirement that one takes one's reason to be evidence) is not a process but a state. My response will be to offer a Rylean alternative, according to which dynamic reasoning involves 'taking', where this is understood not as a kind of belief, but knowledge how to reason. 2.3. then examines Anscombe's argument to the effect that the dynamic approach leads to an account of validity in practical inference that fails to capture what is distinctive of practical reasoning. I raise some doubts about her account of practical inference, suggesting that a practical analogue of the Rylean proposal might remain attractive.

Sections 2.4.-2.7. attempt to provide further motivation for my proposal by showing how it could allow us to capture various insights of otherwise opposing positions on what conclusions of practical reasoning might be. 2.4 and 2.5. discuss Sarah Paul's (2013) argument for the position that the conclusion of practical reasoning is an attitude prior to action (e.g. an intention or belief), while 2.6. and 2.7. focus on Sergio Tenenbaum's argument for the rival view that the conclusion is the act itself. I suggest we can capture central insights of both positions if we understand reasoning as a kind of process that does not terminate until

the action is finished. On such a view, it looks like the conclusion is neither the action itself, nor an attitude prior to it, but the upshot of the action.

2.8. then brings my points to bear on the Guise of the Good, 2.9. concludes.

2.1. *Practical-theoretical parallels*

A major line of objection to Setiya's argument begins with the question whether that argument could also be used to show that believing something for a reason does not involve taking that reason to justify the belief. After all, it is possible to believe q on the basis that p even if p does not in fact justify that belief. Thus, Setiya's reasoning should lead one to conclude that one may believe that p is one's reason for believing q without believing that p justifies q . Setiya accepts this conclusion but argues that in such cases one's belief about what why one holds a belief is false (2010, p.92, note 26). Why can the GG-theorist not use the same escape route in their account of acting for a reason (Marcus, 2013, p.522)?

In answering this question, one needs to consider whether we should expect to be able to offer parallel accounts of theoretical and practical reasoning. Is the relation between reasons and beliefs like that between reasons and actions, or are there significant disanalogies?

Setiya's answer is that there are. "Reasons attach to actions and beliefs in quite different ways", he writes (2010, p.92, note 26). In later work, he summarises his view as being that epistemic and practical reasons "bear contrasting relations to causality and normative thought" (2013a, p.190). He means by this that there is no use for a causal relation in the right account of what it is to believe for a reason, since a normative belief about the relationship of two belief-contents (e.g. ' p is evidence for q ') is allegedly both necessary and sufficient to establish that a belief is one's reason for holding another. We do, however, have grounds for suspecting that p 's being one's reason for *doing* something requires some sort of causal connection, which in turn renders a normative thought superfluous, or so Setiya argues.

I will briefly indicate why he thinks this, beginning with the argument for Non-Causality about epistemic reasons. Setiya is impressed by the supposed Moore-paradoxicality of statements such as " p and the fact that q is evidence that p , but I don't believe that p even partly because I believe that q ." (2013a, p.191). In the standard case, (p but I don't believe p) Moore-paradoxicality depends on the fact that being in a position to assert the first claim guarantees the falsity of the second. However, if there were a causal requirement on believing

for a reason, it would be difficult to see why being in a position to assert ‘p’ and ‘q is evidence for p’ would guarantee that the relevant causal connection between the beliefs that q and that p obtains (2013b, p.503). To Setiya, this suggests that we should take belief in p, q, and that q is evidence for p *to suffice* for believing p at least partly for the reason that q.

Moreover, he takes this view to receive intuitive support from examples (cf. Lehrer, 1971, pp.311-2): Suppose that, say, Susan’s belief in her son’s innocence is causally explained entirely by her motherly feelings: even if all the evidence indicated he was guilty, she would believe him innocent. Nevertheless, if she also possesses conclusive evidence that he has a watertight alibi, perhaps we should allow that this suffices for her to know of his innocence on that basis. For intuitive support of this claim, Setiya draws our attention to the fact that an agent in this type of situation would be able to answer our question ‘why do you believe that?’ by giving conclusive proof. Knowledge should not require any more than this, or so Setiya argues (2013a, p.191).

By contrast, in the case of action, mere belief that some fact is a reason for acting as one does is insufficient to be doing it for that reason. One might think that making one’s grandparents happy is a reason for visiting them, while in fact doing so only for selfish reasons (thinking about their will etc.). This is the key insight behind Davidson’s (2001a, p.9) argument for Causalism about reason-explanation. Note that Setiya does not regard this argument as sufficient by itself to establish Causalism (2011, p.144), but that is irrelevant here. What matters is only that there is no similar motivation for Causalism applicable to belief (given Setiya’s arguments).

I provide only this rudimentary sketch of Setiya’s arguments because my primary concern is not to discuss them directly, but to make a general point about framing the debate between him and his critics. There is a close connection between reasons for action and belief, and corresponding practical and theoretical inferences.² Thus, a natural approach to examining the parallels between epistemic and practical reasons would be to see whether we can give parallel accounts of theoretical and practical inferences.

At this stage it is important to note that there are different phenomena one might have in mind when speaking of ‘inference’. One might be thinking about a dynamic phenomenon,

² Cf. for example: ‘The premises of a valid practical inference [...] state a reason’ (Raz, 1978, p.5); ‘[...] in saying there is a reason one asserts the correctness of a corresponding inference’ (Clark, 1997, p.18); See also Setiya, 2014 and Way, 2018 for more developed accounts of reasons as premises in reasoning.

something like “a person-level, conscious, voluntary mental action” (Boghossian, 2014, p.2). Alternatively, one might be trying to capture not some sort of process or event, but a certain logical structure in a subject’s attitudes. Inference in this latter sense figures, for example, in this exchange:

‘Why do you think it will rain?’

‘I infer it from the way the sky looks’ (cf. Setiya, 2013a, p.185).

Following Setiya (*ibid*) I will speak of ‘dynamic inference’ and ‘inferential judgment’ to mark this distinction. Bearing it in mind, we seem to face a choice: do we focus on dynamic inferences or inferential judgments when looking for parallels and disanalogies between theoretical and practical reasoning?

It seems to me that the debate between Setiya and his critics is primarily focused on inferential judgments. A cornerstone of Setiya’s reasoning is the claim that Davidson’s argument for Causalism shows that there is no practical analogue of theoretical inferential judgment. If there were, then believing an action to be justified by some reason while performing it would be sufficient for one to be doing it for that reason, and Davidson’s point is that this sounds wrong.

Now, there are reasons for being suspicious of Setiya’s account of theoretical inferential judgment, which is anyway extremely controversial. One might well deny his intuitions both about Moore-paradoxicality, and about the examples (Marcus, 2013, p.513). Consequently, it would seem a sensible response to argue that once inferential judgment is properly understood, we can see that the practical analogue is unproblematic. This is essentially Marcus’ (2013) line. He denies that the mere conjunction of beliefs that p, q, and that p is evidence for q suffices for believing q for the reason that p. Instead, he proposes that we need to posit a distinctive kind of representational attitude, which accounts for the rational connections between our beliefs. Simply believing a proposition q is, according to Marcus, representing it as <to be believed>, while believing it for the reason that p is representing q as <to be believed on the grounds that p is to be believed> (2012, p.17). Moreover, he argues that we can think of this connection between beliefs as a sort of causation, albeit one rather different from what philosophers usually have in mind when using the term. In his view, there are relations of what he calls ‘Rational Causation’ between beliefs, which are constituted by our representations of them.

Obviously, doing justice to Marcus' view would require much further unpacking. The important point, however, is simply that if theoretical inferential judgment cannot be analysed merely in terms of conjunctions of beliefs, then there is no reason to suspect that a practical analogue of it would be analysable in terms of conjunctions of beliefs and actions. Thus, there are no obvious grounds for taking Davidson's point to reveal an important disanalogy between theoretical and practical reasoning. The response to Setiya may then be completed by showing how the practical analogue of theoretical inferential judgment is after all hospitable to the Guise of the Good (This would be unsurprising. After all, even Setiya agrees that believing for a reason involves some sort of awareness of the reason's normative significance). In Marcus' framework, doing A for a reason (e.g. in order to B) is to be understood as a further instance of Rational Causation, constituted by representing the act-type A as <to be done on the ground that B is to be done> (Marcus, 2012, p.67). Thus, practical inferential judgment involves awareness of normative relations between actions.

It is noteworthy that despite their differences, there is an affinity here between Setiya's and Marcus' thought insofar as they seem to agree that progress in the debate about the Guise of the Good will primarily require getting right our account of static rather than dynamic phenomena. Marcus is very explicit that his account of theoretical reasoning is primarily an account of a special kind of state:

'Reasoning' may bring to mind a process [...] But believing-for-a-reason does not consist in a process at all. Rather, it is *a state* that can result from such a process. (2012, p.28, my emphasis)

Note that something similar is true of Marcus's account of acting for a reason. While one may naively expect action to contrast with belief in that it ought to be thought of as dynamic, i.e., as a kind of process or event rather than a state, Marcus devotes a section of his book to arguing that there is a distinctive ontological category of events-in-progress whose instances are relevantly like states and different from processes (2012, pp.212-22). Obviously, I cannot go into Marcus' views about temporal ontology here. All that matters for my purposes is emphasising that in his view, because acting for a reason is a matter of rational-causal relations obtaining between actions-in-progress, the account of practical reasoning, just like that of theoretical reasoning, is in the first instance an account of relations obtaining between state-like entities.

I do not wish to suggest that the kind of line against Setiya pursued by Marcus is not legitimate (though I discuss some problems with his view in chapter 3). My aim is merely to illustrate that in following Setiya's focus on inferential judgment, one risks leaving unexplored an opportunity for defending GG by examining analogies between dynamic inferences in theoretical and practical reasoning.

One potential danger of shifting emphasis away from dynamic inference seems to me to be that it might lead us to look for analogies in the wrong places. I can most easily explain what I mean by briefly going back to my remark that while Marcus has arguments to back up his account of practical reasoning as concerning rational-causal relations between state-like entities (actions-in-progress), one might find that idea surprising. To my mind, when asked to envisage what an account of practical reasoning as parallel to theoretical reasoning would look like, it seems most natural to start by thinking about theoretical reasoning as a mental process, constituted by events such as Setiya's dynamic inference, through which we are taken from some mental states to others. One might think of practical reasoning as an analogous mental process, distinguished by the kinds of events constituting it (decisions, bodily actions, etc. rather than just (theoretical) judgments, inferences, etc.).

However, if such an approach were attractive, its possibility would be obscured by the focus on inferential judgment. Consider the most obvious candidates for the role of conclusion in a practical inferential judgment: normative beliefs (Raz, 2011, ch.7), intentions (e.g. Broome, 2001; Clark, 1997³; Paul, 2013), or actions themselves (Tenenbaum, 2007; Fernandez, 2015, 2016). None of these options are compatible with the idea sketched above. This is because dynamic inferences usually have inferential judgments as their upshot: the conscious mental event of inferring p from q often results in a state of believing p for the reason that q . But if actions are constituents in the mental process of practical reasoning, they cannot also be upshots of that process. The same problem arises for the other views because they take practical conclusions to be something upstream of action (i.e. intentions or beliefs about possible future actions). By contrast, if actions were analogues of dynamic inferences, practical conclusions would need to be something downstream of them, something that could be the upshot of action in the way an inferential judgment can be the upshot of a dynamic judgment.

³ Clark's view is that the practical analogue of theoretical inference is not properly called 'inference', but that is irrelevant to my purposes here

I suspect that this kind of thought might easily tempt one into supposing that practical reasoning, construed as a process, consists in psychological episodes preceding action. However, it is possible at least in principle to think of it as a process that extends throughout the execution of actions.⁴

The possibility of this kind of view seems to be sometimes overlooked. For example, in a recent discussion of the conclusion of practical reasoning, Patricio Fernandez sets up the issue as one of choosing between two main options:

Burleyanism: To act on the basis of reasoning is to knowingly perform a movement in conformity with a normative judgment that has been reached through reasoning, where the action and the judgment are “distinct existences.” (2016, p.872)

Aristotelianism: The conclusion of practical reasoning is an action. (*ibid*, p.873)

(Note that we should understand Burleyanism as also covering views according to which practical reasoning concludes in intention).

However, these options do not seem exhaustive. Consider a further statement he makes in characterising Aristotelianism:

Aristotelianism entails that an agent’s practical reasoning has not concluded until she acts: no antecedent representation of an action she should or will perform [...] can be the proper conclusion of her reasoning (*ibid*)

Basically, my point is just that one might agree with Fernandez on this point (contrary to his Burleyans) without being an Aristotelian in his sense. That would be one’s position if one thought of the action as partly constituting the process of reasoning *which is yet to conclude*. The event of completing the action might then count as an event of *concluding* the practical reasoning. The *conclusion*, however, would then most naturally be understood as the upshot of one’s practical reasoning and therefore of one’s action (metaphorically put: the reasoning is traveling towards a destination, concluding is arriving, and the conclusion is being there).

Again, I realise that this way of thinking runs contrary to how Aristotelians understand what they are trying to offer an account of, namely not a kind of process but a certain structure in our theoretical and practical thought. My point is merely to illustrate that by shifting our focus back on reasoning as a process, we might bring into view further theoretical possibilities which are not always explicitly discussed.

⁴ See Jenkins, 2020 for arguments that bodily actions can be constituents of reasoning in the dynamic sense.

I do not suppose that there are no difficulties with the option I gestured at. I see at least two main problems. First, what exactly would it mean to say that a practical conclusion is somehow the upshot of action? Second, there are independent arguments to suggest that focus on dynamic inference could not help the Guise of the Good, which I discuss in the next section. The first question will be addressed briefly at the end of this chapter and in the next.

2.2. Dynamic inference and the ‘taking-condition’

In his (2013a) Setiya argues against the view that inference in the dynamic sense requires ‘taking’ one’s reasons to justify the belief one forms. ‘Taking’ is usually understood as a representational attitude (the most obvious candidate for which is belief⁵) with normative content. If that were right, then we should not expect to be able to find support for the Guise of the Good in a practical analogue of dynamic inference. Inferential judgment would be the only sensible place to look.

Similarly, Markos Valaris (2014, 2019) argues that the ‘taking-condition’ does not fit into an account of theoretical reasoning understood as a process and consequently suggests that we abandon this processive conception of reasoning. Such arguments, if successful, clearly motivate pursuing Marcus’ line of attack on Setiya rather than the alternative I suggested. I therefore discuss them in what follows.

My main aim will be to indicate that we face a choice here. Instead of revising our conception of reasoning, we might revise our conception of the ‘taking-condition’. The objection to the taking-condition I discuss is that it leads to two kinds of regress. However, it is interesting to note that Ryle originally employed a similar argument to show that we should not conceive of what we now call ‘taking’ in terms of propositional knowledge, but as know-how. When we say that a given conclusion is justified by a premise, we normally express our knowledge how to reason, rather than knowledge of a normative proposition. I will suggest, that this Rylean approach provides an alternative to Valaris’ and Setiya’s shift in focus away from dynamic inference, while preserving a chance of finding support for the Guise of the Good in an account of reasoning as a dynamic phenomenon.

The first regress comes into view once we ask how the taking-belief is supposed to figure in the inference. As Lewis Carroll (1895) points out, it could not be a premise next to the others, for to infer one’s conclusion from this expanded set of premises would then require a new

⁵ For reasons of space, I will not discuss other possible attitudes (cf. McHugh and Way, pp.319-20)

taking-belief to the effect that the original premises and original taking-belief together support the conclusion etc.

We face a further regress when trying to explain how one can justifiably form the taking-belief. If this required performing a further dynamic inference, we would once again need a further taking-belief, which sets off the regress. However, it is not obvious how else we might come by justified taking-beliefs. Importantly, these beliefs concern epistemic support relations between specific propositions ('this follows from that'), rather than general principles such as modus ponens. While it may be more defensible to claim that knowledge of the latter is non-inferential, this approach appears less promising when it comes to the former (Setiya, 2013a, p.186).

Valaris (2014) suggests that Carroll's regress is especially problematic for views that treat inference as a kind of mental process in which a conclusion-belief is caused by some premise-beliefs. For how might a taking-belief feature in such a causal process? If it were just another member of the set of beliefs that jointly cause the conclusion-belief, next to the premise-beliefs, then it looks like we will end up saying that it is an additional premise, and Carroll's regress shows that it cannot do that. Moreover, it is not obvious which other role the taking-belief could have within a causal framework. One idea he considers is that it is an enabling condition of the premise-beliefs causing the conclusion-belief. Another is that it causes a higher-order belief about what one ought to believe, which then explains the adoption of the conclusion. Neither thought seems promising. Regarding the first, it is not clear why we should accept that the causal connection between one's two beliefs would depend on the presence of the taking-belief (2014, p.108). Regarding the second, he suggests that it simply relocates the problem: if the taking-belief causing the higher-order belief counts as an instance of reasoning, we are back to the regress argument. If it does not, then it is unclear why the same relation that holds between taking- and higher-order beliefs could not simply hold between premise- and conclusion-beliefs (*ibid*).

In Valaris' view, the lesson to learn from this is that we should not think of reasoning as a mental process. The taking-belief is not part of a process that leads to the acquisition of the conclusion-belief, but it is partly *constitutive* of that conclusion-belief. Under standard conditions, believing that p and that q follows from p *simply is* believing that q. As he puts it:

non-basic reasoning just *is* believing that one's conclusion follows from one's premisses, and thereby believing one's conclusion (2014, p.112 his emphasis)

By restricting the claim to non-basic instances, Valaris makes clear that he is talking about cases where the taking-belief is itself held on the basis of further reasoning. However, not all reasoning can be like that, so he recognises the need to give an account of the basic case as well. In the 2014 paper, he suggests that in such cases the belief that the conclusion follows from the premise is explained simply by our reasoning from the premise to the conclusion. Since we are self-conscious, our doing so involves, according to Valaris, non-inferential and non-observational knowledge of believing our conclusion on the grounds of the premise. Moreover, he suggests that it is generally true of us that when we believe that we hold one belief because of reasoning from another, we also believe that the latter is a normative reason for the former (*ibid*, p.113).

One potentially puzzling aspect of this account of basic reasoning is that it appears to reverse the order of explanation of the non-basic account. In the case of the latter, it seemed as if Valaris' point was that we count as reasoning from p to q in virtue of taking q to follow from p. In the basic case, however, things look to be the other way around: we satisfy the taking condition in virtue of reasoning a certain way (McHugh and Way point this out as well in their 2016, note 11).

In more recent work, Valaris' view seems to take a different shape. In his 'Reasoning and Deducing', he aims to call into doubt the common view that reasoning is both i) a mental process and ii) a way of acquiring or revising one's attitudes. Instead, he suggests that there are really *two separate* phenomena, 'deduction' (as he calls it), which is dynamic but not a way of revising one's attitudes, and 'reasoning' (or inference), which does fulfil the latter function but refers to a state. To those who might balk at the claim that 'reasoning' refers to a state, Valaris offers some sample sentences like my earlier example to illustrate the use he has in mind:

'Scientists infer facts about climate in the distant past from Antarctic ice cores', or 'The detective reasons that Alf committed the murder from the fact that the murder weapon was found in his possession' (2019, p.866)

What is the relationship between reasoning and deduction? Deduction is construed as a mental process by which we figure out what follows from what. We putatively do this, very roughly, by considering possibilities and ruling them out if they are inconsistent with our premises. That some possibilities are inconsistent is something we may "just see" (this ability is taken to be a 'basic feature of our capacity for contentful thought', 2019, p.868), or something we need to "work out by deduction" (*ibid*, 869). Importantly, since deduction is

concerned only with consistency relations between contents, being engaged in this process does not essentially involve taking any specific attitudes such as belief to these contents. Using Valaris' example: deducing whether the existence of God is consistent with evil is the same, regardless of one's theological stance (*ibid*, p.864).

By giving us knowledge of what follows from what deduction thus supplies us with the materials for a taking-belief which can together with a premise-belief come to constitute belief in a conclusion. The picture thus seems to be that mental processes do not constitute reasoning, but merely prepare it by making it possible that certain constitutive relations obtain among beliefs (see Boyle, 2011, pp.16-7 for a similar view).

The difference between the later and the earlier version of the view seems to be that given what Valaris says about deduction, he would now distinguish basic and non-basic reasoning by claiming that the former involves 'just seeing' the inconsistency between the premise and negation of the conclusion (*ibid*, p.880), while the latter involves some kind of further mental work.

It is also noteworthy that Valaris denies the applicability of a distinction analogous to that between theoretical and practical reasoning at the level of deduction. This is for the simple reason that deducing is merely a matter of working out relations between contents and it does not make a difference which attitude is taken to these contents. Consequently, it does not matter whether the attitudes one might take are belief- or intention-like (2019, p.870).

Valaris' view is an example of one attitude one might have to how we should theorise about reasoning, and I have explained what motivates it. However, my aim here is to illustrate an alternative.

According to Jennifer Hornsby (2011), a core insight behind Ryle's reaction to Carroll's regress is that any propositional knowledge of the premises and their epistemic relation to the conclusion would amount to merely "durative conditions of mind". What we are trying to capture however is how we rationally pass from one thought to another, a "modification of mind" (2011, p.85). Giving an account of this dynamic aspect of reasoning requires the inclusion of something other than propositional knowledge, namely know-how. Very crudely, we might say: both Valaris and Ryle take Carroll's regress to reveal a tension between a

dynamic conception of reasoning and the requirement of a taking-belief. Valaris holds onto the latter at the expense of the former, Ryle modifies the latter while preserving the former.

Note that Ryle seems to accept that there is a sense in which drawing an inference requires knowledge of what follows from what. He writes:

If he is to merit the description of having deduced a consequence from premisses, he must know that acceptance of those premisses gives him the right to accept that conclusion (2009a, p.274)

However, this knowledge is, according to Ryle, not of a proposition (as proponents of the taking-belief construe it), but it is knowledge of a rule of inference. Knowing such a rule is, according to him, not “a case of knowing an extra fact or truth; it is *knowing how* to move from acknowledging some facts to acknowledging others” (2009c, p.227). It is worth emphasising that Ryle does not merely mean to deny that abstract logical principles are premises in “inferences that are made in accordance with them”, but he suggests that the same thing is true of the most “meaty” and determinate hypothetical statements like “If today is Monday, then tomorrow is Tuesday” (2009b, p.249).

Now, the knowledge-how one needs for drawing an inference is primarily manifested in performances that non-accidentally conform to a rule. However, having this knowledge also involves a capacity for recognising correct or faulty performances in others, and it will usually put one in a position to articulate at least approximately⁶ some of the rules one is applying. The need for exercising this latter kind of capacity arises most commonly when teaching others, defending one’s own operations, or in attempting to work out how a practice might be improved (2009b, p.249-50).

Interestingly, Ryle suggests that the kinds of expressions we might use to describe the conclusions of normative reasoning about belief⁷, such as judgments about epistemic support relations, apply not to anything that happens *before or during* but *after* the drawing of a conclusion. Moreover, they serve to characterise a spectator’s or critic’s perspective on a given piece of reasoning, rather than that of the reasoner herself. Ryle writes:

the fact that ability to use an argument carries with it the ability to ‘see’ the implication, when someone else presents the argument to him, does not require that he is causally bound to do such a piece of ‘seeing’ just before, or just while, he himself uses the argument. The

⁶ this qualification is necessary because Ryle holds that the rules of many practices are not codifiable

⁷ Ryle speaks of ‘seeing implications’, ‘accepting proofs’, ‘acquiescing in statements’ etc. I take it that his points would generalise to judgments such as ‘E provides epistemic support for H’ and so on

contemplative metaphor of ‘seeing’ implications or jokes, which is perfectly appropriate to certain special situations is, for that very reason, inappropriate to others. [...] Seeing jokes is the role of the audience, whereas making them is the job of the jester. The audience can be described in contemplative metaphors, but the jester must be described in executive terms (2009a, p.278-9)

And thereafter:

Only [...] delivered arguments can be examined and only when an inference has been at least mooted, can an implication be seen or missed. We do not first see an implication and then go on to draw a conclusion, [...]. Multiplications have to be done before they can be marked ‘correct’ (*ibid*, p.279)

However, I think it is important to read these remarks while bearing in mind Ryle’s general view that instances of know-how are acquired by internalising the perspective of one’s teacher, who observes, criticises, and helps to correct one’s performance (he speaks of learning “to double the roles of instructor and pupil” as well as becoming one’s own “referee”, 2009a, p.130). Thus, it is no accident that skills are manifested self-consciously, so that we are usually able to say what we are doing and why doing it might be correct or appropriate (cf. Small, 2017, p.69-70). In summary then, I take it that Ryle’s view is that our non-propositional knowledge how to move from one judgment to another explains both the formation of our conclusion belief as well as our being aware of that conclusion’s normative status once we have reached it (which may be articulable in propositional knowledge to the effect that the conclusion is supported by the premise or something like that). However, this latter awareness is not part of our making the inference, and it affords us with a perspective on our reasoning that might in principle be shared by another.

If the primary function of normative expressions is assessing inferences that have already been drawn, what role can normative thought have in guiding us while we are still trying to make up our minds? As I understand Ryle, when we are trying to make up our minds, we are in the first instance trying to acquire intellectual know-how. We have some evidence E and there is no conclusion C such that we now know how to move from E to C. Our reasoning aims at closing this gap. We might acquire such know-how through being taught, but obviously we cannot teach ourselves to do something we do not yet know how to do. Nevertheless, Ryle thinks that in such situations we usually can draw on more general know-how we have and try to mimic what a teacher would do in trying to impart the specific inferential competence we lack. This more general knowledge might consist in being able to tell which kinds of hypotheses are promising, seeing mistakes in our tentative steps, and so on (cf. 1971, p.225). But teacher’s knowledge is for Ryle paradigmatically normative knowledge of how one should reason, so on his model the process of making up your mind

will be shot through with more or less general thoughts about what one ought to believe in certain types of situations. It is just that the highly specific normative judgement we are looking for, such as that *this conclusion* is justified by *these premises* becomes available only after the completion of one's reasoning.

I think we can now see how, from the Rylean perspective, Valaris might appear to misconstrue the relationship between reasoning and mental processes. In Valaris' view, the latter (deductions) do not themselves involve the formation of attitudes, but, as it were, merely prepare it by uncovering connections between contents. Our awareness of these connections then makes it possible for certain constitutive relations to obtain between attitudes, and reasoning consists in the obtaining of these relations. From the Rylean perspective, however, one's 'seeing' of connections between contents is a manifestation of a capacity for making valid steps in reasoning. Being able to 'see' rational connections is comparable to being able to see what the best move for white is in some chess-position. Thus, properly understanding the process 'deduction' would require understanding the rational revision and formation of attitudes as a dynamic phenomenon (i.e. dynamic inference in Setiya's terminology), which Valaris tries to eliminate. Consequently, it seems to me that his view would be open to the following complaint of Ryle's:

Formal logic was, unfortunately, taught from the start in the esteemed geometrical manner, with the result that the epistemology of ratiocination and of intellectual work in general continues to be told chiefly in the contemplative idiom [...]. We are given to understand that to 'cognise' is not to work something out, but to be shown something. Had arithmetic and chess been brought into the curriculum before geometry and formal logic, theorising work might have been likened to the execution of calculations and gambits instead of to the struggle for a bench from which the blackboard can be clearly seen. We might have formed the habit of talking of inference in the vocabulary of the football field, instead of in that of the grandstand, and we should have thought of the rules of logic rather as licenses to make inferences than as licenses to concur in them. (Ryle, 2009a, p.279-80)

Taking stock: the section began with regress-arguments against the claim that dynamic inference involves normative beliefs. My aim was to illustrate two reactions one might have to these arguments, given that one is persuaded that reasoning should involve some awareness of normative relations between premises and conclusions. First, one might pursue an approach like Valaris' and deny that reasoning is a process constituted by dynamic inferences. Second, one might hold onto the idea that reasoning is essentially dynamic but revise our understanding of how normative knowledge is involved. I take Ryle's proposal here to be that the relevant knowledge is a kind of know-how.

In the next section, I turn back to the practical domain, and discuss Anscombe's argument to the effect that the dynamic conception of inference threatens to obscure what is distinctive of practical reasoning. In the rest of the chapter, I hope to show that contrary to her suggestions, the Rylean proposal remains viable in both the practical and theoretical domains.

2.3. *Anscombe on practical inference*

Like Valaris, Anscombe (1989) counsels against construing inference as a mental process.

She writes:

Is there something else which one could call not *just seeing* that the second [proposition] follows from the first, but actually inferring it? I take it, no (1989, p.379 her emphasis)

That is the picture of a logical step: an act of mind which is making the step from premise to conclusion [...] So the dispute seemed one between people who all agreed there was such a thing as this 'stepping' for assertions or suppositions; but some thought they could see such a 'step' also in the case of practical inference, while others just couldn't descry it at all. But there is no such thing in any case! (*ibid*, p.393)

However, as her discussion illustrates, in the case of *practical* inference, the move away from the dynamic conception has a further significance, which is my focus in this section. I have in mind the connection Anscombe draws between the question whether practical inference is a process, and whether its validity requires a necessary connection between premises and conclusion.

As I understand her, Anscombe seems to view the putative connection between validity and necessity as presenting us with a kind of dilemma: on the one hand, she appears prepared to acknowledge that "Validity is associated with necessity" (1989, p.278). On the other, she considers with suspicion existing attempts to find necessary connections between practical premises and conclusions (e.g. von Wright, 1972). Moreover, her position in *Intention* can be read as suggesting that at least part of what sets practical reasoning apart from theoretical reasoning is that conclusions of the latter are necessitated in a way in which those of the former are not.

I will start by briefly unpacking this last claim, before outlining Anscombe's later strategy for escaping the dilemma. This will allow us to see how her opposition to 'logical steps' fits into the overall picture. In the remainder of the section, I then sketch some reasons for being dissatisfied with Anscombe's position, which lead me to maintain that an approach to practical reasoning giving a central role to dynamic inference remains attractive.

As already mentioned, Anscombe says that in theoretical reasoning, the conclusion is supposed to be proven, 'shewn true' or 'necessitated' by the premises (1963, p.58-9). The latter term might seem too strong since it is sometimes enough for the premises to make the conclusion sufficiently probable. As I explain, this does not matter to the general point she is making. What matters is what an account of practical inference would look like if it had to fit this model. Even if practical inferences were distinctive in the types of attitudes involved (desires, intentions), one might nevertheless hold that the conclusion is necessitated in the same way here as in the theoretical case. This would be true if practical reasoning only ever moved from ends to necessary means. If one reasons from an intention to go to Oxford and a belief that one will not go unless one buys a ticket to buying a ticket, one's premise attitudes require the conclusion just as they do when one applies modus ponens in theoretical inference.

Anscombe regards the attempt to construe practical and theoretical reasoning as parallel in this way as mistaken. Essentially, her thought seems to be that trying to get a practical analogue of a 'proof' would either require premises that we could not reasonably ascribe to people, or else could not lead to 'actionable' conclusions (Schwenkler, 2019, p.121), i.e., conclusions such that they rationally require one who draws them to perform some specific action. To illustrate, consider an example of Aristotle's, an inference with the premises (Anscombe, 1963, p.58):

dry food is good for a man

this food here is dry

Now, an actionable conclusion, in Schwenkler's sense, would be something like 'I will eat this food here', 'I have conclusive reason to eat this food' etc. By contrast, the conclusions 'I have some reason to eat this food', 'perhaps I'll eat this', or the like are not actionable, for one might draw them and yet fail to act accordingly without irrationality. The problem is that to get an actionable conclusion out of the above inference, one would need a universal premise, something like a commitment to eat any dry food one will ever be able to consume (Anscombe, 1963, p.61). But this would be an insane premise for anyone to accept. A more reasonable thought, such as that eating dry food is always good in *some respect*, however, will not get one a proof of an actionable conclusion via the second premise.

The lesson Anscombe draws from this is that practical and theoretical reasoning differ in form, that is in how their premises and conclusions relate to each other (Schwenkler, 2019,

p.188). Theoretical premises show conclusions to be true. Practical conclusions, by contrast, are “actions whose point is shewn by the premises” (Anscombe, 1963, p.60).

However, this raises the question why both should be understood as kinds of inferences. In her later work, Anscombe says that anything justifiably called inference must be capable of being valid, and validity, in turn is “associated with necessity” (1989, p.378). Thus, we face the dilemma mentioned above: on the one hand, making sense of the validity of practical inference seems to require finding some sort of necessary connection between premises and conclusions, on the other hand what makes practical inference distinctive is that the sort of necessary connection we find in the theoretical case seems absent there. Anscombe writes:

Practical grounds may ‘require’ an action, when they shew that only by its means can the end be obtained, but they are just as much grounds when they merely shew that the end will be obtained by a certain means. Thus, in the only sense in which practical grounds can necessitate a conclusion (an action), they need not, and are none the less grounds for that. (*ibid*, p.384)

Note that what Anscombe says here also explains why the problem would not go away if we required only a probabilistic rather than necessary connection for practical validity.

Inferences to merely sufficient means would often fail to satisfy even this weaker requirement: in theoretical reasoning, the conclusion must at least be sufficiently probable to favour believing it over believing something inconsistent with it instead. In practical reasoning, however, one can often validly infer a means, which one could just as legitimately refrain from taking.

As far as I understand Anscombe’s suggestion in ‘Practical Inference’, it is that this problem is dissolved once we recognise that we have been looking for necessary connections in the wrong place. According to her diagnosis, we are liable to be confused about theoretical and practical inference, because we mistakenly assume that there must be some sort of “logical compulsion” through which particular conclusion-attitudes are forced upon us (*ibid*, p.391). Anscombe’s main target in her paper, von Wright, takes ‘logical compulsion’ to mean that having the premise-attitudes logically entails having the conclusion attitude. In the practical case: intending to A and believing that B-ing is necessary for Aing *entails* having the intention to B (*ibid*). Anscombe notes that this seems confused. The validity of an inference is a matter of necessary, logical (or probabilistic) relations between propositions that are the contents of attitudes, not the attitudes themselves. Moreover, it would be equally confused to suppose that the validity of an inference requires some sort of psychological compulsion to accept the conclusion when thinking through the argument (*ibid*).

This point should be uncontroversial. How does it help with our problem? We can see how it helps, according to Anscombe, once we recognise that the contents of theoretical and practical inferences, as well as the relations between these contents, are exactly the same. What renders the reasoning theoretical or practical is whether we exploit the relevant connections to move between beliefs, or to realise our aims. The validity of theoretical and practical inferences is thus grounded in the same relations between the same propositions. Anscombe writes:

Now can we not say that there is logical connexion [...] between the truth-connexion of p , $p \rightarrow q$, and q on the one hand, and the transmission (1) of belief from p to q and (2) of intention from *Fiat* $q!$ to *Fiat* $p!$? But the logical necessity involved is only the truth-connexion of p , $p \rightarrow q$, and q ; this truth-connexion is common to both kinds of inference. (1989, p.390)

To illustrate: let p stand for 'I call Calum', and q 'I call a friend', and suppose that Calum is my friend, so that $p \rightarrow q$. Anscombe's point is that no matter whether one infers theoretically 'I am calling a friend' from the fact that one is calling Calum, or whether one forms the intention to call Calum because one wants to call a friend, in each case one's inference exploits the same logical relations between the same propositions.

Properly appreciating this point allows us to see that there is, for Anscombe, a distinctive form of practical inference in only a minimal sense, namely "if all we mean by the 'form' is (1) the casting of certain propositions in a quasi-imperative form, and (2) how the matters are arranged" (Anscombe, 1989, p.393).

I think we can now see how Anscombe's rejection of 'practical steps' fits in here. In reflecting on the relationship between her position in *Intention* and her later view she says:

I must therefore make amends to Aristotle, whom I formerly blamed for speaking of practical inference as 'just the same', as theoretical. I wanted to say it was a completely different form. I believe Aristotle might have had a difficulty in understanding the debate that has gone on about 'whether there is such a thing as practical (or imperative) inference'. For what I believe has lurked in some of our minds has been something which his mind was quite clear of. That is the picture of a logical step: an act of mind which is making the step from premise to conclusion. (*ibid*)

Of course, I cannot claim to be sure, but she appears to be suggesting that the "picture of a logical step" is what prevented her from appreciating that the arguments in *Intention* are compatible with Aristotle's position. While she does not further explain why that is, I take it the idea is that if we think of inference as both a 'logical step' and of its validity as a kind of necessity, we will be pushed towards the kinds of views she rejects (i.e., views on which reasoning to sufficient but non-necessary means will not count as inference). This is because the 'step' or transition from an intention to one of several equally acceptable means is not

necessary in any sense. One could just as well make a different step. However, when we are thinking about inference purely structurally, in terms of relations between contents, no special problem arises for the practical case. Even if there is nothing necessary about the *transition* from intending the end to intending *a* sufficient means, it may nevertheless be said that the *contents* of these intentions can stand in the right kinds of necessary, logical, or probabilistic relations.

This is most obvious if the contents figuring in valid practical inference are just the same as those of corresponding valid theoretical inferences, as Anscombe proposes. It should be noted that one might disagree with her on this point. Will Small (2021, p.268), for example, suggests that the contents of practical inferences are actions ('doables'), rather than propositions ('thinkables'). However, such differences between Anscombe's own view and others are tangential to my main concern here. All that matters is the core idea that the validity of practical inferences is taken to be a matter of a subject's attitudes exhibiting a certain logical structure rather than in terms of specific transitions between attitudes being somehow necessitated (in the sense of being uniquely correct in the circumstances).

If Anscombe's position is correct, it suggests that even if there were a place for dynamic inference in the theoretical domain, trying to give an account of practical reasoning in terms of it would lead us astray, for given the association between validity and necessity, we would be led to a mistaken view about which practical inferences are valid.

In the rest of this section, I raise some concerns about Anscombe's account of practical inference. These concerns will help motivate an alternative picture of practical inference as continuous with the Rylean proposal about theoretical inference outlined above, which, as already illustrated, preserves a central role for dynamic inference in reasoning of both kinds.

As explained, Anscombe's view is that the validity of theoretical and practical inferences is grounded in the same relations between contents. "Looked at in this way", she writes, "we find no special form of practical inference; we have a set of propositions connected with one another the same way in the two cases. The difference lies in the different service to which they are put." (1989, p.392) Given, for instance, p , q , and $p \rightarrow q$, one can validly infer a judgment that q from p , or q an intention to bring about that p from an intention to make it the case that q .

However, one might worry that the connection between p , q , and $p \rightarrow q$ is of at best limited utility for the practical purpose of bringing it about that q . To bring this out, consider

Anscombe's remark that the order in which the propositions are arranged is the same in the case of practical inference as in explanatory inference (*ibid*, p.395). Using her example: we might observe spectacular plant growth and use the conditional "if this stuff is in the ground there is spectacular growth" in reasoning towards putting forward the hypothesis "this stuff is in the ground" for further investigation (*ibid*). Again, the same propositions would figure in reasoning from knowledge of the soil towards a prediction about plant growth.

Now, this looks like Peirce's account of abduction, which he takes to be an inference of the form:

The surprising fact, C, is observed

But if A were true, C would be a matter of course

Hence, there is reason to suspect that A is true [5.189.] (Frankfurt, 1958, p.594)

What is the use of such a schema? In some places, Peirce seems to suggest that it is the generation of new ideas, that abduction is a form of reasoning by which we discover new hypotheses (*ibid*). As Harry Frankfurt points out, this cannot be right. The new hypothesis is A, and to formulate the second premise of the inference, one must already have discovered it. Thus, we do not discover hypotheses through abductive reasoning, but in some other way (*ibid*, 595).

Consequently, a more charitable interpretation of Peirce would be that the function of abduction is not the generation of hypotheses, but that of guiding us towards their adoption for further investigation. However, Frankfurt argues that this is hardly more plausible than the first suggestion. The mere observation that A would render C a matter of course could not suffice to warrant our adoption of A. After all there are going to be infinitely many hypotheses that would entail C as well. What we are looking for is a form of reasoning that can guide our selection among several candidate hypotheses. Peirce's form cannot fulfil that function for the only favourable thing that his premises by themselves allow us to say about any given hypothesis (that it would predict our observation) is just the minimal condition a claim must satisfy to count as a hypothesis to begin with. Thus, if we are trying to explain how we come to select hypotheses, we need to look elsewhere. Peirce himself, who anticipated Frankfurt's point, attributes the selection of hypotheses to our power for "intelligent guessing" (*ibid*). However, now we might wonder exactly what the aforementioned inference pattern is good for. How we come up with hypotheses, and how we

choose among multiple candidates are at least two of the main questions we want an account of explanatory reasoning to speak to. But understanding the former requires an examination of our capacity for imagination, the latter an account of our power for ‘intelligent guessing’. In neither case is abductive inference, as defined, of any help.

I suggest that, since Anscombe thinks of explanatory and practical inference in the same way, this last complaint applies to her account of the latter as well. Two of the main things we might want our account of practical reasoning to explain is how we identify means to our ends and how we select one from a range of candidate means. In neither case will the kinds of logical connections she says practical inference exploits be of any help, and this is for the reasons Frankfurt cites.

Note that an extension of the Rylean proposal introduced above might be more attractive in both the explanatory and the practical case. On such a view, the connections exploited would not be ‘just the same’ in these and the deductive case. That is because all three kinds of reasoning will involve exercising different kinds of know-how, i.e., knowledge how to perform deductions, how to move from observations to good explanations, and how to move from facts about one’s situation to good practical responses. In each case, having the relevant know-how will involve being able to make intelligent selections among candidate hypotheses or courses of action.

Now, I expect that Anscombe’s followers would complain that in taking Frankfurt’s objections against Peirce to carry over to Anscombe, I ignore the feature that distinguishes practical from theoretical thought. The function of theoretical thought is to represent what is actually the case, while the function of practical thought is to make actual what it represents (at least according to those who are inclined to agree with Anscombe). One might think that this generates pressure for theoretical thought to be selective in ways that practical thought need not be. If I know that only one of a range of mutually incompatible hypotheses will be the true explanation of the spectacular plant growth, I know that I cannot pick just any of them. By contrast, if a range of mutually incompatible means will each get me what I want, then insofar as getting what I want is my aim, there is no comparable objection to picking one arbitrarily (cf. Fix. 2021, p.9).

This thought also explains why Anscombeans are likely to be unimpressed by arguments from ‘Buridan’s Ass’-cases, aimed to show that practical reasoning concludes in an attitude

short of intentional action itself (e.g. Paul, 2013, p.296). It is, I take it, what Sebastian Rödl is getting at when he writes:

Practical reasoning proceeds from something general and its office is to arrive at a specification. It is in the nature of the case that there may be more than one way of doing this. Inferring from this fact that practical reasoning fails to reach a definite action is refusing to consider the idea of *practical* reasoning” (2007, p.22, his emphasis)

However, if this is what is going on, then it strikes me as a bit quick. In a nutshell, I want to argue that the thought that practical reasoning does not conclude with anything short of action (the rejection of Fernandez’s Burleyanism) is entirely compatible with the thought that it must be selective in the way Frankfurt says explanatory reasoning must be. Thus, emphasising the practicality of practical reasoning in the aforementioned manner is, on its own, insufficient to address objections from the failure of Anscombean inferences to explain how we choose among available means. The other side of this point, however, is that we cannot straightforwardly argue, as Paul does, from the selectivity of practical reasoning to the view that it must reach its conclusion prior to action. I explain this point in the following section, after setting out the debate about ‘Buridan’s Ass’- cases in more detail.

2.4 *Buridan’s Ass*

In a ‘Buridan’s Ass’-case, one is rationally required to perform one of several actions between which one is rationally indifferent (e.g., choosing one of twenty identical bottles). Such cases are supposedly illustrative of a general problem with Aristotelianism, and thereby provide support for some version of Burleyanism.

A recent development of this line is offered by Paul (cf. Bratman, 1985). She maintains that “[...]to suppose that the particular execution of a token action should be subsumed under the process of reasoning [...] seems to me to modify the notion of reasoning beyond recognition.” (Paul, 2013, p.295) In making her case, she reminds us that “practical reasoning is regulated by some standard regarding the relation between means and end”. This captures the uncontroversial thought that good practical reasoning enables us to realise our ends by taking appropriate means. However, it is crucial to her argument that this form of reasoning does not merely aim at identifying *any* means by which an end might be realised, but *good* means. In her words, “the question [of practical reasoning] is not 'how at all?' but rather 'how well?’” (*ibid*, her emphasis). Such an understanding of the aim of practical reasoning supposedly implies that “if a choice is to be understood as a transition in practical reasoning

at all, it must at least implicitly involve the application of a principle of *choiceworthiness*” (*ibid*, her emphasis).

The problem for the Aristotelian view is supposed to be that forming an intention to perform some particular action can usually not be understood as an application of a principle of choiceworthiness. This is because there are normally numerous equally good ways of realising the same objective: it does not matter which shoe I put on first, which of the twenty identical bottles I take, and so on. Cases such as these illustrate that “[i]n general, the performing of particular actions bottoms out in large part with mere plumping for indifferent, adequate ways of getting it done” (*ibid*, p.296). Paul recommends that we desist from calling such ‘plumping’ an exercise of reasoning, since it does not involve a “further judgment of choiceworthiness” (*ibid*).

It seems to me that the evaluation of this argument is likely to be complicated by the fact that Burleyans and Aristotelians might differ in their attitudes to the question what role dynamic phenomena should occupy in our theorising about reasoning. Therefore, it seems a good idea to start by clearly distinguishing two claims Paul might make. First, that Buridan’s-cases show that certain events (deciding to take *this* bottle, etc.) are not part of practical reasoning *understood as a process*. Second, that in Buridan’s-cases, the means-action does not stand in the conclusion-premise relation to another action or intention.

Much of what Paul says seems aimed at arguing for the first claim. ‘Plumping’ appears to refer to a type of event, and her argument putatively establishes that such events are not part of the process of reasoning. Her speaking of what is necessary for a choice to count as a *transition* in reasoning makes that clear. This creates a dialectical difficulty, for Aristotelians might, like Valaris, be altogether disinclined to think of reasoning as a mental process. The main concern of such theorists (I take Marcus to be an example) is what we should say about the state-like phenomena of believing-for-a-reason and acting-for-a-reason. But this question is not directly about events of any type, so it is not immediately obvious how what we say about plumping-events is relevant.

Early in her paper, Paul seems to connect these concerns. She writes:

The significance of what we deem the terminus of reasoning lies in how we are to understand the rationalizing relation [...]. Is the particular performance of an action the kind of thing that can stand directly in the rationalizing relation to the premises, such that the transition to that event is a step in reasoning? Or are there elements of what it is to realize a token action that are not properly characterized as an exercise of reason, such that reasoning must terminate in an intention to perform a general type of action? (2013, p.289, my emphasis)

As I understand it, the thought is that if two attitudes (or actions) stand in the premise-conclusion relation to one another, then the event of transitioning between them will be a possible ‘step’ in reasoning. Thus, if the transition from an intention to get *a* bottle to the formation of an intention to get *this* bottle would be an event of ‘mere plumping’ then this tells us that the two intentions (or actions) are not related as premise and conclusion.

However, I am not sure that all Aristotelians would be prepared to accept this way of connecting their question (about the premise-conclusion relation) and Paul’s (about events and processes of reasoning). This is because they would reject the assumption that we can meaningfully distinguish events of forming or transitioning between attitudes that are exercises of reasoning (e.g. judging the streets are wet because it is raining) from those that are not (e.g. ‘plumping’ for this bottle). On their view, *no* such events will themselves be exercises of reasoning, for only the states in which they result may be that. Think back to Valaris’ distinction between reasoning and deducing here. The latter is a mental process, but it does not involve the formation of attitudes for reasons (a ‘reasoned change in view’ as he puts it). It merely involves becoming aware of relations between contents. This process enables reasoning but reasoning itself is a state not a process. In this kind of framework, there seems to be no room for asking whether Paul’s plumping-events count as reasoning. Thus, we cannot use that question to decide whether the relation between two actions that may come to obtain as the result of our plumping should count as reasoning or not. Instead, we must find some way of addressing the latter question directly.

For the Aristotelian, the relevant question is whether premise-conclusion relations can obtain between actions and attitudes or other actions, and I have suggested that they might not regard reflection on the nature of plumping-events to settle that question. However, one might now ask: does the mere fact that the relation between the contents of premise- and conclusion-attitudes is, in Buridan’s-cases, the same as that between the same premise and any number of incompatible conclusions undermine the claim that one may validly hold the conclusion-attitude because of the premise? In the theoretical case, that is how it looks: if my evidence supports belief in two incompatible propositions to the same degree, it supports belief in neither.

Aristotelians can argue that this feature of theoretical reasoning is explained by its aim. The aim of theoretical reason is to accurately represent what is true, and “[s]ince incompatible propositions cannot both be true, an exercise of that capacity about one of them is legitimate

only if I have a basis to believe one which is thereby a basis to not believe the other. Without such a basis, I cannot believe either. I must suspend judgment.” (Fix, 2021, p.9). Given this explanation, we might expect the reverse to be true in the practical case. Since I will act as I represent only if I choose one of the available sufficient means, we would expect that practical reason *prohibits* taking neither of the equally supported options, where theoretical reason *requires* the analogous response (suspending belief). Moreover, where theoretical reason *prohibits* choice among options in the absence of reasons to prefer one of them, the aim of practical reason grounds no such prohibition (*ibid*).

Now, with these points in mind, it is tempting to suggest that Paul’s argument illicitly assumes that the standards governing theoretical reasoning must govern reasoning in general. This would be problematic, since taking seriously the practicality of practical reasoning requires recognising that it aims not at *reflecting* but *creating* the truth (Velleman, 1992, p12), and we have reason to suspect that these aims ground different requirements and permissions.

This complaint is made explicitly in Fix (2020) who writes:

The claim is not that M is not choiceworthy given the end but that M is not more choiceworthy than N given the end. Paul infers that willing M is thereby not part of an exercise of practical reason. That inference is sound, though, only if you assume that an exercise of practical reason *consists in, or at least tracks, only judgments* about what there is most reason to do, where ‘most’ excludes ties. This assumption is at odds with the correctness conditions of the exercise of a will and thus implicitly separates practical reason from our will. (p.454, my emphasis)

As I understand it, the accusation here is that Paul assumes that practical reasoning is theoretical reasoning about what there is most reason to do when she alleges that any move in such reasoning requires a further judgment of choiceworthiness.

Essentially the same complaint is made by Schwenkler (2021) against Dancy (2018) (Dancy does think that practical reasoning concludes in action but agrees with Paul that it tracks which action is most favoured). Schwenkler reminds us of Anscombe’s complaint that our conception of knowledge is “incurably contemplative”, i.e. such that “[t]he facts, reality, are prior, and dictate what is to be said, if it is knowledge.” (Anscombe, 1963, p.57; Schwenkler, 2021, p.192). He continues to suggest that Dancy’s conception of practical reasoning as aimed at tracking what is most favoured, is guilty of the mistake Anscombe warned against.

On [Dancy's] account, values play the same role in reasoning to action that considerations of truth and probability play in reasoning to belief [...]: in each case there is a response [...] most favored by the shape of the situation one is in, and one's reasoning is successful only if it concludes in a response of this sort [...] this underlying picture should not be taken for granted, especially if we wish to uphold the Aristotelian idea that practical reasoning concludes, *not merely in a description of the normative landscape*, but in something a person actually does (2021, p.192-3, my emphasis)

However, I am not sure that this response, at least as I have presented it, is entirely fair to Paul and Dancy. To bring this out, consider some things which Ryle says about the nature of thinking in general. In essence, his view is that we count as thinking whenever we do something intelligently, that is by exercising a capacity for making our performance conform to certain standards of doing the thing well or poorly (cf. e.g. Ryle, 2009a, p.17). This seems to me to capture at least part of the spirit behind Paul's remarks about the nature of reasoning. However, Ryle notoriously rejects as mythical the idea that exercising such a capacity can be broken down into a theoretical operation, resulting in what Schwenkler would call a "description of a normative landscape" and a practical operation of making one's conduct fit that description.

Thus, we should not rule out the possibility of a view according to which practical reasoning aims not *just* at successfully executing one's intentions, but at doing so *well*, i.e., in accordance with certain standards, as Paul says. On such a view, reasoning practically is not simply a matter identifying and pursuing ways of achieving what we want, but also essentially a matter of *selecting* among a range of possible ways. However, as I suggested above in drawing the analogy with Frankfurt's critique of Peirce, the problem with accounts of practical inference like Anscombe's is that they derive a conclusion from the premise that it is *a way* of achieving the end. We might put the problem by saying that such inferences are not *selective*. They support picking any sufficient means to the same extent and thus could have no use in settling *which way* to pick.

To summarise, the problem with the aforementioned response to Paul is that to demand that practical inference should be selective is not the same thing as to suppose that its function must be essentially descriptive or theoretical. One will have reason to think otherwise insofar as one agrees with Ryle that capacities for picking appropriate responses to situations are not analysable into further theoretical and practical components. Carroll's regress already illustrates why this should be true of our capacities to respond to theoretical reasons by forming judgments. It is thus not ad hoc to suspect the same will be true in the practical case.

2.5. A concessive response

I do not suggest that these last points count decisively in favour of Paul. Rather, my aim is to indicate that there is a threat of the debate reaching an impasse here. From the Aristotelian's point of view, it may seem that Paul's conception of how to theorise about reasoning is just too different from their own for her arguments to really speak to them. Paul, by contrast, might complain that the Aristotelians' accusation that she supposes practical and theoretical reasoning to be essentially the same misunderstands the general point about reasoning she is trying to make.

To help overcome this impasse, I therefore try to offer a more concessive response to Paul. Before I do, a brief note on how this all fits in with the overall aim of this chapter: First, there is more pressure for me to be concessive in response to Paul than there might be for some Aristotelians because I am trying to explore a line of response to Setiya which gives dynamic inference a central place, while construing reasoning as a kind of process, as Paul does. If Paul's argument established that this process terminates in the formation of an attitude upstream of the performance of the action,⁸ there would be no room for the view I am trying to put on the table. The point of that view was to illustrate the possibility of construing acting for a reason not as a kind of inferential judgment and hence a state-like phenomenon (as Marcus does), but as a kind of dynamic judgment. If Paul were right, however, the only place for such dynamic judgments would be in a mental process that terminate prior to action.

Nevertheless, my aim in responding to Paul is not purely defensive. If successful, my response would show that we could follow Aristotelians in rejecting Burleyanism while generally continuing to think of practical reasoning as someone like Paul might (i.e. as a kind of process, as aimed at selecting good rather than merely sufficient means etc.). Thus, I am trying to motivate the view by advertising it as attractively ecumenical.

My response begins with granting Paul's main conclusion: when I have no reason for preferring whether to C by doing A or B, I simply 'plump' for one of them. To form the intention to do A under such circumstances is not to draw the conclusion of an inference. Superficially, however, it seems to me that one can say this without yet taking any stance on the question whether my reasoning about how to do C has now *terminated*, or indeed on the

⁸ Cf. Paul (2013, p.): "[...] on the Attitude View, it is compatible with an unblemished episode of practical reasoning coming to completion that no action is initiated"

question whether it can terminate anywhere short of actually doing C in some particular way, as Aristotelians allege.

Perhaps the following analogy is helpful: picking up Anscombe's example, let us say you are trying to figure out why there is spectacular plant growth. You consider all the evidence and come to think that two hypotheses H1 (say stuff X is in the ground) and H2 (stuff Y is) are equally well supported. At this stage you might simply 'plump' for one of these hypotheses and adopt it for further investigation. That is permissible, because to adopt a hypothesis is not yet to form any judgment. However, we need not think that this plumping marks the end of your reasoning about why there is plant growth. Instead, you might now proceed to work out what would follow from your hypothesis and come to form judgments such as: if p were the case, then H1 would be highly plausible, upon which you examine whether p and so on.

The point is just to illustrate that while figuring out why p may involve 'plumping' at some stages, this does not yet tell us anything about the proper conclusion of such reasoning. My suggestion will be to argue that we should apply this lesson to practical reasoning, where we could treat action as occupying the role of judgment, while plumping for an intention would be akin to adopting a hypothesis in the service of advancing one's reasoning.

It is possible to flesh this idea out and to provide a preliminary case for it by considering some points Matthew Soteriou (2013, ch.12) makes about mental acts of deciding. Very roughly, on his account, the act of deciding can be thought of as making an assumption about one's future actions. To make an assumption is, in turn, to impose a constraint upon one's reasoning. One achieves this by reasoning in recognition of the constraint. When I assume that it is going to rain tomorrow, I impose on myself the constraint of reasoning as if that were the case, and I manage to do that by, for instance, drawing appropriate conclusions such as the conditional judgment that I will need a coat if it will rain (*ibid*, p.264). This accounts for the difference between assumption and judgment: in the latter case, I am bound by some constraints because of the way the world is, in the former case I am thus bound because I reason as if I were, and am aware of how the bindingness of the constraint depends on how I reason (Soteriou, 2013, p.288).

Now, Soteriou proposes that to decide is likewise to self-impose a constraint by reasoning in recognition of it. In his view, the constraint one imposes is factual: in deciding to A, I impose on myself the constraint of reasoning as if it were true that I will A. Moreover, I recognise that I am bound by this constraint, because I reason in recognition of it. A crucial difference

with merely assuming I will A, however, is that I treat the decision as to be discharged not merely by making conditional judgments such as ‘if I A, B will be true’, but by actually doing the thing (*ibid*).

It is not my primary aim here to discuss Soteriou’s proposal in detail. Rather, I would like to draw attention to a possible variant of it, which I hope will still be in the spirit of what he says. Moreover, I want to show how one might draw on this variant in responding to Paul. In a nutshell, the proposal is to take from Soteriou the idea that deciding is a matter of self-imposing a constraint by reasoning in recognition of it, but to plug a slightly different constraint into the account. Instead of saying that in deciding to A one self-imposes the *factual* constraint of reasoning as if one will do A, I suggest that one instead self-imposes the constraint of *reasoning in accordance with the rules that constitute one’s knowledge how to A* under the given circumstances (‘rules of Aing’ for short). Moreover, I suggest that in cases where one’s practical reasoning is non-defective, its being constrained by one’s imposing the rules of Aing on oneself *makes it the case* that one is either Aing or going to A, as well as grounding one’s practical knowledge of this fact.

To illustrate: we can approximately represent knowledge how to play the Sicilian defence with black as knowledge of a system of rules telling one what to do under which circumstances (if white opens with e4, play c5; if white plays d4 on move two, play c5xd4, and so on. Note, however, Ryle’s point that most know-how is not codifiable by giving lists of propositions like this, thus the illustration is just an approximation). The thought is that in deciding to play the Sicilian, one imposes on one’s practical reasoning about which chess-moves to make the constraints constitutive of this know-how. Moreover, to successfully impose these constraints on one’s practical thought, *is* to be playing the Sicilian, or at least to be going to play it.

How does all of this relate to Paul’s argument? Let me stick with the chess-example. Say you are playing the black pieces, white has just opened with e4, and you now face the decision whether to play the Sicilian or the Scandinavian (I assume, for simplicity, that you have ruled out all other options). In line with Paul’s stipulations, I assume that you have no reason for preferring one line to the other. Now what happens when you ‘plump’ for the Sicilian? Given the proposed variant of Soteriou’s account of deciding, you impose on yourself the rules of the Sicilian. At this stage, it is important to note that to successfully impose the relevant rules on yourself, you must now reason in recognition of them, which in this case means reasoning

from the premise that white played e4 to playing c5 in response. Consequently, it appears that even if plumping for an intention is not drawing an inference, it cannot be the point where practical reasoning terminates. For successful plumping requires further reasoning in recognition of the constraints which one thereby imposes on oneself. Sometimes, one recognises the imposed constraints by reasoning to further, more specific intentions. I respect the self-imposed rules of going-to-Oxford by forming the intention to buy a ticket (i.e., self-imposing the rules of ticket-buying). It cannot be like this all the way down, however. That would lead to a vicious regress. At some point, my reasoning must reach constraints which I can respect not just by imposing further constraints, but by *actually doing* what the imposed rules prescribe. Hence, there is motivation for the view that reasoning must reach down all the way to action.

To summarise the proposal: whenever I am Aing (or going to A) and can do so by taking one of a range of equally adequate means M_1, \dots, M_n , in deciding to go for one of them, I impose on myself an optional specification of the rules that already govern my conduct in virtue of the fact that I am Aing (i.e. the rules of Aing), namely the rules of Aing-by- M_1 ing, or Aing-by- M_2 ing and so on. In the example above I am already constrained by the rules of chess and impose on myself the more specific rules of the Sicilian. To introduce such more specific constraints is not to bring one's practical reasoning to a close, but to alter the manner in which it unfolds.

In the section 2.7., I try to illustrate how my proposal may be able to accommodate some central Aristotelian insights, by examining Tenenbaum's (2007) argument for Aristotelianism. In doing so, I will also suggest that this argument is rendered more compelling given the kind of picture of practical inference that I am trying to advocate than it would be on an Anscombean alternative. To set the stage for this argument, I will first briefly summarise the key points of divergence between my favoured view and the alternatives.

2.6. *Comparing the proposals*

As explained Anscombeans suggest that the principles of practical inference allow inferring conclusions on grounds that equally support incompatible alternatives, whereas theoretical reasoning does not allow such inferences. I have sided with Paul in questioning whether we should take practical and theoretical reasoning to differ in these respects, while also suggesting that contrary to appearances, this point would not entail that practical reasoning terminates prior to action.

What about the contents and attitudes involved in reasoning? According to Anscombeans, both the premises and conclusions of practical inference are desires, intentions, or intentional actions (e.g. Fernandez, 2016; Tenenbaum, 2007; Wiland, 2013). I propose a different view, according to which the premises are theoretical attitudes, whereas the conclusions are intentions or actions. This point of divergence is bound up with a difference in which contents are taken to be involved. On the Anscombean view, the content of the practical inference's first premise⁹ is the agent's aim, and it is given in a desire, intention, or intentional action. The second premise, by contrast, is a theoretical attitude about means to that end. Anscombe herself maintains, furthermore, that the second premise contains the *same kind of conditional* that could also figure in a theoretical inference by which an observer might make a prediction about what the upshot of my action will be (e.g. 'if I take an umbrella, I will stay dry').

To introduce the alternative,¹⁰ recall Ryle's account of theoretical inference. The important point there was that the conditional connecting the premise and conclusion need not, for the purposes of the inference, be an object of propositional knowledge for the reasoning subject. Rather, what is required is knowledge *how* to move from the claim that I take an umbrella to the claim that I will stay dry. Likewise, we might try having the work of the conditional premise in the practical case done by know-how as well. In that case, what matters is that taking an umbrella is the result of exercising knowledge how to stay dry. If the job of know-how is connecting premises and conclusions both in the theoretical and practical case, what is connected to what in practical inference? I think the most natural proposal is that know-how connects knowledge of facts about one's situation to the actions one takes or intentions one forms in response to these facts. This seems natural if we think of know-how as approximable in terms of rules governing some activity: such rules seem to tell one what to do given certain relevant facts (e.g: if white plays d4 on move two, play c5xd4 etc.).

On the emerging picture, the premises of theoretical and practical reasoning are both theoretical attitudes, and their contents need not be means-ends conditionals in the practical case. Moreover, Anscombe's claim that the same conditionals connect theoretical and

⁹ Though note Anscombe's hesitation to speak of a 'premise' rather than 'objective' here. For my purposes the difference does not matter. See also Müller (1979), who argues that the aim should be understood as figuring in a premise while defending what I consider a paradigmatic example of an Anscombean view

¹⁰ I think my proposal may be similar to Dancy's in some respects, but he does not like putting things in terms of inference (cf. 2018, p.24)

practical premises to their respective conclusions turns out false. In both cases different kinds of know-how do that work.

This last point is important for it allows us to say that practical reasoning is selective, as Paul suggests (it aims at taking good rather than merely sufficient means), while still allowing us to say that it is nevertheless not just theoretical reasoning about what one ought to do. For practical reasoning essentially involves exercising knowledge how to perform some action, theoretical reasoning about normative matters need not do that.

It is noteworthy that the proposed view does not give one's aim the role of a premise in practical inference: to do B in order to do A is not to infer B-ing from the aim of A-ing. Rather, the fact that one aims at A-ing is manifested in the way one's practical reasoning towards B-ing is constrained by the rules of A-ing. Thus, saying 'I'm going to the shop to get milk' is the direct analogue of saying 'I judge that the streets are wet, because they are *if* it is raining', whereas the direct analogue of 'I judge that the streets are wet because it *is* raining' would be something along the lines of 'I'm going to the shop because they have milk there'. In the one case, a response is justified by citing a rule that prescribes it, in the other by citing the fact in virtue of which that rule applies.

The proposal that the aim of an action might figure not as a premise but as a principle of inference in practical reasoning is considered and rejected by Anselm Müller. He writes:

Could wanting, or intending, a certain end be considered as a principle of inference—on the following analogy: for someone to infer 'Q' from 'P' he has to believe that if P then Q; for someone to 'infer' action A from (I) he has to want B? The analogy is not convincing: First, for the theoretical reasoning to be correct, the belief concerning the entailment must be true; there is no comparable condition on the primary wants of practical reasoners (whence the 'principles of inference' can here vary from person to person and from time to time). Secondly, if a piece of practical reasoning can be assessed as valid or invalid, an assessment of it as valid will be based on a judgement about the logical relations between its premises and its conclusion; and it is this judgement which has a *prima facie* claim to being the analogue of a judgement (or implicit belief) that one's premises entail the conclusion of one's theoretical argument. (Müller, 1979, pp. 94-5)

However, against the first point, I suggest that to know that p entails (or, in the non-deductive case, is sufficient evidence of) q is to know that a rule permitting reasoning from p to q is a good rule (for instance because following it will be conducive to the aim of knowing the world). Likewise in the practical case: to know that the shop's having milk supports going there is to know that the system of rules recommending this transition (one's knowledge of how to get milk) is a good system of rules to be bound by in this situation (for instance because following it will be conducive to some more general aims one has, such as always

having enough milk, being well-nourished etc.). With respect to the second point, my hope is to provide an alternative to the view that the validity of practical inference should be understood in terms of the logical relations between propositions about action. Instead, I suggest that validity in practical inference is a matter of making transitions in accordance with the rules constraining one's practical reasoning at the time.

2.7. Tenenbaum's argument for Aristotelianism

With this stage-setting in place, I can now explain how Tenenbaum's argument for Aristotelianism works, and why I think it is helped by the conception of practical inference I recommend. Very roughly, the idea is that since intentions represent general ways of acting rather than a particular action (though see Wilson, 1989, pp.120ff.), there are usually going to be many rationally unacceptable ways of executing an intention. To use his favourite example: one might switch on the light by walking to the switch in a perfectly straight line, thus knocking over the computer in one's way. This raises two interrelated problems: firstly, we might wonder how the intention could be justified by my reasons for action if successful execution of it is consistent with acting very badly. As Paul notes, the point here is similar to one made by Davidson in 'Intending'. Davidson thought that it was unclear how unconditional or all-out judgments of goodness could apply to act-types rather than act-tokens. The latter could be good without qualification, but since the types will usually cover bad as well as good act-tokens, it seems that the best thing that could be said for them is that their instances are good *in some respect* (insofar as they have the property that defines the type) (Davidson, 2001b, p.97).

The second problem is how to explain the irrationality of executing an apparently sound intention in a crazy way, if the selection of a particular way of realising the intention occurs by means of something other than reasoning. Suppose my reasons for turning on the light do justify my forming the intention of doing so, but I do that by walking through the computer. Obviously, this is irrational, but why? A straightforward explanation would be that my reasoning concerning how to execute my intention is defective, but this answer does not seem to be available to Paul.

In response to Tenenbaum, Paul suggests that the rationality of forming an intention depends on the reasoner's being disposed to execute it in sensible ways. In her words:

[...] the rationality of an intention should be understood as conditional on the agent's belief that there will be a way of implementing the intention that is [not?] unduly costly in relation to her other concerns, and a disposition to select only such ways (2013, p.298)

While it is clear how this addresses Tenenbaum's first concern, it is less obvious how it helps with the second. Having the necessary disposition is compatible with failing to manifest it on some occasion. So, we might still get scenarios where I rationally form the intention to switch on the light, since I am disposed to execute that intention sensibly, but nevertheless act irrationally, failing to manifest that disposition. Now, one line of defence here would be to question the possibility of that scenario. Perhaps any condition that could block the manifestation of the relevant disposition would also block the charge of irrationality (e.g. a seizure). I set this question aside here, granting that the argument might be met by making this sort of reply work. Instead, I want to focus on Tenenbaum's claim that Aristotelianism has a more straightforward story about the problematic cases. I shall argue that this does not turn out to be the case if we assume a version of the Anscombean account of practical reasoning that I criticised above, whereas the Ryle-inspired proposal arguably does better.

The simplest way to see this is by noting that if the function of practical inference were to derive an action from an aim and a belief about how to realise it, it looks like the practical inference from an intention to turn on the light towards taking the straight path through the laptop is valid, as Paul also notes. After all, the light will be turned on that way. Thus, simply calling the intention and action the premise and conclusion of an inference does not help explain the irrationality, for the reasoning in question does not discriminate between sensible and crazy sufficient means. Now, I take it that Anscombeans will insist the relevant inference is irrational for it undermines other aims of the agents besides turning on the light. We get a valid inference only if we unrealistically assume that the only premise figuring in the reasoning is the aim of turning on the light. Once other premises are added (the aim of not damaging one's property needlessly etc.), the bad action will indeed be ruled out.

However, the problem now is that the explanation locates the irrationality at the wrong stage of the reasoning. Basically, according to the current explanation, the source of irrationality in the computer case is my reasoning from the wrong starting points: I should have reasoned from the complex aim of switching on the light while avoiding needless damage, while I in fact reasoned from the irrationally narrow aim of simply switching on the light. But now the irrationality is rooted in *how I arrive* at the intentions which are executed. I should have formed the intention to avoid damage to the computer, but I did not. This might very well be a failure in reasoning (e.g., I fail to reason from my standing aim not to waste money to the

relevant intention). The problem is that *this* explanation no longer concerns how I execute the intention that I *in fact* form. But it is precisely here that Aristotelians like Tenenbaum claim Paul cannot get away with arguing that the relevant transition happens through a process other than reasoning.

In short: while no one disputes that reasoning does important work upstream from the intentions on which one acts, we have yet to see how reasoning makes a difference downstream of these intentions. Here is how my proposal helps: on my view, practical inference involves transitions from facts to responses in accordance with one's know-how. It is commonly accepted that knowing how to do something well involves knowing how to modify one's way of doing it given the specific circumstances in which one finds oneself. When I turn on the light, the fact that I hate wasting money is part of the circumstances against the background of which I must select how precisely I am going to act. Manifesting the relevant know-how will thus involve being appropriately sensitive to the presence of the laptop. To illustrate further: your knowing well how to practice your instrument usually involves knowing how to do so while having to adjust to the interests of the neighbours, your knowing well how to play chess might involve the ability to adjust to the skills of your opponent. In general, a higher degree of mastery of an activity usually comes with heightened flexibility and sensitivity in the face of an increasing range of further factors which might need to be considered.

It is perhaps interesting to note that from one point of view, Paul actually appears to come quite close to a version of the thesis I advocate. After all, the kind of disposition she invokes might very well just be knowledge how to execute one's intentions (which, with Ryle, I take to be a special kind of rational disposition. Cf. Small, 2017, p.73). If that were the case, however, I think there would be a principled reason for thinking that selecting the appropriate ways of executing one's intention is a form of inference. Just as theoretical inferences are, at least on the Rylean view, transitions between beliefs that manifest what we might call intellectual know-how, it would make sense to think of transitions made in accordance with practical know-how as practical inferences.

2.8. *Revisiting the Guise of the Good*

Having discussed the relationship between theoretical and practical reasoning, it remains to discuss the relationship between practical reasoning and reasoning about normative

questions, and consequently how my account of practical reasoning fits with the picture of the will as a capacity for knowledge of the good.

I wish to consider two views. First, Normativism (the term is Marcus'), is the view that practical reasoning just is reasoning about what is good to do. Practical conclusions are at once actions and judgments of the good. In contrast, Separatists (e.g. Hieronymi, 2009; Silverstein, 2017) think of practical reasoning as distinct from normative reasoning, and they might construe the latter as concerning how to conduct the former (so that judging that it would be good to do A is judging that correct practical reasoning would conclude in Aing). Two apparent advantages of Separatism over Normativism are its abilities to accommodate Akrasia and to give a unified account of normative conclusions (on this cf. Fix, 2021, p.2) (it is not the case that some are intentions or actions while others are beliefs etc.).

An objection to Separatism is that it looks like it construes practical reasoning as simply recapitulating theoretical reasoning about what to do (*ibid*, p.11). First my normative reasoning addresses the question: what should I do? E.g: I need to get up early tomorrow, so I ought to stop drinking. Then I ask myself what *to do* (cf. Hieronymi, 2009, p.205). How do I answer that question? Well, it seems by reasoning: I need to get up early, so I will stop drinking. But this looks like doing the same thing twice-over (Fix, *ibid*).

I believe that my favoured proposal provides a way for maintaining the advantages of Separatism over Normativism without construing practical reasoning as mere repetition. One potentially important point is the Rylean idea that normative judgments are in the normal case expressions of know-how: the judgment that black ought to take on d4 may be an expression of knowledge how to play the Sicilian. A second Rylean point is that normative reasoning is useful mainly whenever we lack knowledge of this sort. Suppose you have excellent knowledge of how to play the Sicilian as black, but you do not know how to respond to the Yugoslav attack. If that is your situation, you will ask yourself what you ought to do after white plays Bc4 on move nine. If your normative deliberation is successful, your concluding 'I should do X' will express not just propositional knowledge, which a small child might express just as well while simply parroting an expert, but you will have knowledge how to respond to Bc4, which is in turn part of your now improved knowledge how to play the Sicilian as black.

Given these remarks, we can say that in the standard case successful normative reasoning leads to know-how which one then applies in practical reasoning: working out what to do

about Bc4 results in know-how which I can express either by saying ‘I ought to do X’, or by simply making a sensible move. Nevertheless, this knowledge I have of what I ought to do is constituted by the same condition that might also constitute your knowledge of what I should do. It is not, as Normativists claim, an intention in the first-personal and a belief in the third-personal case, but knowledge-how in both instances.

What then accounts for the difference between the case where my normative reasoning concludes in the acquisition of know-how which is immediately applied in action and cases where I merely express it by making an evaluative judgment (such as when I act akratically)? The difference is that in the enkratic case, I conduct my practical reasoning under the self-imposed constraint of the rules constitutive of the know-how which is expressed in my normative judgment, while this is not true of me in the akratic case.

This proposal concerning Akrasia requires further elaboration. Particularly, it seems to me that there is a need to clarify its relationship with the Guise of the Good thesis. In doing so, I take up again the question of validity in practical inference, which I only briefly touched on in discussing Müller’s objection to my preferred account of practical inference.

One question I expect some might ask is this: if deciding is imposing a system of rules on oneself, can we just impose any arbitrary rules on our conduct, or, if not, what constrains which rules we can impose on ourselves? Can one decide to eat a plate of mud simply because it is muddy (Anscombe, 1963, p.71)?

Since, given Soteriou’s proposal, we self-impose rules by reasoning in recognition of them, what we can decide will be constrained by which systems of rules we could reason in accordance with. For what we do to count as reasoning, we need to be able to make sense of the distinction between doing it validly and invalidly. Above, I suggested that to reason validly, both in the theoretical and practical case, is to follow a good rule. In our present context, we can thus say that the self-imposition of rules through reasoning requires some capacity for distinguishing between good and bad practical rules. Earlier, I suggested that one way for a system of rules to be good is for one’s following it to be conducive to some more general aim. For example: depending on the context, following the rules of the Sicilian might be a good way of playing chess (if the opponent is unlikely to find a good response, for instance). However, one way of unpacking this is just as saying that the more specific system of rules (Sicilian) is, given context, required, or at least permitted, by a more general system (chess). But it cannot be like this all the way down. We should expect there to be some

systems of rules the goodness of which is not purely derivative of further, more encompassing, systems of rules.

I suspect one might try to pursue several different ideas at this stage, and I here want to merely sketch one possible approach that strikes me as attractive on account of providing a natural elaboration of the idea that the will is a capacity for evaluative knowledge. My starting point are some remarks made by John McDowell (1979, p.332-3) in defence of the claim that virtue is knowledge. He argues that possession of a virtue requires a capacity to recognise certain values. Kindness, for instance, requires a capacity to recognise the property of calling for a certain type of response, which McDowell suggests, we might understand as a “kind of perceptual capacity” (*ibid*). This sort of capacity might itself count as a kind of knowledge, such as knowledge of “what it is like to be confronted with a requirement of kindness”. Slightly more controversially, McDowell also takes the perceptual capacity to be sufficient for virtue. When one acts virtuously, one’s action is, according to McDowell, not the joint upshot of one’s recognition that the situation is such as to demand a certain action and a desire to act as the situation demands, but one’s being motivated to act as demanded *just is* one’s recognition of the situation as calling for the relevant response.

What is the relationship between the evaluative knowledge that is virtue (according to McDowell) and know-how? It seems to me that we could try to understand both in terms of knowledge of systems of rules (though, again, these will usually not be codifiable). There are, however, at least two key differences. First, as illustrated above, the goodness of following the rules constitutive of know-how is derivative of other rules which constrain one’s practical reasoning at a time. In contrast, the goodness of following the rules of kindness, say, is not derivative of anything other than the value of kindness, whose understanding consists in knowledge of these rules. Second, it is often up to us whether we are bound by the rules corresponding to know-how, for we may impose these rules on ourselves when deciding how to pursue a given aim. In contrast, if one has the kind of capacity for recognising the demands of kindness which is constitutive of virtue, one is already bound by the relevant rules. This would be a consequence of McDowell’s claim that virtue cannot be broken down into a capacity to recognise what is good to do and some further desires for acting accordingly, which may or may not accompany that capacity.

To illustrate: the considerate person can recognise in situations the property of calling on her to respond considerately. Moreover, her recognition of that property already involves being

prepared to respond in that manner. One way to think about this seems to me to be that being confronted with a situation's demand for considerateness will function as a constraint on her practical reasoning, just as seeing an Elephant in the room will constrain her theoretical reasoning. Such constraints can then provide the background against which further rules, to be introduced by making decisions and exercising one's relevant know-how, can be evaluated as good or bad. However, since one's knowledge of values like considerateness will usually only provide fairly general constraints on how one reasons practically, this sort of knowledge will not by itself suffice for us to know which particular action to perform in a given case. Being prepared to act considerately is different from being prepared to perform some specific considerate action. To identify such particular actions, we need to impose further, more specific constraints on our reasoning by making decisions and exercising our know-how.

For example: her knowledge of the value of considerateness might enable Susanne to perceive that Sally's lying sick in bed requires her to respond considerately. She might live up to this requirement by acting in all kinds of ways, e.g. by being quiet around the house, keeping her entertained, or by seeing what she might do to help her out. To actually get to do any of these things, Susanne will likely need to make some decisions, and consequently exercise, say, her knowledge how to find out what she can do to help Sally, which may ultimately result in her beginning to walk up the stairs with the intention of asking Sally if she needs something from the shop. Thus, know-how takes us from good systems of rules to good particular actions, sometimes via good specifications of these systems (the former if we can reason directly to an action, the latter if we reason to an intention). By itself, however, know-how does not explain the goodness of the system from which we start. That role may instead be played by the evaluative knowledge that constitutes virtue, as McDowell suggests.

I intend this sketch as an illustration of how value might restrict which kinds of constraints we can impose on our practical reasoning when making decisions. Basically, the thought is just that we can only impose a system of rules by reasoning in recognition of it if there is a possibility of doing the latter well rather than poorly, that is a possibility that what we ultimately end up doing results from exercising our knowledge of a good rule. Since the goodness of these rules ultimately derives from the values we are confronted with, it is hard to make sense of a decision until some relevant value is identified: we cannot make sense of one's deciding to eat a plate of mud simply because it is muddy, since muddiness is not a value-concept, and thus trivially not one from which distinctions among better and worse ways of eating mud might be derived (cf. Sussman, 2009, p.620).

At this stage, it is worth mentioning that my view leaves room for the possibility that there might be a significant disanalogy between theoretical and practical reasoning, insofar as it might turn out that there is a single standard of goodness for rules of theoretical reasoning (perhaps truth, knowledge, credibility, etc.), but there are many incommensurable values relevant to practical reasoning (see e.g. Raz, 2002, ch.3). If that kind of pluralist story were correct, we would frequently face situations in which there is no system of rules we could impose on our reasoning (i.e. no decision we could take), which serves all relevant values well.

The question of pluralism is relevant to how we account for Akrasia. On my view, akratic agents will often know how to decide (and be able to express this know-how in judgments of what they should do), while failing to exercise this know-how. It would be nice to have an explanation of why this happens. One thing to say at this stage might be that in deciding as one thinks one should, one is still going to miss out on some value to be realised by reasoning in accordance with a different (and, as one judges, worse) system of rules (Raz, 2011, p.42). According to this suggestion, practical akrasia would be significantly less puzzling than theoretical akrasia. If there is only one relevant concern in theoretical reasoning, there is nothing to be gained by governing one's reasoning in accordance with an inferior system of rules, there is no concern the inferior system serves which the superior system would not serve better (*ibid*).

While I find this suggestion attractive, I cannot - and need not - commit to defending it here. There are accounts of Akrasia that do not invoke a pluralist story, and which treat it as a phenomenon equally characteristic of theoretical and practical reasoning, explicable in analogous ways in both cases (e.g. Tenenbaum's, 1999; 2007; 2018). If this sort of story were correct, I would see no reason why whichever explanation is ultimately offered for why people sometimes fail to exercise their reasoning capacities in the way they think obligatory should not be available on my favoured account.

2.9. Conclusion

To recapitulate what I have tried to do in this chapter, let me return to the beginning. Setiya accepts that believing for a reason requires taking one's reason to justify the belief. However, this taking-condition holds only for inferential judgments, which are states, and not for conscious mental events of judging. Since Davidson's argument for Causalism allegedly

shows that the state of inferential judgment lacks a practical analogue, Setiya's argument against the Guise of the Good allegedly goes through after all.

Setting things up this way allows us to see at least two possible strategies for Setiya's critics to pursue. The first, Marcus', focuses on contesting the claim that theoretical inferential judgment lacks a practical analogue. The second strategy is to show that there is after all a sense in which dynamic inference involves taking one's reason to justify one's response.

Since Marcus' strategy seems to me more developed in the literature, and since at least some of Setiya's critics are likely to share his scepticism regarding the significance of dynamic judgments and a processive conception of reasoning more generally, my main aim here has been to develop the second option. On my view, both conscious judgments in theoretical reasoning and actions are occurrences that involve exercising know-how. Correct exercises of know-how, manifest one's knowledge of good rules of theoretical or practical reasoning. Moreover, since such exercises are, as Ryle argues, normally self-conscious, they will put their agents in a position to say why their judgment or action was justified.

Hopefully I will have provided some motivation for considering this view a genuine theoretical possibility. None of what I have said amounts to a definitive argument for preferring my proposal to the alternatives. Nevertheless, I have argued that it has at least some appeal insofar as it accommodates core insights of various approaches to practical reasoning, such as Paul's and Tenenbaum's, that are otherwise opposed to each other.

One question I have so far given little attention to is what the conclusion of practical reasoning is. I suggested that the action is not itself the conclusion, but part of a process by which that conclusion is eventually reached. That would suggest that the conclusion of practical reasoning is somehow the upshot of action. I expect that this proposal will sound puzzling. Therefore, the next chapter will cover some material which might help make it more palatable. In a recent discussion, Matthias Haase defends the idea that the distinctively practical knowledge agents have of what they are in the process of doing actually depends on their capacity to know what they affect in so acting in the same way. Haase's position is interesting to me for two reasons: first he takes it to create problems for Marcus' Normativism. If that turns out to be correct, it provides further motivation for my alternative. Second, I take it that the idea that acting gives us a special way of knowing the world by making it a certain way might soften us up to the idea that it is this knowledge of what we have achieved that constitutes the conclusion of practical reasoning.

Chapter 3: The Reach of Practical Knowledge

As was already touched upon in the first chapter, practical knowledge of our own intentional actions is distinctive in that it is non-observational and non-inferential. The question of this chapter is one recently taken up by Matthias Haase, namely how knowledge of what we are doing relates to knowledge of what we have done. Is practical knowledge limited to action-in-progress, or does it extend to what we have accomplished (Haase, 2018, p.412)? Following Haase, I refer to the view that endorses the former alternative as ‘Presentism’. The truth or falsity of Presentism may have implications for the Guise of the Good. As Haase suggests, Marcus’ Normativism might seem to commit one to Presentism. Hence, difficulties with that view would give us grounds for exploring alternatives such as the proposal raised in the previous chapter.

The plan is as follows: 3.1. briefly introduces the Normativist account of practical knowledge and explains why it might commit one to Presentism. Thereafter, 3.2. outlines Haase’s reasons for finding Presentism problematic and discusses Marcus’ attempt to block the implication from Normativism to Presentism, arguing that it fails. This will lead to the question of 3.3, namely whether Marcus’ might respond instead by embracing presentism. I try to raise some doubts about the attractiveness of this move. Finally, in 3.4. I briefly address possible connections between Marcus’ problems with Haase’s challenge and the suggestions made in the previous chapter. 3.5. concludes.

3.1. *From Normativism to Presentism*

Following Marcus, I take *Normativism* to be a position claiming that intentional actions are normative judgments, and that practical knowledge is consequently knowledge of such a judgment. It is thus to be explained along the same lines as self-knowledge of our beliefs. Representing a proposition as true puts one in a position to know that one believes it, similarly representing an action-type as to-be-done allegedly puts one in a position to know that one is in the process of performing an action of that type.

The plausibility of Normativism appears to hinge partly on the observation that being in the process of doing something, in contrast to having done it, requires relatively little: if you are struck by a lightning on your way to the shops, it will not be true that you went to the shops, but this does not show you were not going there before being struck (cf. Falvey, 2000, p.23; Thompson, 2011, p.206). Thus, even though the thesis that representing going to the shops as to-be-done suffices to put you in a position to know that you will indeed go there looks

highly unpromising, the proposal that the same representation might be enough for knowledge of what you are *doing* might not be as hopeless.

There is a question about how far to push this point. In some passages, Marcus seems to be pushing it very far. In *Rational Causation*, he claims that to say ‘I am Aing’ is not to express a belief about what one is doing, but rather *simply to express the sui generis attitude of representing Aing as to-be-done* (just like saying ‘I believe that p’ is usually expressing the first-order attitude of representing p as to-be-believed). “An action-expression”, he writes “is not the agent’s attempt to describe the world” (2012, p.71). Likewise, early in the paper that contains his response to Haase, he says with respect to the question of whether I need to perceive myself doing what I intend to know that I am actually doing it:

[...] the insertion of the ‘actually’ changes the subject. To know that I’m actually doing it, I do rely on perceptual knowledge. But ‘actually’ signals that what’s at issue is a judgment that my plan is working or perhaps that my chances of success exceed a certain threshold. It is the sort of judgment that an observer of my action might make; it is a theoretical judgment (2018, p.327)

The foregoing remarks might be taken to suggest that on Marcus’ view, practical knowledge is really just knowledge of what one intends to do. However, to say that would be to misunderstand how Marcus thinks of the relationship between intending, acting, and representing-as-to-be-done. That relationship is, according to him analogous to the relationship of an acorn, a mature oak tree, and a biological species such as *Quercus robur* (2012, p.87). We use the names ‘acorn’ and ‘oak’ to refer to different developmental stages in the life of an instance of *Quercus robur*. Likewise, intending and doing are supposed to be just earlier and later developmental stages of one and the same phenomenon, representing as to-be-done.

There are two main occasions for saying that in representing Aing as to-be-done one is merely intending to A rather than simply Aing. The first is when pursuit of the aim is still at an early, preparatory stage, where all one is doing is waiting for the time to act, seeing to it that one will not be prevented from acting etc. For example: suppose that on Monday I represent going to the shop on Tuesday as to-be-done. At that stage, my attitude is manifested mainly in omissions (refraining from making incompatible plans etc.), while on Tuesday the same attitude will be manifested in my taking my keys, opening the door etc. Note that Marcus denies that there is a deep distinction between merely intending to act and being in the midst of acting. The latter requires that one is sufficiently far down the line, but how far

will often depend on context. As Marcus says, “It is not clear why we must identify a clean break between before and after the [acting] has really commenced.” (2012, p.89)

The second occasion for speaking of mere intention is when there is a major defect in one’s performance. One might, for instance, try to commit a murder using a voodoo doll. Since, unbeknownst to one, there is no possibility of achieving the end using this means, it will be correct to describe one as merely intending to kill rather than actually killing anyone.

Thus, it is important to note that for Marcus, there is nothing more to action than one’s being at an appropriately late stage of *non-defectively* representing an action as to-be-done. What makes a normative representation non-defective? That it is the result of good practical reasoning: one selects means to one’s end which one knows to be good means, guided by knowledge of one’s circumstances. One must also have the skills required to correct and adjust the performance as needed (cf. Marcus, 2018, p.329).

Given this exposition of the view, it should be obvious why one might take it to be a version of Presentism. The quotation above from Marcus (2018, p.327) appears to suggest that view, and in *Rational Causation* Marcus also comes close to endorsing Presentism:

Another important distinction is between what an agent represents as to be done and what an agent actually succeeds in doing, i.e., what actions she completes. But that is not relevant here, since *my claims regarding what agents can just say is limited to actions-in-progress*. (2012, p.90. my emphasis. Note that when speaking of what we can ‘just say’ about our actions, Marcus is speaking of our practical knowledge.)

This tendency towards Presentism is all but surprising. After all, the plausibility of the claim that non-defectively representing an action as to be done puts me in a position to know I am performing it trades, as explained above, on the observation that being in the process of doing something entails very little about how much one actually gets done (cf. Thompson, 2011, p.206).

What, then, is the supposed problem with Presentism?

3.2. *Normativism without Presentism?*

Haase’s argument begins with the observation that the difference between merely intending and acting is that the latter entails having made at least some progress. In Marcus’ words: “to be in the middle of doing something, as opposed to being in stasis, is for some part of what I’m doing to already have been done” (2018, p.333). From here, Haase makes two further steps. First, he argues that we know that we are acting only if we know that we have

completed some phase of Aing (e.g. that we have B-ed in order to A) (Haase, 2018, p.228), and second, that the former knowledge is practical knowledge only if the latter is (*ibid*, pp.231-6). Thus, if we have any practical knowledge, we must have at least some practical knowledge of what we have achieved, contrary to Presentism.

This argument appears to make trouble for Marcus, because it is unclear how our knowledge of having B-ed in order to A could be knowledge of what is to-be-done, and hence how it could be practical. However, we might wonder whether one could not agree with Haase that knowing that one is Aing entails some knowledge of what one is done, without agreeing that the former knowledge is practical only if the latter is. Marcus seems ambivalent on this question. In his initial introduction of Haase's challenge, he writes:

The problem is not simply that an agent's knowledge of her action includes knowledge of the past. Were that the problem, it would be tempting to try to mitigate the damage by dividing this knowledge into two components, a practical component (concerned with what I'm doing right now), and a theoretical component (concerned with what I've already done). [...] But not even this backstop is available, for the nature of my knowledge of what I've already done has an ineliminably agentive character (2018, p.332)

This passage might look like it is conceding Haase's second step, but as I understand it, Marcus main point is merely to deny that *all* knowledge of past action, that is both knowledge of past actions-in-progress (one's having been Aing), and of their successful completion, is theoretical. This much is suggested by the fact that after emphasising the agentive character of knowledge of past action, Marcus devotes a substantial part of his paper to arguing that our knowledge of having been Aing is practical.

However, the practicality of knowledge that one has been Aing is consistent with the claim that knowledge of actually having Aed requires theoretical knowledge that one's activity had the intended upshot. Marcus considers this possibility. He writes:

it looks as if the memory of having built the penguin would be some hybrid of remembering being in the midst of doing it—a practical judgment—and a belief, based on perception, that it was in the end done—a theoretical judgment. (*ibid*, p.338)

Echoing Marcus' description, let us call this the 'Hybrid Account'. It is obvious that as Haase sees things, a commitment to the Hybrid Account would be a clear liability, but Marcus' stance on it seems less clear. Anticipating that his opponents will object to his view on the grounds that it leaves us saddled with the Hybrid Account, Marcus argues that there are at least three possible replies available to him:

First, one might reject the idea that a purely practical judgment is necessary to answer Haase's challenge. Second, one might argue that the judgment that the action was complete is itself a practical judgment. And third, one might argue that it is a mistake to think that

knowledge of completion requires more than knowledge of progress. (Marcus, 2018, pp.338-9)

I discuss all three suggestions below. For now, however, it is noteworthy that while the second and third reply appear to aim at showing that Marcus can after all allow that knowledge of having A-ed, rather than merely having been A-ing, can be practical, the first reply looks like it is questioning Haase's assumption that he is obliged to show how that is possible. On one way of understanding it, the first reply is, in a nutshell, that it is acceptable to answer Haase's objection by retreating to the Hybrid Account.

In what follows, I will begin by examining Marcus' arguments that knowledge of completed action can be practical, before discussing the prospects of retreating to the Hybrid Account in the next section.

Marcus' first suggestion in defence of the claim that one can practically know of being finished doing something builds on Anton Ford's observation that intentional actions are not just marked out by the applicability of Anscombe's 'why?'-question, but also by the fact that they raise a corresponding 'how?'-question for their agent. To answer this question is to say what remains to be done. Moreover, a non-defective answer will constitute practical knowledge that one is now doing that which is left to do. Thus, on Marcus view, "[t]he unfolding of the action is made up of the agent's answers to a series of how-questions, informed by what has been done already and also what remains to be done." (*ibid*, p.339)

Now, judging that there is nothing left to do is allegedly a "limiting case" (*ibid*, p.340) of answering the how-question. Just as answering 'Y-ing' in response to 'how to A?' is, according to Marcus a way of knowing that one is now Y-ing in order to A, answering 'nothing' to the same question is supposedly a way of knowing that one is done. This is meant to show that practical knowledge of having finished is continuous with other practical knowledge about which means I am taking at which times: in both cases, I know the relevant facts because I recognise that my intentional A-ing raises a question of how to A. It is just that when I am finished the answer is that there is nothing left to do.

The trouble with this suggestion is that Marcus seems wrong in arguing that "Just as "no reason" does not show that the "why?" question was misplaced, neither does "nothing" show that the "how?" question was misplaced" (*ibid*). There is a significant disanalogy. Saying that I am A-ing for no reason does not entail that I am not A-ing. Saying there is no more A-ing left to do, however, entails that my aim is realised, and thus that I can no longer be engaged in

realising it. To answer the how-question by saying ‘nothing’ is consequently more like answering that one cannot A. The answer would show that when I raised the question, I was falsely taking myself to be Aing. Hence, knowing that one is finished by seeing there is nothing left to do is not continuous with knowing that one is doing what remains by representing that as to-be-done. In the former case, one discovers that a presupposition of one’s question was mistaken, in the latter the question is apt and answered.

However, Marcus does not lean on the foregoing line of response, suggesting that it is rendered unnecessary by the third option mentioned above (*ibid*, p.342). The core idea of that response consists in the conjunction of two thoughts. First, that completion of an event does not require anything over and above a sufficient amount of progress, and second that in the case of intentional action, one is at any moment making progress only if one knows that one is. Thus, Marcus claims “Just as completion itself does not require more than more progress, knowledge of completion does not require more than more knowledge of progress” (2018, p.341).

The task of properly understanding this response is, in my view, further complicated by the connections Marcus draws between the first and third response. He states that he sees the latter as supplementing the former, a statement he explains as follows:

The reason why an observation can be the ground of my stopping without rendering my knowledge of completion theoretical is because it is built-in to the character of the relevant sort of progress that it is facilitated by observation. To make enough of this facilitated-by-observation progress is to complete the action, and so to know (by knowing all of the progress one has made) that one has completed the action (*ibid*, 342)

This passage echoes his earlier characterisation of the first response:

if knowing that I was putting [the spatula in the drawer] is distinctively practical, then it’s not clear why, even if my knowledge that I’m finished is theoretical, my knowledge of the action would not vindicate Haase’s formula. After all, practical knowledge is [...] generally facilitated by theoretical knowledge. That I only know to pick up that Lego on the basis of perception does not undermine the practicality of my knowledge that I am building a Lego penguin. Similarly that I know that the model is complete on the basis of perception does not obviously render my knowledge that I built it theoretical. It is just the last in a series of perception-based theoretical judgments on which I rely in performing the action. (*ibid*, p.339)

To my mind, these passages illustrate the need to clarify the relationship between practical and perceptual knowledge. Getting clear on this relationship is supposed to help with Haase’s challenge, but there seems to be a weaker and a stronger version of that idea in Marcus’ paper.

The challenge, I take it, is to account for the truth of Haase's formula (i.e. practical knowledge that one is Aing entails that one knows of some B that one has B-ed in order to A). As Marcus explains, perceptual knowledge facilitates practical knowledge insofar as it grounds our non-defective representations of what is to-be-done. Because I see that this is the Lego I need, I can represent picking it up as <to-be-done because building the penguin is to-be-done>, and this latter representation constitutes my knowledge that I am picking it up in order to build the penguin. Now, the weak version of Marcus' proposal might just be that one's transitioning from representing, say, B-ing as <to-be-done because Aing is to-be-done> to representing C-ing (let us assume C-ing is the next step after B-ing) in the same way is non-defective because one has perceptual (or other theoretical) knowledge that one has successfully B-ed. Thus understood, Marcus' point is simply that since intentional action and practical knowledge depend on all sorts of perceptual, theoretical knowledge, the fact that it depends on perceiving that one has finished specific parts of one's project poses no special difficulty. This is a version of the Hybrid Account.

Marcus' third reply, however, appears more ambitious. In the first quotation above, he seems to be suggesting that a perception that, say, the model penguin is fully assembled could enable one to know practically that one finished building it just as a perception that this Lego fits could enable practical knowledge that one is grabbing it.

However, it is difficult to see how this works. Marcus says that "[j]ust as completion itself does not require more than more progress, knowledge of completion does not require more than more knowledge of progress." (2018, p.341) But this is not obviously right. Suppose you are battering someone to death. While you are battering, you know that you are. However, in your rage you may fail to notice for a while that the offender has indeed been neutralised, at which point you falsely take yourself to still be engaged in killing the already deceased. Nevertheless, it appears right to say that you had knowledge of killing the person for the whole time that you were doing so, right up until you had in fact killed him. Obviously, self-consciously making some progress up to the point at which it is *in fact* sufficient for your having finished, does not entail that you know it is enough (Haase points this out as well, see Haase,2018, p.246).

In general, the problem for Marcus' third reply seems quite straightforward. Act-types are individuated partly by some characteristic end-state, their telos. For some event to be a completed instance of an act-type, is for its telos to be brought about through rational

causation (Marcus, 2012, p.246), which is to say that the other parts of the action explain the obtaining of the end-state because the agent represents the normative relations between them. Now, given this framework, it seems clear that the observation that the telos of my Aing is realised cannot be part of my Aing. I cannot bring about an end-state that already obtains. So even if my Aing is generally facilitated by theoretical knowledge, this knowledge can only play that role so long as it does not entail that the telos is already realised.

In light of these difficulties, it would seem advisable for Marcus to abandon the more ambitious lines of response that try to make room for practical knowledge of completed action and just endorse the Hybrid Account. So, what would be wrong with that?

3.3. *A problem for Presentism*

According to the Hybrid Account, knowing that one has A-ed is a matter of having practical knowledge that one is A-ing, and theoretical knowledge that one's aim in A-ing is realised as a result of one's activity. Haase objects that the epistemic order of priority is reversed in at least some cases. That is to say, sometimes we know that the telos of Aing obtains because we know we are done A-ing, rather than the other way around (Haase, 2018, p.239).

According to him, knowing by doing is a *sui generis* way of knowing. I can know that something is the case by self-consciously making it the case, rather than by perceiving or being told that it is. If this is right then the Hybrid Account does not cover all our knowledge of what we have done and is thus inadequate.

One way of bringing out the appeal of Haase's position is by noting that at least some pursuits appear to leave it up to the agent to decide mid-action whether a given state is a good stopping point. Whether she is still acting will then depend on her judgement that what she has achieved so far is good enough. Indeed, it is often part of knowing how to do something well that one can make these kinds of judgments on the spot. Consider writing an essay. If I am content with the incoherent mess that is my first draft and leave it there, I will have written an essay. But it is not true that as soon as I have something I could technically decide to call my finished work, I am no longer writing the paper. It appears up to me to decide when my essay-writing is over in a way in which it is not up to me when my killing is over.

This seems to me to be an intuitive picture of how we can know practically that we are finished. Not by *discovering* but by *deciding* that there is nothing left to do. In this sort of case, my judgment that I am finished does not reflect the fact that the state that now obtains

realises the telos of my action, but rather makes it the case that the state that will obtain as the result of my stopping is this particular action's end state. Generalising, we might say that I can have practical knowledge of completion where the telos of my action is that a state of some general type obtains (e.g. I produce *an* essay), and I can determine by deciding on the spot which specific state (e.g. the state of my writing right now) will realise that type on this occasion.

It is difficult to see how the Hybrid Account might accommodate cases like the essay-writing example. When I decide the time is right to stop, it is not the case that the state of my work at that time suffices for the completion of my action. This is a significant point of contrast with the paradigm example of theoretical knowledge of what is done, namely when I notice after a little while that my battering has finally resulted in killing the man. For me to have theoretical knowledge of having written my essay, that knowledge must somehow be sensitive to my decision to stop when I do (since it is underdetermined by everything else I know prior to so deciding). Perhaps I have practical knowledge of forming the intention to stop but need to infer from knowledge of this intention that I will actually stop, or perceive that I am no longer working on it. But all of this seems implausible. It is much more natural to say that deciding to stop immediately gives me practical knowledge that I am done.

Might Marcus say instead that I know I finished writing in virtue of representing the to-be-doneness of stopping now as a consequence of the to-be-doneness of writing an essay? The problem with this is that the relevant notion of stopping does not fit into Marcus' account of intentional action and practical knowledge. As said before, what is represented as to-be-done is that some as of yet unrealised telos is brought about through my thought about how to do so. Deciding to stop now, however, is supposed to give me knowledge that the way things are at this very moment is sufficient for my being finished. But the way things are at this very moment cannot be represented as to-be-brought about through practical thought.

3.4. *Practical inference revisited*

I want to note a final problem for Marcus' position because it is connected to his conception of practical inference, and I believe it illustrates an advantage of the rival conception I recommended in the previous chapter. Towards the end of his paper, Marcus anticipates the following objection:

It might seem as if [...] what is remembered by someone who built a Lego penguin lacks the unity that such a thing in fact possesses. For all I've said, what I remember might be a disconnected hodgepodge of the various phases of penguin-building and a single moment of finding the penguin fully assembled. (Marcus, 2018, p.342)

He addresses this concern by considering how it is that we perceive ordinary non-agential processes, such as the falling of a tree, as unified wholes. According to him,

[...] my experience of the tree falling is intimately connected to my understanding of the event as a causal process. My experience of the connection between each phase of the fall is informed by my seeing each phase as occurring because of the prior one. [...] The crucial point is that my apprehension of the unity of the event is a function of my seeing or 'seeing' it as a manifestation of a particular species of mechanistic causality. (Marcus, 2018, p.343)

Analogously, we supposedly see the different phases of our actions as forming a unified whole, because we understand the principle that unifies them. It is just that in this case that principle is not ordinary causation, but Rational Causation:

[...] the causal connections between the parts of actions are themselves constituted by normative judgments about how the to-be-done-ness of the end-action confers to-be-done-ness on actions that facilitate the agent's achieving the end. Furthermore, my knowledge of the unity is itself an aspect of the reality of that unity [...]. For those *rational-causal connections between the phases of action depend on my having in view my own aims and drawing practical inferences in light of them.*

My memory of an action is thus not of a series of instants [...] finally capped by seeing that the end had been reached. Rather I remember what was, while it was unfolding, always already a unity—one constituted by my possession of the locally governing intention. [...] Often, if not always, my memory of an action includes the memory of having completed it as a culmination of the various actions taken in support of its end. (*ibid*, 343-4, my emphasis)

However, there seems to me to be a significant disanalogy between the falling of the tree and one's building the model. In the case of the tree, there seems to be no problem in understanding how the completion of the fall could stand in the same causal relations to earlier phases of that process in which those earlier phases stand to each other. By contrast, as illustrated earlier, it is not clear how the state of the model's being built or the event of it becoming complete could stand in rational-causal relations to anything that belongs to the action of building it. That is for the simple reason that practical inferences from the premise that an act is to-be-done (the model is to-be-built) could not take us to a conclusion that entails that the act is done (the truth of the conclusion would entail that the premise is incorrect), and Rational Causation consists, according to Marcus, in such inferences (see the aforementioned quotation).

The source of the problem seems to be the conception of the premises of practical inferences as including a representation of the act I am performing. In the previous chapter I suggested that we might instead think of the premise as simply a fact about the world and take one's

aim in acting to provide the inference's principle of validity. Moreover, I proposed a dynamic rather than a structural approach to inference. This picture seems advantageous in the present context, since it allows us to say, as Marcus seems to want to say, that the same kinds of connections obtain between the different stages of one's A-ing, and one's A-ing eventually being finished, as well as that these connections are constituted by practical inferences. On my preferred view, one's being finished might be the conclusion of a practical inference whose premise is a fact about the current stage of one's performance (in virtue of which it is a good time to stop). Moreover, to draw such an inference is to complete the action.

3.5. *Conclusion.*

This chapter has focused on whether we can have practical knowledge not only of what we are doing but also of what we have done. I suggested that contrary to what he argues, Marcus' view seems to leave little room for practical knowledge of completed action. I have also followed Haase in arguing that it would be desirable to make room for such knowledge. Finally, I proposed that my favoured account of practical reasoning and inference potentially has the advantage of being able to accommodate one of Haase's main insights: on Haase's view doing something can be a distinctive way of knowing the world. On my suggested view, acting involves drawing practical inferences which conclude in knowledge of what we have accomplished. Note that regardless of Haase's views on this issue, my way of fleshing out the idea that acting is a way of knowing the world is congenial to the Guise of the Good: just as in theoretical reasoning, when all goes well, drawing a practical inference usually puts us in a position to cite a justification for the conclusion, i.e., an explanation why it is good that we did what we did.

Conclusion

In this thesis, I examined three variants of the claim that the will is a capacity for knowledge. First, Setiya's non-evaluative account of acting for reasons, second the family of views which consider actions (or at least intentions) to be at once conclusions of practical reasoning and judgments of the good, and finally my proposal according to which actions are constituents in practical reasoning (understood as a process) which, if all goes well, concludes in knowledge of what we have done and what good there was in doing it.

As stated at the outset, I do not take myself to have provided decisive grounds for favouring any of these proposals over the others. Nevertheless, I hope that some of my results will be relevant to how we are inclined to choose between them. I suggested, for instance, that Setiya's non-evaluative view relies on a problematic conception of practical knowledge and its relationship to know-how, and that the Rylean conception of theoretical and practical inference might offer a more natural picture of this relationship.

A second major theme was the choice of whether to prioritise dynamic or structural phenomena in our theorising about reasoning. I attempted to show that at least some of the reasons for reluctance towards the dynamic approach, such as Valaris' and Anscombe's arguments, are not decisive, and that it might be possible to develop such an approach in an attractively ecumenical way. It may be possible to agree with authors like Paul on how we should think about reasoning in general, while nevertheless insisting, as Fernandez does, that practical reasoning is not over until action is over.

Finally, I suggested that my preferred account allows us to say that practical reasoning concludes in knowledge of what we have done, rather than what we are in the process of doing. If Matthias Haase is right, this would be an advantage over the more familiar view that the normative judgment with which reasoning concludes is the action itself.

Bibliography

- Anscombe, G. E. M. (1963). *Intention*. Cambridge: Harvard University Press.
- Anscombe, G. E. M. (1989). Von Wright on Practical Inference. In *The Philosophy of Georg Henrik von Wright*, ed. P. A. Schlipp and L. E. Hahn (pp. 377–404). La Salle, Ill.: Open Court.
- Boghossian, P. A. (2014). What is inference? *Philosophical Studies*, 169 (1), 1–18.
- Boyle, M. (2011). “Making Up Your Mind” and the Activity of Reason. *Philosophers’ Imprint*, 11 (17).
- Bratman, M. (1985). Davidson’s Theory of Intention. Reprinted in *Faces of Intention*, Cambridge: Cambridge University Press, 1999 (pp. 209–224)
- Broome, J. (2001). Normative Practical Reasoning. *Proceedings of the Aristotelian Society, Supplementary Volumes 75* (2001), 175–216.
- Carroll, L. (1895). What the Tortoise Said to Archilles. *Mind*, 4(14), 278–280.
- Clark, P. (1997). Practical Steps and Reasons for Action. *Canadian Journal of Philosophy*, 27 (1), 17-45.
- Clark, P. (2010). Aspects, Guises, Species and Knowing Something to be Good. In S. Tenenbaum (Ed.), *Desire, practical reason, and the good* (pp. 161– 201). Oxford: Oxford University Press.
- Dancy, J. (2018). *Practical Shape: A Theory of Practical Reasoning*. Oxford: Oxford University Press.
- Davidson, D. (2001a). Actions, Reasons, and Causes. In *Essays on actions and events* (pp. 1–20). Oxford: Oxford University Press.
- Davidson, D. (2001b). Intending. In *Essays on actions and events* (pp. 1–20). Oxford: Oxford University Press.
- Falvey, K. (2000). Knowledge in Intention. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 99 (1), 21–44.
- Fernandez, P. (2015). A Realistic Practical Conclusion. *American Philosophical Quarterly*, 52, 115– 128.
- Fernandez, P. (2016). Practical Reasoning: Where the Action is. *Ethics*, 126(4), 869–900.
- Frankfurt, H.G. (). Peirce’s Notion of Abduction. *Journal of Philosophy*, 55, 593–596.
- Fix, J.D. (2020). The Instrumental Rule. *Journal of the American Philosophical Association*, 6(4), 444–462.
- Fix, J.D. (2021a). Practical Cognition as Volition. *European Journal of Philosophy*
- Fix, J.D. (2021b). The Unity of Normative Thought. *Philosophy and Phenomenological Research*

- Gregory, A. (2013). The Guise of Reasons. *American Philosophical Quarterly*, 50 (1), 63-72.
- Haase, M. (2018). Knowing What I Have Done. *Manuscrito*, 41 (4), 195-253.
- Hieronymi, P. (2009). The Will as Reason. *Philosophical Perspectives*, 23, 201– 220.
- Hornsby, J. (2011). Ryle’s Knowing-How, and Knowing How to Act. In J. Bengson and M. Moffett (eds.) *Knowing How: Essays on Knowledge, Mind, and Action* (pp. 80-98). Oxford: Oxford University Press.
- Jenkins, D. (2020). Reasoning’s Relation to Bodily Action. *Ratio*, 33, 87-96.
- Lehrer, K. (1971). How Reasons Give Us Knowledge, Or the Case of the Gypsie Lawyer. *Journal of Philosophy*, 68 (10), 311-313.
- Marcus, E. (2012). *Rational Causation*. Oxford: Oxford University Press.
- Marcus, E. (2013). On the Parallels Between Theoretical and Practical Reasoning. *Analysis*, 73 (3), 512-525
- Marcus, E. (2018). Practical Knowledge as Knowledge of a Normative Judgment. *Manuscrito*, 41(4), 319–347.
- McDowell, J. (1979). Virtue and Reason. *The Monist*, 62 (3), 331-350.
- McHugh, C. and Way, J. (2016). Against the Taking Condition. *Philosophical Issues*, 26(1), 314–331.
- Moran, R. (2001). *Authority and Estrangement*. Princeton: Princeton University Press.
- Müller, A. (1979). How Theoretical is Practical Reasoning? *In Intention and Intentionality: Essays in Honour of G. E. M. Anscombe*, C. Diamond and J. Teichman (eds.) (pp. 91–108). Ithaca, N.Y.: Cornell University Press.
- Paul, S. K. (2009). Intention, Belief and Wishful Thinking: Setiya on “Practical Knowledge”. *Ethics*, 119 (3), 546-557.
- Paul, S. K. (2013). The Conclusion of Practical Reasoning: The Shadow Between Idea and Act. *Canadian Journal of Philosophy*, 43(3), 287–302.
- Raz, J. (1978). Introduction. In J. Raz (ed.) *Practical Reasoning* (pp.1-17), Oxford: Oxford University Press.
- Raz, J. (2002). *Engaging Reason*. Oxford: Oxford University Press.
- Raz, J. (2010). On the Guise of the Good. In S. Tenenbaum (Ed.), *Desire, practical reason, and the good* (pp. 111– 135). Oxford: Oxford University Press.
- Raz, J. (2011). *From Normativity to Responsibility*. Oxford: Oxford University Press.
- Rödl, Sebastian (2007), *Self-Consciousness*. Cambridge: Havard University Press.
- Ryle, G. (1971). Thinking and Self-teaching. *Journal of Philosophy of Education* 5 (2), 216-228.
- Ryle, G. (2009a). *The Concept of Mind*. London: Routledge.

- Ryle, G. (2009b). 'If', 'so', and 'because'. In *Collected Essays* (Vol. 2, pp. 244- 260). London: Routledge.
- Ryle, G. (2009c). Knowing How and Knowing That. In *Collected Essays* (Vol. 2, pp. 244-260). London: Routledge.
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford: Oxford University Press.
- Schwenkler, J. (2019), *Anscombe's Intention: A Guide*. Oxford: Oxford University Press.
- Schwenkler, J. (2021). How Can there be Reasoning to Action? *Analytic Philosophy*, 62 (2), 184-194
- Setiya, K. (2008). Practical Knowledge. *Ethics*, 118 (3), 388-409.
- Setiya, K. (2009). Practical Knowledge Revisited. *Ethics*, 120 (1), 128-137.
- Setiya, K. (2010). Sympathy for the Devil. In Sergio Tenenbaum (ed.), *Desire, Practical Reason, and the Good* (Oxford: Oxford University Press), pp. 82-110.
- Setiya, K. (2011). Reasons and Causes. *European Journal of Philosophy*, 19 (1), 129-157.
- Setiya, K. (2012). Knowing How, *Proceedings of the Aristotelian Society*, 112 (3), 285-307.
- Setiya, K. (2013a). Epistemic Agency: Some Doubts. *Philosophical Issues*, 23(1), 179–198.
- Setiya, K. (2013b). Causality in Action. *Analysis*, 73 (3), 501-512.
- Setiya, K. (2014). What is a reason to act? *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* ,67 (2), 221-235.
- Setiya, K. (2016). Akrasia and the Constitution of Agency. In his *Practical Knowledge: Selected Essays*. Oxford: Oxford University Press.
- Silverstein, M. (2017). Ethics and Practical Reasoning. *Ethics*, 127 (2), 353-382.
- Small, W. (2017). Ryle on the Explanatory Role of Knowledge How. *Journal for the History of Analytical Philosophy*, 5 (5)
- Small, W. (2021). The Practicality of Practical Inference. In A. Haddock and R. Wiseman (eds.) *The Anscombean Mind* (pp. 253-290). New York: Routledge.
- Soteriou, M. (2013). *The Mind's Construction*. Oxford: Oxford University Press.
- Sussman, D. (2009). For Badness' Sake. *The Journal of Philosophy*, 106 (11), 613-628.
- Stanley, J. and Williamson, T. (2001). Knowing How. *Journal of Philosophy*, 98 (8), pp. 411-444.
- Tenenbaum, S. (1999). The Judgment of a Weak Will. *Philosophy and Phenomenological Research*, 59 (4), 875-911.
- Tenenbaum, S. (2007). The Conclusion of Practical Reasoning. In S. Tenenbaum (ed) *New Trends in Philosophy: Moral Psychology*. Rodopi. 323-343.
- Tenenbaum, S. (2009a). Knowing the Good and Knowing What One is Doing. *Canadian Journal of Philosophy*, 39 (1), 91-117.

- Tenenbaum, S. (2009b). In Defense of ‘Appearances’. *Dialogue*, 48 (2), 411-421.
- Tenenbaum, S. (2018). The Guise of the Guise of the Bad. *Ethical Theory and Moral Practice*, 21, 5–20.
- Thompson, M. (2011). Anscombe’s Intention and Practical Knowledge. In A. Ford, J. Hornsby, & F. Stouland (Eds.), *Essays on Anscombe’s Intention* (pp. 76– 104). Cambridge: Harvard University Press.
- Valaris, M. (2014). Reasoning and Regress. *Mind*, 123 (489), 101–127.
- Valaris, M. (2019). Reasoning and Deducing. *Mind*, 128 (511), 861-885.
- Velleman, J. D. (1989), *Practical Reflection*. Princeton: Princeton University Press.
- Velleman, J. D. (1992). The Guise of the Good. *Nous*, 26 (1), 3-26.
- Von Wright, G. H. (1972). On So-Called Practical Inference. *Acta Sociologica*, 15 (1), 39-53.
- Way, J. (2017). Reasons as Premises of Good Reasoning. *Pacific Philosophical Quarterly*, 98 (2)
- Wiland, E. (2013). In the beginning was the doing: The premises of the practical syllogism. *Canadian Journal of Philosophy*, 43 (3), 303-321.
- Wilson, G. M. (1989), *The Intentionality of Human Action*. Stanford University Press.