

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



## How inference isn't blind: Self-conscious inference and its role in doxastic agency

Jenkins, David Russell

*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](mailto:librarypure@kcl.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.

**How inference isn't "blind":  
Self-conscious inference and its role in doxastic agency**

*A thesis submitted to King's College London for the degree of*

*Doctor of Philosophy*

2018

David Russell Jenkins

## Abstract

This thesis brings together two concerns. The first is the nature of *inference*—what it is to infer—where inference is understood as a distinctive kind of conscious and self-conscious occurrence. The second concern is the possibility of *doxastic agency*. To be capable of doxastic agency is to be such that one is capable of directly exercising agency over one’s beliefs. It is to be capable of exercising agency over one’s beliefs in a way which does not amount to mere self-manipulation. Subjects who can exercise doxastic agency can settle questions for themselves.

A challenge to the possibility of doxastic agency stems from the fact that we cannot believe or come to believe “at will”, where this in turn seems to be so because belief “aims at truth”. It must be explained how we are capable of doxastic agency despite that we cannot believe or come to believe at will.

On the orthodox ‘causalist’ conception of inference for an inference to occur is for one act of acceptance to cause another in some specifiable “right way”. This conception of inference prevents its advocates from adequately seeing how *reasoning* could be a means to exercise doxastic agency, as it is natural to think it is. Suppose, for instance, that one reasons and concludes by inferring where one’s inference yields belief in what one infers. Such an inference cannot be performed at will. We cannot infer at will when inference yields belief any more than we can believe or come to believe at will. When it comes to understanding the extent to which one could be exercising agency in such a case the causalist conception of inference suggests that we must look to the causal history of one’s concluding act of acceptance, the nature of the act’s being determined by the way in which it is caused. What results is a picture on which such reasoning as a whole cannot be action. We are at best capable of actions of a kind which lead causally to belief fixation through “mental ballistics”.

The causalist account of inference, I argue, is in fact either inadequate or unmotivated. It either fails to accommodate the self-consciousness of inference or is not best placed to play the very explanatory role which it is put forward to play. On the alternative I develop when one infers one’s inference is the conscious *event* which is one’s act of accepting that which one is inferring. The

act's being an inference is determined, not by the way it is caused, but by the self-knowledge which it constitutively involves. This corrected understanding of inference renders the move from the challenge to the possibility of doxastic agency to the above ballistics picture no longer tempting. It also yields an account of how we are capable of exercising doxastic agency by reasoning despite being unable to believe or come to believe at will. In order to see how such reasoning could amount to the exercise of doxastic agency it needs to be conceived of appropriately. I suggest that paradigm reasoning which potentially amounts the exercise of doxastic agency ought to be conceived of as primarily *epistemic agency*—agency the aim of which is *knowledge*. With inference conceived as suggested, I argue, it can be seen how to engage in such reasoning can *just be* to successfully exercise such agency.

## Acknowledgments

This thesis' completion has depended upon many people and in many ways. First, I thank my supervisors. Early in my PhD I worked with Clayton Littlejohn. Clayton was great fun to work with and instilled in me an approach which I am sure has led to crucial successes. Bill Brewer was my primary supervisor for my first two years and was also a joy to work with. He is a brilliant philosopher and guided me to the heart of whatever issue I was working on without ever telling me what to think. Few things were more motivating than the enthusiasm with which Bill would engage with philosophy and with my attempts to produce it. This was so even when my attempts were crude and even when Bill was swamped with other duties. I was influenced by Matt Soteriou's work, meanwhile, some time before he arrived at King's. I ran a reading group on *The Mind's Construction* and tried to wrap my head around its central insights by doing my best to articulate and defend them to more sceptical members of the group. The influence of his work on mine will be evident throughout. I was delighted to hear when Matt was employed by King's, and extremely lucky to be supervised by him for the rest of my PhD, not just because he is a brilliant philosopher, but because he could not be a more dedicated, conscientious and supportive supervisor. He has put more work in to helping me produce this thesis, publish papers and acquire a post-doc than I could ever have reasonably expected of him. I will miss working with him a great deal.

The main thing which has made studying at King's enjoyable is its large and supportive graduate community. I have learnt an enormous amount from my fellow graduate students and have always benefited from their good humour and support during more difficult periods. I would like to thank Matthew Bodle, Vanessa Brassey, Kimi Chen, Mike Coxhead, Alex Franklin, Andrea Giananti, John Heron, Sam Kimpton-Nye, Clare Moriarty and Cecily Whiteley. Special thanks are due to Paul Doody, Jørgen Dyrstad and Jakob Schieder-Hestermann. Talking philosophy and hanging out with them in London is a big part of what has made my time here so fun.

London as a whole also has a great community of philosophers of mind. I have enjoyed discussing philosophy with all those who attend the London Mind

Group and learnt a lot from doing so. Special thanks there are due to Michael Markunas. I have also gained enormously from discussing my ideas with Lucy Campbell, Joe Cunningham, Nico Koziolok, David Owens, Matt Parrott and Mark Textor, among countless others.

I would also like to thank all my family for their love and support throughout my PhD. Finally, I want to thank my partner Danielle. More than anything else, it has been because I have been lucky enough to share my time in London with her that I have enjoyed it so much. This thesis is dedicated to her.

## Contents

Abstract .....	2
Acknowledgments.....	4
Introduction .....	9
Summary of chapters .....	13
1. Inference, self-consciousness and an ontological quandary .....	17
1.1. What is inference?.....	17
1.2. Inference as a conscious act.....	18
1.3. Inference and belief fixation .....	20
1.4. Inference and judgment.....	23
1.5. The possibility of doxastic agency.....	27
1.6. The self-consciousness of inference as a constraint in what follows 33	
1.7. Overintellectualization .....	34
1.8. An ontological quandary.....	35
1.9. Theoretical and practical inference .....	40
1.10. Conclusion.....	41
2. Causalist accounts of inference.....	43
2.1. The causalist conception of inference .....	44
2.2. The causalist conception's motivation.....	45
2.3. The Taking Condition.....	49
2.4. Two versions of causalism .....	54
2.5. The case for the Taking Condition.....	55
2.6. The case against the Taking Condition: overintellectualization...59	
2.7. The case against the Taking Condition: Carroll's regress .....	60
2.8. Trouble for the causalist: inferences beyond one's ken.....	64
2.9. Causalism's motivation undermined.....	67

2.10.	The Taking Condition’s motivation undermined .....	74
2.11.	Conclusion.....	76
3.	Constitutive accounts of inference.....	78
3.1.	An alternative to the causalist conception .....	78
3.2.	Valaris’ account.....	79
3.3.	Valaris’ failure to characterise inference.....	83
3.4.	Rödl’s account.....	87
3.5.	Conclusion.....	92
4.	Inference and self-consciousness as its “silent partner” .....	93
4.1.	What we have so far.....	93
4.2.	Koziolak’s causalist account of inference .....	94
4.3.	How Koziolak fails to accommodate the self-consciousness of inference.....	99
4.4.	Self-consciousness as a “silent partner” .....	105
4.5.	My account of inference.....	108
4.6.	Our ontological quandary resolved.....	112
4.7.	The problem of deviant causation resolved.....	114
4.8.	Taking beliefs and the intelligibility of inference .....	114
4.9.	The account in sum.....	120
4.10.	Developing the account further .....	121
4.11.	Taking beliefs as neither epistemically nor causally prior .....	124
4.12.	Conclusion.....	127
5.	Self-conscious reasoning as doxastic agency and “making up our minds”	128
5.1.	The challenge to the possibility of doxastic agency again .....	128
5.2.	Reasoning as action and the token identity view .....	130
5.3.	Concerns for the token identity view .....	138



5.4.	Reasoning and mental ballistics .....	141
5.5.	“Making up our minds” and doxastic versus epistemic agency	144
5.6.	An alternative conception of doxastic agency.....	153
5.7.	The role of self-consciousness in doxastic agency.....	155
5.8.	Incidental agency and the aim of inquiry .....	160
5.9.	Conclusion.....	163
	Bibliography .....	166

## Introduction

My concern is *inference* and its role in *doxastic agency*. By ‘inference’ I mean the kind of *conscious occurrence* of a sort which is apt to occur as a constituent of conscious reasoning. I operate throughout on the presumption that inference of the sort of concern is a distinctive kind of *self-conscious* occurrence—a presumption which is ultimately to be motivated by the work which the account of inference which it yields can be put to. To say that we are capable of doxastic agency, meanwhile, is to say that we are capable of exercising a direct kind of agency over our beliefs. We are capable of being active with respect to our beliefs in a way which does not amount to mere self-manipulation. We are, one might put it, capable of settling questions for ourselves.

When it comes to explaining how we are capable of doxastic agency a puzzle stems from the fact that there are distinctive reasons to doubt that we are capable of believing or coming to believe “at will”—of believing or coming to believe where our doing so is up to us. We cannot do so because belief “aims at truth”, where what is true is not up to us, such that when one believes  $p$  one must treat  $p$  as true where one does not see the constraint of treating  $p$  as true as self-imposed. To believe at will, meanwhile, would be to treat  $p$  as true whilst seeing the constraint of treating  $p$  as true as self-imposed. How are we capable of doxastic agency, we must then ask, given that we cannot believe or come to believe at will? How are we capable of direct agency over belief given that when we believe and come to believe our doing so is not up to us? Via appeal to the nature of inference, on the view of inference to be developed, it can be explained how we can exercise doxastic agency in the face of the above challenge.

It is standard to accept what I label the ‘causalist’ conception of inference. Advocates of this conception have it that for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  in some specifiable “right way”. When one makes such an inference one’s inference is often then identified with the *process* beginning with one’s acceptance of  $p$  and terminating with one’s acceptance of  $q$ . This conception is motivated by the thought that appeal to it yields the best explanation of why various cases of non-inference fail to amount

to cases of inference. I argue that the causalist conception is either inadequate, due to its being unable to accommodate the self-consciousness of inference, or is in fact unmotivated, due to its failing to be best placed to play the very explanatory role which it is put forward to play.

A natural response to the challenge to the possibility of doxastic agency is to say that we can exercise it by *reasoning*. Although we cannot believe or come to believe at will, we can exercise doxastic agency by actively reasoning. In paradigm cases of such reasoning, however, one will conclude by inferring, coming to believe what one infers by doing so. And when one does so one does not infer at will. This is often taken to suggest that when such reasoning yields belief it cannot, as a whole, be action. For such an inference to occur, according to the causalist, is for an act of acceptance to be a causal product of prior acts in an appropriate way. The causalist is thus led to the causal history of one's concluding act when it comes to identifying the extent to which one can be exercising agency in coming to believe in the way in question. What results is a picture on which we are not really capable of doxastic agency and are at best capable of actions of a kind which lead causally to belief fixation—of “mental ballistics”.<sup>1</sup>

On the alternative account of inference which I develop inference is a kind of essentially self-conscious *event* and *act of acceptance*. When one infers one's inference is identical to the conscious and self-conscious event which is one's act of accepting that which one infers. Inference is self-conscious in that when one infers a state of awareness of one's inference *constitutively depends* upon the occurrence which is one's inference. Furthermore, I suggest, we should appeal to the self-consciousness which inference constitutively involves in order to say what it is to infer, rather than to the supposed “right way” in which an act of acceptance must be caused for one to infer. Whether one is inferring when consciously accepting, on the account suggested, does not depend upon whether one's act of acceptance has been appropriately caused by prior

---

<sup>1</sup> I have not quite managed to identify a *mandatory* line of thought from the causalist conception of inference to this ballistics picture of the extent to which we are capable of agency over belief. I rather show how, for an advocate of the causalist conception, the move from the challenge to the possibility of doxastic agency to the ballistics picture is natural and tempting. The two are made for each other—natural companions.

acceptance or upon whether one's act of acceptance will appropriately cause further acceptance. It rather depends upon whether in so acting one is manifesting awareness that one is inferring.

When inference yields belief in what one infers, on the account developed, one's inference does not cause the belief. The event which is one's inferring is rather *token identical* to the event which is one's coming to believe. I argue that for one's inferences to be intelligible to oneself when one infers, as they in fact are, one must have appropriate "taking beliefs" when one infers. When one infers  $q$  from  $p$ , for instance, one must believe that  $q$  follows from  $p$  such that one's making the inference is intelligible to oneself. Such beliefs need not cause one's inferences, as advocates of the "Taking Condition" tend to insist. In fact, I suggest, just as inference can yield belief in what one infers inference can yield belief that what one is inferring follows from what one is inferring from. Again, when this occurs the event which is one's inferring is *token identical* to the event which is one's coming to believe.

Once we endorse the above account of inference, I suggest, the move from the challenge to the possibility of doxastic agency to the ballistics picture on which we are best capable of actions which lead causally to belief is no longer tempting. Once we understand a concluding act of inference which yields belief as having its nature determined by the self-consciousness which it constitutively involves, rather than by its causal history or by what it causally involves, there should be no temptation to look to the causal history of the act when it comes to identifying the extent to which one is exercising agency in so coming to believe. Furthermore, I argue, the account of inference offered renders inference just as it needs to be in order for it to occur as a constituent of reasoning which amounts to doxastic agency. Reasoning can be action and terminate with inference where to infer *just is* to come to believe that which one infers. To reason can thus *just be* to exercise doxastic agency. In explaining how such reasoning can amount to action as a whole despite terminating with inference where one does not infer at will the reasoning needs to be conceived of appropriately. I suggest that such reasoning is best understood as primarily *epistemic agency*—as action the aim of which is *knowledge*. To infer, on the view suggested, can just be to come to know. To reason can thus just be to actively

achieve one's aim of knowledge and in turn amount to doxastic agency. There should be no temptation to see this as requiring that one's reasoning terminate with inference where one infers at will. One's not concluding such reasoning by inferring at will does not prevent one's reasoning from amounting to the active acquisition of knowledge.

I take it that the suggested account of inference, along with the presumption that inference is self-conscious which leads to it, is well-motivated insofar as it yields the above account of doxastic agency. We are capable of doxastic agency, as it is natural to think, because we can actively reason and inquire. We need not be able to believe, come to believe or infer at will in order to do so.

## Summary of chapters

### *Chapter one*

My main aim in chapter 1 is to get clear on what the target phenomenon is. Of concern is *inference* where by ‘inference’ I mean the kind of conscious occurrence which is apt to occur as a constituent of conscious reasoning. I shed further light on what inference of the sort of concern is by making clear what it is *not*. One can infer  $q$  from  $p$  without believing  $p$  and without coming to believe  $q$  or that  $q$  follows from  $p$ . One can make familiar inferences, after all, and can infer from mere suppositions and assumptions. Inference is thus not merely a kind of belief fixation, nor is it merely a kind of judgment. I suggest that inference of the sort of concern is a distinctive kind of self-conscious occurrence. Thus understood, I contend, appeal to inference will be apt to help explain how we are capable of doxastic agency despite that we cannot believe or come to believe “at will” (Williams 1970). I explain why there are distinctive reasons to think that we cannot believe or come to believe at will, stemming from belief’s “aiming at truth” (Williams 1970). I likewise explain what it would be to be capable of doxastic agency. I also suggest that we face a quandary when it comes to saying what one’s inference is identical to when one makes an inference. Accordingly, I propose that rather than beginning with the question ‘What is an inference?’ we should begin by asking ‘What is it for an inference to occur?’, the latter question being neutral on what one’s inference is identical to when one infers. The question to ask is ‘What is it for an inference to occur?’, where inference is understood as a distinctive kind of conscious and self-conscious occurrence.

### *Chapter two*

My aim in chapter 2 is to assess the orthodox ‘causalist’ conception of inference. On this conception, for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  in some specifiable “right way” (Boghossian 2014). Advocates of the “Taking Condition” have it that this requires that one’s taking it that  $q$  follows from  $p$  play an appropriate causal role in one’s coming to accept  $q$ . The Taking Condition’s having been the focus of much of the recent literature on inference, I consider its motivation and the objections its advocates

face, finding the objections wanting. A simple and seemingly adequate reply, for instance, is available to the common suggestion that endorsing the Taking Condition leads to a regress of the sort to be found in Carroll (1895). I go on to argue that the causalist conception of inference allows that inference can fail to be self-conscious and thus fails to capture inference of the sort of concern. In response the self-consciousness of inference might be built into the causalist conception. But to do so, I argue, is to undermine the causalist conception's motivation. The conception is either inadequate or unmotivated. A distinct response would be to suggest that the causalist conception is motivated insofar as the Taking Condition is, the Taking Condition's yielding a version of the causalist conception. I thus go on to reveal why the most promising route to motivating the condition fails. It might be thought that we must have it that whenever one infers  $q$  from  $p$  one's taking it that  $q$  follows from  $p$  must play an appropriate causal role in bringing about one's acceptance of  $q$  such that we can capture the respect in which one's making the inference must be intelligible from one's point of view. On reflection it can be seen that the Taking Condition cannot be motivated by the apparent need to capture the way in which one's inferences are intelligible to oneself in this way. If we are to arrive at an adequate account of inference of the sort apt to help explain how we are capable of doxastic agency then an alternative to the causalist conception is required.

### *Chapter three*

In chapter 3 I consider existing accounts of inference which amount to rejections of the causalist conception. On what I label 'constitutive' accounts what it takes for an inference to occur is specified in terms of what inferring constitutively involves rather than in terms of how inferences, or their constituent occurrences, are caused. One such account is offered by Valaris (2014, 2016a, 2016b, 2016c). According to Valaris there can be constitutive relations between beliefs. Believing  $p$  and that  $q$  follows from  $p$  can constitute believing  $q$ . Similarly, believing  $q$  on the basis of believed  $p$  can constitute believing that  $q$  follows from  $p$ . I argue that what Valaris offers fails to yield an account of inference of the kind of concern: he at most characterises a kind of potentially non-conscious belief fixation. On Rödl's (2013) account, meanwhile,

to infer  $q$  from  $p$  one must judge  $p$  and then judge  $q$  on the basis of  $p$ . When one does so one's inference is identical to one's judgment of  $q$ . And the latter judgment is also *token identical* to one's knowledge of one's judging and to one's (at least apparently) realising that  $p$  is sufficient grounds to judge that  $q$ . Rödl's account is too narrow in scope, ruling out the possibility of inference without judgment, and involves unnecessary commitment to a revisionary ontology. Rödl has it that when one infers one's knowledge of one's inference is identical to the *event* which is one's inferring and thus that one's knowledge is an *event* rather than a *state*.

#### *Chapter four*

In chapter 4 I develop my own account of inference via consideration of the shortcomings of Koziol's (2017) causalist account. On Koziol's view inference is self-conscious because when one infers one's inference causes awareness of one's inference. Much like the causalist conception of inference, I argue, Koziol's account of inference's *self-consciousness* is either inadequate or unmotivated. The inadequacy of Koziol's account of the self-consciousness of inference motivates an alternative on which the self-consciousness of inference is not a causal product of one's inferences but a cotemporaneous "silent partner" (O'Shaughnessy 2000, p. 106). When one infers one knows that one is inferring where this knowledge of one's inference constitutively depends on the occurrence which is one's inference. This suggestion yields a way to say what it is for an inference to occur without appeal to what causes or is causally involved in inference. For an inference to occur one must simply accept what one infers from and then what one infers where in doing so one knows that one is inferring, the latter knowledge being a silent partner to one's inference. On the account which results inference is a kind of conscious *event*. When one infers one's inference is identical to the event which is one's acceptance of *that which one is inferring*. And whether one is inferring when accepting depends, not upon whether one's act of acceptance is appropriately caused, but upon whether in so accepting one is manifesting awareness that one is inferring. I aim to shed light on what is involved in knowing that one is inferring via appeal to the need for appropriate taking beliefs in order for an inference one makes to be



intelligible. I also explain how this account avoids the difficulties faced by those considered in previous chapters, how the account can be developed, and what it suggests concerning the place of taking beliefs in inference.

#### *Chapter five*

I finish, in chapter 5, by applying the account of inference developed in order to yield an account of doxastic agency and an explanation of how we can exercise such agency despite that we cannot believe or come to believe at will. We can do so, I suggest, by *reasoning* because reasoning is a kind of action, because reasoning involves inferences as constituents, and because to infer can *just be* to come to believe. Given this, to reason can *just be* to exercise doxastic agency. I explain how the temptation to respond to the challenge to the possibility of doxastic agency by moving to a view on which we are at best capable of actions of a sort which lead causally to belief revision is removed once the causalist conception of inference is abandoned. I likewise explain why, with the suggested alternative account of inference on the table, and with the relevant reasoning conceived of appropriately, there should be no temptation to think that to exercise doxastic agency by reasoning one would need to be able to infer at will where one's inference yields belief. The reasoning by which we exercise doxastic agency is best conceived of as primarily *epistemic* agency—as agency the aim of which is *knowledge* (rather than *judgment* or *belief*). The fact that a process with this aim fails to conclude with one's inferring at will should not prevent us from seeing it as amounting to the active achievement of one's aim. I conclude by showing how the account offered sheds light on the role of self-consciousness in doxastic agency. Self-consciousness is required for and exploited in the exercise of doxastic agency, not because we must know what we believe in order to be active with respect to our beliefs, but because the exercise of doxastic agency is *itself a self-conscious activity*.

## 1. Inference, self-consciousness and an ontological quandary

**Chapter abstract** My aim in this chapter is to make clear what phenomenon is of concern and why. Of concern is *inference* of the sort apt to occur as a constituent of conscious reasoning. I make clear what inference of the sort of concern is *not*. It is not merely a kind of belief fixation. Nor does judging for reasons suffice for the occurrence of inference. I suggest that inference of the kind of concern is a distinctive kind of self-conscious occurrence. With inference thus understood it will be become possible to explain how we are capable of exercising *doxastic agency* by actively reasoning—how our capacity to reason makes us capable of being active with respect to our beliefs. It will also become possible to explain why, as is often suggested, self-consciousness is required for and exploited in the exercise of such agency. I also suggest that we face an ontological quandary when it comes to saying what one’s inference is identical to when one makes an inference. Accordingly, I propose that rather than beginning with the question ‘What is an inference?’ we should begin by asking ‘What is it for an inference to occur?’, the latter question being neutral on what one’s inference is identical to when one infers. The question to ask is ‘What is it for an inference to occur?’, where inference is understood as a distinctive kind of conscious and self-conscious occurrence.

### 1.1. What is inference?

In an influential article, Boghossian (2014) asks “[w]hat is inference?” He suggests that the phenomenon has been “understudied by philosophers” (2014, p. 1).<sup>2</sup> It is, meanwhile, common to appeal to inference in giving accounts of other phenomena. Inference, for instance, is appealed to in accounts of meaning and our grasp of it (e.g. Brandom 1994; Dummett 1991), in accounts of self-knowledge (e.g. Cassam 2014) and in accounts of the a priori (e.g. Boghossian 1997, 2000; C. Peacocke 1993). My aim in what follows is to answer Boghossian’s question by offering an account of inference. I aim to give an

---

<sup>2</sup> Although see Brewer (1995), Ryle (2009a, Chapter 9), Stroud (1979), Thomson (1965) and White (1971).

account of what inference is, of how we make inferences, and of how inference plays its role in our self-conscious and active reasoning. This will in turn shed light on how we are capable of exercising agency over what we believe by reasoning and can be responsible for what we believe accordingly—on how we are capable of “doxastic agency”.

## **1.2. Inference as a conscious act**

First, it is crucial that we get clear on what phenomenon is of concern when discussing inference. The term ‘inference’ is used in various distinct ways. It is sometimes used to refer to a kind of abstract object: an ordered pair of sets, one which is a set of one or more “premises” and one which is a set containing a “conclusion”. The term can also be used to refer to sub-personal information processing rather than, or in addition to, personal-level occurrences. For instance, we often make judgments and form beliefs about the emotions of others, prompted by subtle facial and behavioural cues (Johnson-Laird 2008, Chapter 5). Johnson-Laird is happy to call such judgments and beliefs the products of inference. Marr meanwhile uses ‘inference’ when discussing the computational processes involved in visual perception. According to him “the true heart of visual perception is the inference from the structure of an image about the structure of the real world outside” (2010, p. 68). Both Johnson-Laird and Marr are happy to use the term ‘inference’ to refer to transitions from sub-personal to personal level occurrences or states. They are similarly happy to use the term to refer to transitions which are wholly sub-personal.

On other occasions ‘inference’ is used to refer to what are taken to be processes of two distinct kinds or to processes which are taken to take place in two distinct cognitive systems (e.g. J. S. B. T. Evans 2003; Kahneman 2011). “System 1” inferences are taken to be processes which are unconscious, automatic, quick, low effort, and not accompanied with a sense of voluntary control. In contrast “System 2” inferences are taken to be processes which are conscious, attention hogging, slow, effortful, and typically accompanied with a sense of voluntary control (Kahneman 2011, pp. 20–21).

It is not my aim to suggest that any of the above uses of ‘inference’ are illegitimate or that the sense of ‘inference’ in question below is in any way

primary or proprietorial. I claim only that inference of the sort to be discussed is a phenomenon which is worth discussing. In particular, I will claim, focussing on inference of the kind in question sheds light on the nature of our conscious and active reasoning and on how it is that we are capable of exercising doxastic agency. Given this aim, inference of the sort of concern is that which is apt to occur as a constituent of *conscious reasoning*. It is thus a kind of *conscious* and dateable *occurrence*—a kind of constituent of the stream of consciousness. It is also a kind of conscious *act*. By saying this I do not mean to commit to the stronger claim that inference is a kind of *action* (e.g. Boghossian 2014, p. 11). I mean only that inference is something which we consciously *do*. As Soteriou notes, “[a]rguably, not everything one does counts as an action one performs” (2013, p. 227).

The above makes clear that of concern is a personal-level phenomenon. It is less clear how to place inference of the sort of concern on the System 1/System 2 distinction. What is of concern is not a sub-personal or non-conscious process that would be appropriately classified as a System 1 phenomenon. Yet it is far from clear that of concern is something which is in all cases slow, effortful and attention hogging such that it is clear that a System 2 phenomenon is in question either (Boghossian 2014, p. 2). In fact, it is far from clear that it is even *typical* for conscious inferences to be attention hogging, slow and effortful.<sup>3</sup> Boghossian tells us that what is of interest is what could be called “System 1.5 reasoning” and up (2014, p. 2). But given the above difficulties it seems that rather than getting mixed up in these controversies it is better to simply leave the System 1/System 2 distinction aside. The distinction is not needed to get clear on the phenomena of concern and plays no role in what follows.

Of concern, then, is inference where inference is understood as a kind of conscious occurrence and a kind of conscious act. Saying this much, however, does nothing to distinguish inference of the sort of concern from other

---

<sup>3</sup> As the case of conscious inferences suggests, the properties used to distinguish System 1 from System 2 in fact seem to crosscut one another (Mugg 2015, sec. 2). Mugg (2015) uses such considerations to argue that the System 1/System 2 distinction cannot do the explanatory work it is posited for.

conscious acts. We can begin to do so by noting something distinctive of inference as a conscious act. Whenever one infers there is both *something which one infers* and *something which one infers from*. When one makes an inference one infers something and infers from something else. Whilst playing Sudoku, for instance, I might judge that the answer is either 3 or 5. I might then realise that the answer cannot be 3 and infer that the answer is 5. I infer *that the answer is 5* from *that the answer is not 3* (and perhaps also from *that the answer is 3 or 5*). In sum, so far, we thus have the following to say about what it is to infer in the sense of concern. Inference is a kind of conscious occurrence and a kind of conscious act such that whenever one infers there is something which one infers and something which one infers from.

### 1.3. Inference and belief fixation

A way to bring into focus further what is of concern is to make clear what inference of the sort of concern is *not*. Some have it that one's making an inference entails one's coming to believe what one infers (e.g. Rumfitt 2015, pp. 34–8; White 1971, pp. 290–2). Others have it that one's making an inference entails one's coming to believe what one infers *or* one's revising one's beliefs about the matter in question. Boghossian does so when he claims that when an inference occurs

you start off with some beliefs and then ... end up either adding some new beliefs, or giving up some old beliefs, or both. (2014, p. 2)

Others still have it that making an inference entails coming to believe what one infers *or* that an appropriate inferential relation obtains between what one infers and what one infers from (Valaris 2016b).

Inference of the sort of concern here is a kind of conscious act which does not always yield belief in what is inferred. One can make an inference without coming to believe what one infers. Inferences can, for instance, be suppositional. As Wright puts it

[o]ne may infer from a supposition, or from a proposition one takes to be false (for instance, in the course of looking, perhaps unsuccessfully, for consequences that bring out its falsity.) (2014, p. 28)

In such cases we infer without coming to believe that which we infer.

The same is true in cases in which we infer from that which we merely take for granted—in cases in which we infer from mere assumptions made for practical purposes. As Bratman (1992) notes, we can take things for granted without believing them. To take one of Bratman’s examples, in planning a construction project one might take for granted that the cost of each job will be at the top of one’s estimated range. This cautious approach may be wisest given the risks of underestimating costs. It is clear that what one takes for granted here is not simply what one believes. One does not in fact believe that cost of each job will be at the top of one’s estimated range. If one did then one would not be assuming for practical purposes and would rather be going by what one believes each job will cost.<sup>4</sup> It is also clear that we can infer from such assumptions and that when we do so we will not thereby come to believe that which we infer. As we can infer from mere suppositions without coming to believe what we infer, we can infer from assumptions made for practical purposes without doing so.

Just as we can infer without coming to believe what we infer we can infer without coming to believe that what we infer stands in some appropriate inferential relation to what we infer from. In fact, we can make inferences without this resulting in change in any of our beliefs concerning what we are reasoning about or concerning what inferential relations obtain (Ryle 2009a, pp. 273–4). Consider, for instance, a logic teacher going through a familiar piece of reasoning. She might use the same example every time she teaches her students logical explosion, that is, that anything follows from any given contradiction.<sup>5</sup> And she might go through this familiar reasoning in teaching her students whilst already holding all the beliefs that she hopes her students will acquire. She might, for instance, begin by supposing that

(1) All swans are white and it is not the case that all swans are white.

She can then infer from this that

---

<sup>4</sup> As Bratman (1992, p. 4) notes, what one will take for granted, unlike what one takes to be the case, will vary depending on practical context. In the case in question, for instance, what one takes for granted will vary depending on whether one is planning cautiously or guessing what the actual cost will be.

<sup>5</sup> I do not really need to assume that this is in fact correct, only that the teacher believes so.

(2) All swans are white (by conjunction elimination).<sup>6</sup>

Our logic teacher can next infer from (2) that

(3) All swans are white or the moon is made of cheese (by disjunction introduction).

She can also infer from (1) that

(4) It is not the case that all swans are white (again, by conjunction elimination).

And from (3) and (4) she can then infer that

(5) The moon is made from cheese (by disjunctive syllogism).

Finally, she can discharge the supposition that (1) and infer

(6) If all swans are white and it is not the case that all swans are white, then the moon is made of cheese (by conditional proof).

In going through this familiar reasoning our logic teacher may not change any of her beliefs about the subject matter at hand or change any of her beliefs about the what inferential relations obtain.

The above examples make clear that we can make inferences without changing our beliefs about or acquiring beliefs about the subject matter at hand and without changing our beliefs about or acquiring beliefs about what inferential relations obtain. With inference understood as a kind of conscious act as above it similarly becomes clear that belief fixation can occur and inference not, where the belief fixation in question is of a sort which could have been the product of inference. One might, for instance, come to believe that one's friend's train is delayed and come to believe that they will be late accordingly. As Moran observes, there seems to be no reason to think that such belief fixation "could not take place below the threshold of consciousness" (2001, p. 110). In fact, as Broome (2013 77-8, 206-8) and McHugh (2013, p. 134) suggest, it seems that it commonly does.<sup>7</sup>

---

<sup>6</sup> Some, myself included, will take this to require that she comes to believe, if she does not believe already, that (2) follows from (1), or that some other appropriate inferential relation obtains between (1) and (2) (see sec. 2.3 - 2.10 and 4.8 below). We can suppose that our logic teacher already has all such relevant beliefs. A similar belief about what inferential relation obtains might be taken to be required for each inference to follow.

<sup>7</sup> Moran (2001, pp. 110–11) suggests that things need to be this way if we are to be capable of conscious reasoning whatsoever. It is hard to see how the kind of conscious reasoning which

#### 1.4. Inference and judgment

As the occurrence of belief fixation of the sort which inference can yield is neither necessary nor sufficient for the occurrence of inference, there are grounds to see the occurrence of judgment *for reasons* as neither necessary nor sufficient for the occurrence of inference. It is thus a mistake to equate judging for reasons with inferring—to see inferring  $q$  from  $p$  as one and the same as judging that  $q$  for the reason that  $p$  (e.g. Rödl 2013).<sup>8</sup>

By ‘judgment’ here I mean a kind of conscious occurrence, as by ‘inference’ I likewise mean a kind of conscious occurrence. Specifically, by judgment I mean a kind of conscious *event* in contrast to belief, which is a kind of *state*.<sup>9</sup> Judgment that  $p$  can yield belief that  $p$ . That said, it does not always do so. For instance, judging that  $p$  does not yield belief that  $p$  when one already believes that  $p$  (Soteriou 2013, p. 237; Toribio 2011, p. 346). It also seems that judgement can fail to yield belief even in the absence of prior belief (e.g. Cassam 2010, pp. 81–2; Shah and Velleman 2005, p. 507). A nervous flyer, for instance, might judge that their plane will not crash and yet it may be clear immediately after their doing so that they lack the corresponding belief (Shah and Velleman 2005, p. 507). This appears to be a potential case in which a judgment fails to yield a corresponding belief, and not just because one already holds the relevant belief.<sup>10</sup> Even if that is correct, we can still plausibly think of judgments as conscious acts of a kind that at least typically yield corresponding beliefs (unless we already hold those beliefs).

Judgment, as I use the term, is also a kind of conscious *act of acceptance*, other such acts including supposing, assuming and guessing. One can thus consciously accept  $p$  without judging that  $p$ . And when one judges that  $p$  this

---

we in fact engage in could ever get going if all belief fixation of the sort which inference *can* yield in fact required conscious inference.

<sup>8</sup> Although Rödl does equate inferring with judging for reasons it is not clear that he in fact thinks that it is strictly speaking correct to do so. He might instead be read as aiming to give an account of the central and successful cases of inference. Doing so makes available the disjunctivist move of then going on to understand other cases via reference to this central case.

<sup>9</sup> Belief is typically and perhaps always non-conscious (Crane 2013). Belief, for instance, is a state of a kind which can persist during dreamless sleep.

<sup>10</sup> Boyle (2009, pp. 130–2), Peacocke (2017, sec. 3.1) and Setiya (2013, pp. 184–5) resist this interpretation whilst Golob (2015, pp. 245–6) defends it.



conscious occurrence will typically yield belief that  $p$ , assuming one does not already believe that  $p$ . With judgment thus understood it is clear that it is not necessary for one to judge that  $q$  for a reason in order to infer  $q$  from  $p$ . As seen, we can infer from mere suppositions and from assumptions made for practical purposes. One might infer  $q$  from assumed  $p$  or from the supposition that  $p$  without judging that  $q$ . There are also grounds to see judging that  $q$  where  $p$  is one of one's reasons to do so as not sufficing for the occurrence of inference to  $q$  (Broome 2013, p. 222; Valaris 2014, p. 105).<sup>11</sup> The point is best illustrated with a case. Suppose I look out of the window and see a goldfinch. I consciously judge that there is a goldfinch. We can suppose that one of my reasons for judging is that the bird in question has a bright red face (another might be that it has yellow wing patches). I would, for instance, cite this fact as one of my reasons for judging that the bird is a goldfinch if prompted and would not have judged that it was a goldfinch had I not believed that the bird has the feature in question. The case is one in which I consciously judge that there is a goldfinch and do so for a believed reason: that the bird in question has a bright red face. But doing this does not seem to suffice for making an inference on one natural way of categorising cases. This is clearest when we consider how things might otherwise have gone. I might have noticed that the bird outside has a bright red face. I might then have judged that no other birds in this region have this feature and concluded that the bird is a goldfinch. In the latter case I *infer* that the bird is a goldfinch from that it has a bright red face and that no other birds in the region do. In doing so I engage in a minimal bit of conscious reasoning. In the former case, meanwhile, I may not engage in reasoning and may rather judge that the bird is a goldfinch outright—without considering the basis on which I judge. I simply exercise my ability to *recognise* goldfinches without engaging in any reasoning whatsoever. The two cases are both cases in which I consciously judge that there is a goldfinch and do so for reasons. A natural way to capture

---

<sup>11</sup> For the sake of simplicity I am assuming here that whenever one judges on the basis of some belief one judges for a reason (e.g. Hornsby 2008). In fact, it may well be right to see many cases in which one does so as cases in which one judges on the basis of *merely apparent* reasons. It may, for instance, be that to judge that  $q$  for the reason that  $p$  it must be a *fact* that  $p$  (e.g. Alvarez 2010). Genuinely judging on the basis of a reason may even require *knowledge* of the basis on which one judges (e.g. Littlejohn forthcoming; Williamson 2000).

the potential difference between the two cases is to say that it is only in the latter case that I clearly *infer* that the bird is a goldfinch. In the latter case I judge that the bird is a goldfinch by inferring that it is from prior grounds, whereas in the former case I may merely judge that it is a goldfinch outright, which is not to say that I do not do so for reasons. Distinguishing between merely judging that  $q$  where  $p$  is a reason of mine for doing so and doing so *by inferring*  $q$  from  $p$  allows us to distinguish between cases in which one merely judges for reasons and cases in which one judges on the basis of explicitly considered reasons and qualifies as engaging in conscious reasoning accordingly. It is the latter and not the former which is of concern here given that the aim is to understand inference of the sort which plays a central role in our conscious reasoning.<sup>12</sup>

If we count the case in which one merely exercises one's ability to recognise goldfinches as one in which an inference occurs then we may face counterexamples when it comes to giving an account of the phenomenon exemplified by the case in which one concludes that the bird is a goldfinch by engaging in reasoning. We might, for instance, find ourselves saying things like 'Inference must not essentially involve  $x$  because there need be no such thing in the former goldfinch case'. We will then be led to an account of inference which may leave out crucial features of the phenomenon exemplified by the latter case. This is precisely what I aim to show often happens in the recent literature on inference. My aim is to give an account of inference where this is understood as a kind of conscious act exemplified by the latter and not the former goldfinch case. If we ignore what is distinctive about this phenomenon by endorsing a more permissive conception of inference then we will miss out on an account of a conscious act where this account is apt to yield an explanation of how we can exercise doxastic agency. I do not doubt that much good work is done whilst operating with the more permissive conception of conscious inference. The suggestion is only that restricting ourselves to the more demanding conception of inference on which judging for reasons does not suffice for inference will yield an account of inference apt to explain how

---

<sup>12</sup> Mere judgment also plays a crucial role in our conscious reasoning, but a distinct role to that played by inference.

doxastic agency is possible. Since the ultimate aim is to understand how we can exercise doxastic agency by consciously reasoning it is inference of this sort which will be of concern in what follows.

It is worth noting at this point that authors interested in a personal-level phenomenon which they label ‘inference’ can appear to not be interested in the same phenomenon which is of concern in the discussion of inference to follow. Koziolok (2017), for instance, has it that an inference occurs iff one comes to believe something on the basis of something else which one believes. There is thus a respect in which his concerns are broader than my own. For Koziolok it must be that all cases in which one comes to believe on the basis of prior beliefs are cases of inference rather than just those in which one does so *by consciously inferring*. There is also a respect in which Koziolok’s concerns are narrower than my own. Koziolok cannot count cases in which one infers from mere suppositions or from assumptions made for practical purposes as cases of inference.

Harman (1986), meanwhile, is interested in what he calls “reasoned change in view”—a kind of attitude revision. And Boghossian claims to share his topic. He states that by inference he means

the sort of “reasoned change in view” that Harman (1986) discusses, in which you start off with some beliefs and then ... end up either adding some new beliefs, or giving up some old beliefs, or both. (2014, p. 2)

Boghossian, like Koziolok, is thus committed to having it that inference need not be conscious.

That said, both Boghossian and Koziolok explicitly state that their concern is a *conscious* phenomenon. Boghossian, for instance, states that the phenomenon of concern for him is “person-level, conscious and voluntary” (2014, pp. 2–3). Neither Boghossian nor Koziolok is consistent in their characterisation of what phenomenon they mean by ‘inference’. At times it is characterised such that what is of concern must be a potentially non-conscious phenomenon, at other times they are explicit that inference of the sort which concerns them is conscious. Surveying the recent literature on inference as above suggests that it is often simply not made clear precisely what phenomenon is of concern. It is, for instance, frequently unclear whether the phenomenon in question is

supposed to be conscious or whether it can occur unconsciously and likewise unclear whether it requires the occurrence of appropriate belief fixation.

Again, to be clear, I am not suggesting that inference of the sort of concern is necessarily what we mean whenever we speak of inferences nor that there is not good work in which ‘inference’ is used to refer to phenomena other than that which is of concern here. What I am warning against is taking conclusions drawn from discussion in which ‘inference’ is used in a permissive sense—say to refer to sub-personal processes (e.g. Johnson-Laird 2008) or belief fixation of the kind I claimed may not amount to inference (e.g. Parrott forthcoming)—and drawing conclusions about the conscious phenomenon of interest. The criticisms of accounts of inference to follow thus need to be understood as potentially restricted in the following way. I will argue that the accounts considered are inadequate insofar as they are intended to either just be or to yield accounts of the phenomenon of concern here. The accounts to be considered are often presented in just that way.

### **1.5. The possibility of doxastic agency**

I have suggested that a proper account of inference of the sort of concern will shed light on how we are capable of exercising agency over what we believe and can be responsible for what we believe accordingly—on how we are capable of “doxastic agency”. I am not the first to suggest that a proper account of inference will be apt to play such a role. Boghossian (2014) does so. But I will go beyond what Boghossian says in explaining *how* this is so—I will explain how a proper account of inference allows us to see how we can exercise doxastic agency by *self-consciously reasoning*.

It has not yet been made clear what doxastic agency is. The account of doxastic agency to follow will yield a clear answer, but for now we can say the following. To exercise doxastic agency is to exercise agency over one’s beliefs. And it is to do so in a “direct way” (Boyle 2009; Lee forthcoming; see also Moran 2001, pp. 116–20). We can bring about beliefs and changes in our beliefs indirectly, via a kind of self-manipulation. Knowing how forgetful I am in mornings, for instance, and worried about being late for a train I might set my watch half an hour fast. By doing so I may bring it about that at 8 the next

morning I believe that it is 8:30 a.m. I will leave the house in good time as a result. I can, that is, deliberately *cause* myself to have certain beliefs and in this indirect way exercise agency over my beliefs. However, the thought is, it is natural to see ourselves as capable of exercising agency over our beliefs in a more direct way than this—of being capable of being active with respect to what we believe without having to self-manipulate. As Hieronymi (2014) puts it, we can “settle questions” for ourselves and thus are capable of exercising agency over our beliefs in a way which does not amount to mere self-manipulation. To exercise such agency over one’s belief is to exercise doxastic agency.

It might be wondered why there is any difficulty in understanding how we are capable of exercising doxastic agency. In fact, however, many are sceptical about the very possibility of doing so (e.g. Alston 1988; Kornblith 2012; Lee forthcoming; Owens 2000; Setiya 2013; Valaris 2016b). This is because the suggestion that we are capable of agency over belief of the direct kind in question faces a distinctive challenge. The challenge itself stems from the fact that there are distinctive reasons to think that we cannot believe or come to believe “at will” (Williams 1970). These reasons are distinctive in that analogous considerations do not present a challenge to the suggestion that the likes of raising one’s arm, thinking about whether *p*, supposing that *p*, and even deciding to  $\varphi$  can all be done at will. One cannot believe or come to believe at will, as one can raise one’s arm and think about whether *p* at will, because “beliefs aim at truth” (Williams 1970, p. 136). Williams’ suggestion that belief aims at truth involves the claim that belief is correct only if true, that to express belief is to express what one takes to be the truth, and, most crucially, that “to believe that *p* is to believe that *p* is true” (1970, p. 137).

There is controversy when it comes to how to correctly characterise the thought that belief aims at truth (e.g. Shah and Velleman 2005; Mchugh 2011a, sec. 1). Likewise, there is controversy when it comes to explaining why this should be taken to prevent us from being capable of believing or of coming to believe at will (e.g. Bennett 1990; Hieronymi 2006). The key thought behind

Williams' suggestion appears to be as follows.<sup>13</sup> To believe that  $p$  is to believe what one takes to be the truth about whether  $p$ . Crucially, whether  $p$  is true is a matter of how things are in the world and is thus not up to oneself.<sup>14</sup> Accordingly, when one believes  $p$  one regards  $p$  as true where one does not regard the constraint of treating  $p$  as true as self-imposed (unlike when one is merely supposing or assuming that  $p$ , or when one's guess is that  $p$ ) (Soteriou 2013, Chapter 11). Given this, to take an attitude to  $p$  where one sees the constraint of treating  $p$  as true as self-imposed, such that one's doing so can really be up to oneself, cannot be to believe  $p$  (c.f. O'Shaughnessy 2008a, pp. 60–7; see also Mchugh 2011a, sec. 1). We cannot believe or come to believe at will accordingly. To believe or to come to believe  $p$  at will one would have to do so where one's doing so is up to oneself and thus such that one does regard the constraint of treating  $p$  as true as self-imposed.

There is no parallel difficulty when it comes to seeing the likes of arm raisings and even decisions as being performed at will. There is rather distinctive pressure to see us as unable to believe or come to believe at will since belief aims at truth. When one raises one's arm at will, for instance, one's doing so is up to oneself. But the matter of whether to raise one's arm, unlike the matter of what is true, is up to oneself. Similarly, supposing one decided to  $\varphi$  at will one's doing so does not seem to be prevented from being up to oneself as believing  $p$  is. The matter of whether to  $\varphi$  when one decides to  $\varphi$ , unlike the matter of whether  $p$  is true when one believes  $p$ , can be up to oneself.<sup>15</sup>

---

<sup>13</sup> Williams in fact trades on the claims that (i) being able to believe at will would require knowing oneself to be capable of doing so, and (ii) one could not believe at will and not know oneself to be doing so. It is far from clear that Williams' precise line of thought in fact leads to anything as strong as the conclusion which he draws.

<sup>14</sup> There may seem to be countless exceptions. If  $p$  is substituted for 'the lights are on', for instance, then the matter of whether  $p$  may be up to oneself since one can turn them on or off. But such cases are cases in which one can make it the case that  $p$  and *by doing that* make it that case that one believes  $p$ . And doing that is not believing at will.

<sup>15</sup> I set aside the matter of whether we can in fact decide at will. It appears to be orthodox to think that we can and that decisions are or at least can be actions (see Pink 2003, p. 255; Shepherd 2015, p. 335). Others doubt that decisions are really actions and doubt that they need to be in order to be apt to play their role in deciding for ourselves (e.g. O'Shaughnessy 2008b, pp. 543–7; Strawson 2003, p. 244).

As noted, in what way belief can be said to “aim at truth” and precisely how this prevents us from being able to believe or come to believe at will is a matter of controversy. Nonetheless, as Boyle notes, although

[t]he explanation of the impossibility of believing "at will," and the precise nature of this constraint, are matters of controversy ... that there is a truth here that needs explaining is not terribly controversial. (Boyle 2012, pp. 119–20 note 4)

There are, it seems, distinctive reasons to think that we cannot believe or come to believe at will. We cannot do so because belief aims at truth and because what is true is not, in the required way, up to us. As Boyle (2009, p. 120) also notes, this yields a distinctive challenge for defenders of the possibility of agency over belief of the direct kind in question (see also Boyle 2011a, pp. 3–4; Kornblith 2012, p. 85; McHugh 2017, sec. 1.1). If we really can exercise doxastic agency, the thought is, it needs to be explained how this is so despite the fact that we can neither believe nor form particular beliefs at will. How can we exercise doxastic agency despite that we can neither believe nor come to believe at will? Boyle assumes that in order to do so we must appeal to some “other”, *special* kind of agency or control—that we must appeal to agency or control which is different in kind to that in question when one, say, raises one's arm (see also Moran 2001, pp. 114–5). As he puts it

defenders of the application of agential notions to belief ... must give a clear account of what *other* notion of agency or control is at issue here. (2009, p. 120)

The challenge, in sum, is that if we are to take ourselves to be capable of exercising doxastic agency then it must be explained how this is so despite that we cannot believe or come to believe at will. How are we capable of being agents with respect to our beliefs in any way other than via self-manipulation given that we cannot believe or come to believe at will?<sup>16</sup> Some have it that this challenge should push us to accept a view on which agency is connected to belief only in that beliefs can be the effects of something agential. We are best

---

<sup>16</sup> In presenting the challenge I have focused on the case in which doxastic agency is allegedly exercised in believing or coming to believe. A parallel challenge casts doubt on the possibility of exercising doxastic agency in revising one's beliefs.

capable of actions of a sort which can lead causally to beliefs via “mental ballistics” (Strawson 2003, p. 241) (see sec. 5.4 below). As advocates of this suggestion acknowledge, to accept such a picture would be to give up on the idea that we are capable of exercising doxastic agency. To accept the ballistics picture is to accept that we are misguided if we take ourselves to be capable of exercising doxastic agency.

Some allege that there are counterexamples to Williams’ claim that we cannot believe or come to believe at will. Ginet, for instance, presents the following case:

Before Sam left for his office this morning, Sue asked him to bring from his office, when he comes back, a particular book that she needs to use in preparing for her lecture the next day. Later Sue wonders whether Sam will remember to bring the book. She recalls that he has sometimes, though not often, forgotten such things, but, given the inconvenience of getting in touch with him and interrupting his work and the thought that her continuing to wonder whether he’ll remember it will make her anxious all day, she decides to stop fretting and believe that he will remember to bring it. (2001, p. 64)

Such cases, Ginet suggests, can genuinely amount to cases in which one decides to believe. Even if Ginet is right here, which is doubtful in light of the above reflections on why we cannot believe at will, appeal to the capacity to believe or to come to believe in such cases will not be apt to explain how we are capable of doxastic agency to the extent sought. Acknowledging that we are capable of doxastic agency in such cases, for instance, will not tell us how one can qualify as having exercised doxastic agency in a case in which one has figured out the answer to a puzzle (where once one has done so the answer is evident) or in a case in which one has found one’s keys (again, where once one has done so it is evident where they were). One does not and cannot believe or come to believe at will in cases of the latter sort. Yet such cases are paradigm cases of the alleged exercise of doxastic agency—paradigm cases of the sort where one appears to be or to have been active with respect to one’s beliefs in a way which can make one apt to be due credit or blame for the relevant beliefs.



When it comes to understanding how we are capable of doxastic agency it is typical to have it that our being self-conscious is crucial (Korsgaard 2009; Leite 2016; Moran 2001; O’Shaughnessy 2000; Rödl 2007; Soteriou 2013). O’Shaughnessy, for instance, suggests that we are “aware of our own minds” where this allows us to “rationally govern our existences” (2000, p. 112). He insists that our being self-conscious is part of what makes us capable of “[t]ranscending the condition of ‘animal-immersion’” and of “mental freedom” where doxastic agency is one kind of manifestation of such freedom (2000, pp. 111–2). The natural thought here is that to be capable of doxastic agency—to be a doxastic agent—one must be self-conscious. Furthermore, to exercise doxastic agency is to exploit one’s self-consciousness.

The typical contemporary version of this view is advanced by McGinn when he claims that

[i]f a person were not aware of his beliefs, then he could not be aware of their inconsistency; but awareness of inconsistency is (primarily) what allows normative considerations to get purchase on beliefs (1996, p. 22).

McGinn suggests that

the rational adjustment of beliefs one to another seems to involve self-consciousness, that is, knowledge of what you believe. (1996, p. 22)

Moran sums up the typical view as being one on which

[i]f some attitude of mine is out of reach of the light shone by consciousness, then I cannot “see” it in order to check it out and adjust it to the rest of my beliefs and desires. As with the intelligent, directed control of other things, the rational adjustment of one’s beliefs and other attitudes requires, at the very least, *awareness* of what is being subjected to control. (2001, p. 110)

The suggestion is that self-consciousness is required for and exploited in the exercise of doxastic agency because only the self-conscious can be aware of their beliefs. Furthermore, the thought is, one must be aware of one’s beliefs in order to exercise agency over them.

It is in fact far from clear why knowing one’s beliefs should be taken to make direct agency over one’s beliefs of the sort in question possible (Hieronymi 2014, pp. 4–8, 20–1; Owens 2000, pp. 81–7). Knowing what one

believes might make it possible to act in ways which lead causally to intended changes in those beliefs. But why should knowing what one believes make it possible for one to be active with respect to what one believes—to settle questions for oneself? Knowing what one believes, after all, does not make believing or coming to believe something which one can do at will. As Hieronymi notes, mere knowledge of something, such that one can reflect upon it, does not suffice for the capacity to exercise control over that thing (2014, p. 8). Nor does it suffice for the capacity to directly exercise agency over it.

In sum, I will argue that self-consciousness indeed is required for and exploited in the exercise of doxastic agency. But why this is so has been misunderstood. I will argue that we are capable of exercising doxastic agency because we are capable of *self-conscious reasoning*. This is so because self-conscious reasoning is a kind of *action*, because such reasoning *constitutively involves making inferences*, and because to infer can *just be* to come to believe what one infers. It is typical to think that such an answer to the above challenge cannot be correct (e.g. Boyle 2009, 2011a; Kornblith 2012; Lee forthcoming; Strawson 2003; Valaris 2014). However, a corrected account of inference, of its role in reasoning, and of its relation to belief will reveal how the case against the suggested picture is in fact misguided.

#### **1.6. The self-consciousness of inference as a constraint in what follows**

My aim is thus to show how, in the way suggested above, a proper account of inference yields an explanation of how doxastic agency is possible and of how self-conscious reasoning is a means to exercise it. Given this, I will proceed as follows when it comes to giving my account of inference and criticising alternatives. I will *assume* that inference is self-conscious. An adequate account of inference apt to play the above explanatory role must capture that inference is self-conscious. Furthermore, it must capture that inference is self-conscious and that this is not simply so because the inferences of the sort of concern are a subset of the inferences in a more permissive sense which *happen* to be self-conscious. The thought here is as follows. Appeal to inferences in a more permissive sense which one happens to know are occurring will not be apt to

explain how we are capable of doxastic agency despite being unable to believe or come to believe at will, nor will such appeal be apt to explain how to exercise such agency is to exploit one's self-consciousness. Suppose, for instance, that we used 'inference' so as to include events of belief fixation which one does not know are occurring. It seems clear that when such events occur without our knowledge they do not amount to exercises of agency. Furthermore, supposing that an inference in this sense occurs and one does happen to know that it is occurring it becomes no clearer why we might see the event as an exercise of agency, nor why we might think that it is itself an exploitation of one's self-consciousness. We thus should not appeal to inferences in such a permissive sense when it comes to explaining how we exercise doxastic agency. Inference of the sort appeal to which can illuminate (i) how we are capable of doxastic agency, and (ii) how to exercise such agency is to exploit our self-consciousness, is a distinctive kind of self-conscious occurrence. It is a kind of essentially self-conscious occurrence, rather than a subspecies of inference in some more permissive sense which need not occur self-consciously.

Imposing the above constraint on our account of inference will yield an account apt to explain how doxastic agency is possible as well as why to exercise doxastic agency is to exploit one's self-consciousness. The constraint will also help to reveal what is wrong with other recent attempts to give accounts of inference of the sort of concern. The constraint is thus motivated by its facilitating such explanatory work. It is to be vindicated by the explanatory work which the account of inference it yields is apt to be put to (see ch. 5 Below).

### **1.7. Overintellectualization**

Inference of the sort of concern, I have suggested, is a kind of conscious and self-conscious occurrence. I will ultimately argue that inference's being self-conscious requires that when one infers one believes that a relevant inferential relation obtains between that which one infers and that which one infers from (see sec. 4.8). It is common to object to accounts with such commitments on the grounds that they "overintellectualize" (e.g. Boghossian 2001, pp. 637–8, 2014, pp. 6–7; McHugh and Way 2016, p. 319; Taylor 2015, p. 2991; Thurow 2008, p. 279). On the account I will offer to infer  $q$  from  $p$  one must know that

one is inferring and believe that a relevant inferential relation obtains between  $p$  and  $q$ . However, the objection goes, it seems clear that animals and young children lack the conceptual sophistication needed to have the required knowledge and beliefs. Furthermore, the claim that animals and young children can make inferences is well supported, both empirically and by everyday observation and interaction. In standard inference by exclusion tasks, for instance, subjects need to “infer” from the absence of a cue that another location should be investigated to find a reward. Primates of various species as well as young children (19 to 25 month) succeed in such tasks (e.g. Call 2001, 2004; Petit et al. 2015).

To object to the claim that inference is self-conscious and to the account of inference that results in the above way is to equivocate on the meaning of ‘inference’. Inference of the sort of concern here is conscious inference of the sort which is apt to occur as a constituent of self-conscious and active reasoning. Inference, understood as a kind of constituent of such reasoning, will be appealed to when it comes to explaining how we are capable of exercising doxastic agency by reasoning. The claim that animals and young children can make inferences of this sort, rather than inferences in a more permissive sense, is not well supported. I will argue that with inference understood as a kind of self-conscious occurrence we can see how we can exercise doxastic agency by reasoning and how to do so is to exploit one’s self-consciousness. Being self-conscious in the required way will turn out to require one to also be capable of forming beliefs about relevant inferential relations (see sec. 4.8 below). The claim that such conceptual sophistication is required in order to be capable of making inferences and of exercising doxastic agency by reasoning is not in conflict with the well supported claim that animals and young children are capable of making inferences in a more permissive sense.

### **1.8. An ontological quandary**

We began with Boghossian’s (2014) question ‘What is an inference?’, the ultimate aim being to explain how we are capable of doxastic agency. The question is understood as concerning the nature of a distinctive kind of conscious and self-conscious occurrence. A full answer will be informative

when it comes to the ontology of inference. A full answer, that is, will tell us what kind of occurrence an inference is. When one infers  $q$  from  $p$  the occurrence which is one's inference might be alleged to be identical to an event of belief fixation. It might instead be alleged to be identical to an occurrence involving acceptance of  $p$  and then acceptance of  $q$ —to what we might call the *process* beginning with one's acceptance of  $p$  and terminating with one's acceptance of  $q$ . Or it might be alleged to be identical to the occurrence which is one's accepting  $q$ —to the conscious *event* which is one's acceptance of  $q$ . However, as will be detailed below, if we try to answer this ontological question directly each possible answer appears problematic. We are left in a quandary when it comes to saying what one's inference is when one infers  $q$  from  $p$ . In light of this difficulty I suggest that our question should be 'What is it for an inference to occur?' The latter question can be answered without first saying what occurrence ought to be identified with one's inference when one infers  $q$  from  $p$ . Furthermore, with the latter question answered it will become clear what kind of occurrence inference is in a way which resolves the apparent difficulties for the resultant answer to the above ontological question.

One natural suggestion when it comes to saying what kind of occurrence inference is would be to say that inference is a kind of belief fixation. After all, the thought might go, how could inference play a central role in reasoning which is itself a means to exercise doxastic agency if it is not a kind of belief fixation? In the above Sudoku case, for instance, I come to believe that the answer is 5 through reasoning by inferring that the answer is 5. How could inference be a means to do so and play a role in doxastic agency in turn if it is not a kind of belief fixation?

It has already become clear why this suggestion cannot be correct. Inference cannot merely be a kind of belief fixation. Inference can occur without the occurrence of belief fixation of the sort which inference can yield. The above examples of inference from assumptions and suppositions make this clear.

Matthew Parrott suggested a response on behalf of my opponent here.<sup>17</sup> It could be claimed that inference is a kind of belief fixation and that suppositional

---

<sup>17</sup> It is not a response which he actually endorsed.

inference and inference from assumptions are kinds of *pretend* inference. To infer under the scope of a supposition is really to pretend to infer proper, where inference proper is a kind of belief fixation. An analogous suggestion could be made when it comes to inference from assumptions. But the suggestion cannot be correct. Suppositional inference is inference proper. As Anscombe remarks:

if you say ‘suppose  $p$ , and suppose  $q$ , then  $r$ ’; or if, being given ‘ $p$ ’, you say: ‘suppose  $q$ , then  $r$ ’; you are just as much inferring, and essentially in the same way, as if you are given ‘ $p$ ’ and ‘ $q$ ’ as true and say ‘therefore  $r$ ’.  
(1971, p. 116)

Broome (2013, p. 265) observes that this must be correct as it is only by making suppositional inferences in accordance with the method of conditional proof that we can come to know conditional propositions, as in the case of explosion reasoning above. We in fact acquire such knowledge by reasoning suppositionally where this involves making suppositional inferences. To deny that these are real inferences is simply to change the subject.

Inference does not require the occurrence of belief fixation of the sort which inference can yield. It does, however, require the occurrence of conscious acceptance where, as noted (see sec. 1.4 above), by ‘acceptance’ I mean a kind of event and conscious act where one can accept  $p$  without judging that  $p$ . Thus understood the likes of judging, supposing and guessing are all kinds of conscious acts of acceptance. With acceptance thus understood it is clear that to infer  $q$  from  $p$  one must consciously accept  $q$ . Furthermore, one must also first consciously accept  $p$  (Boghossian 2015, p. 43; Broome 2013, p. 222). As seen, if we deny this then we will fail to distinguish inference from mere judgment for reasons. Judging that  $q$  where believed  $p$  is one’s reason does not suffice for the occurrence of an inference from  $p$  to  $q$  of the sort of concern. To infer  $q$  from  $p$  one must *first* consciously accept  $p$  and *then* consciously accept  $q$ . Suppose, for instance, that I believe  $p$  and infer  $q$  from it. To do so I must first believe  $p$ . Furthermore, that  $p$  is the case (I take it at least) must occur to me—I must in this way consciously accept  $p$ . Finally, an inference from  $p$  to  $q$  will not have occurred unless I consciously accept  $q$ . There must be a conscious occurrence with content  $q$  for me to qualify as having inferred  $q$ . The same holds true in the supposition case. I might, for instance, suppose that  $p$  and infer  $q$

from it. To do so I must first suppose that  $p$  and in this way consciously accept  $p$ . Again, an inference from  $p$  to  $q$  will not have occurred unless I consciously accept  $q$ . I must suppose that  $p$  and then accept  $q$  “in the same context of supposition” (Wright 2014, p. 29). That is not to say that in order to infer  $q$  from  $p$  one must accept  $p$  in a *distinctive way*. One might simply notice that  $p$  and then  $q$  just occur to one in such a way that one qualifies as having inferred  $q$  from  $p$ . One need not, for instance, accept  $p$  *in order to infer*. To infer  $q$  from  $p$ , that is, one must first accept  $p$  and then accept  $q$ , but one can accept  $p$  in any old way, rather than needing to do so in a way distinctive to cases in which one infers from  $p$ .

Although to infer  $q$  from  $p$  one must first accept  $p$  and then accept  $q$  it is clear that doing so does not suffice for the occurrence of an inference. One might judge that  $p$  and then unrelatedly judge that  $q$ . Clearly one will not thereby have inferred  $q$  from  $p$ . This observation naturally prompts the suggestion that for an inference from  $p$  to  $q$  to occur one must accept  $p$  and then accept  $q$  where the former event *causes* the latter event. This in turn makes the following the most common proposal when it comes to the ontology of inference. To infer  $q$  from  $p$  one must accept  $p$  and then accept  $q$  where one’s acceptance of  $p$  causes one’s acceptance of  $q$ . With this said one’s inference is then taken to have both the event which one’s acceptance of  $p$  and the event which is one’s acceptance of  $q$  as constituents. One’s inference from  $p$  to  $q$ , on this view, is an occurrence which has as constituents the event which is one’s acceptance of  $p$  and the event which is one’s acceptance of  $q$  and which the former event causes. Boghossian (2014, p. 5) seems to be endorsing this view when he suggests that inference is a kind of “causal process”, whilst it is the norm to be at least implicitly committed to the view (e.g. Brewer 1995; Broome 2013, 2014; Chudnoff 2014; Fumerton 2004; Huemer 2016; Koziolok 2017; Siegel forthcoming). Following Boghossian, it is natural to describe the view as one on which inference is a *process*. When one infers  $q$  from  $p$  one’s inference is the process beginning with one’s acceptance of  $p$  and terminating with one’s acceptance of  $q$ .<sup>18</sup>

---

<sup>18</sup> I say that inference is a process, on this view, only so that we have a succinct way of discussing the proposal. I do not want to take a stand on any issues concerning the metaphysics of

The suggestion now in question is that when one infers  $q$  from  $p$  one's acceptance of  $p$  appropriately causes one's acceptance of  $q$  and one's inference is the occurrence which has these events as constituents—what we might call the process beginning with one's acceptance of  $p$  and terminating with one's acceptance of  $q$ . Again, this proposal concerning the ontology of inference appears to be problematic. The suggestion that when one infers  $q$  from  $p$  one's inference is the process which begins with one's acceptance of  $p$  and terminates with one's acceptance of  $q$  seems to entail that inference is *interruptible*. Consider a case in which I infer  $q$  from  $p$ . The occurrence which is my inference, on the proposal in question, begins when I accept  $p$  and is not over until I have accepted  $q$ . That seems to imply that had I been interrupted after having accepted  $p$  but before accepting  $q$  I would thereby have been interrupted *whilst* inferring  $q$  from  $p$ . However, as several authors have observed, it seems clear that we cannot be interrupted *whilst* inferring (Hlobil 2016, p. 4; Ryle 2009a, pp. 275–6; White 1971, p. 289). One cannot be interrupted whilst inferring  $q$  from  $p$  having begun to do so.

To be clear, I do not intend for the above problem for the suggestion that inference is a process to be decisive or even convincing. The point I want to make is just that if we try to answer the question of what kind of occurrence inference is directly then each potential answer seems problematic in its own way. I will recommend that we begin with a different question accordingly.

We have considered the suggestion that inference is a kind of belief fixation and seen that it cannot be correct. The suggestion that inference is a process also seems problematic. This might lead us to the suggestion that when one infers  $q$  from  $p$  one's inference is instead the *event* which is one's acceptance of  $q$ . To infer  $q$  from  $p$ , as seen, one must accept  $p$  and then accept  $q$ . Furthermore, one has not inferred  $q$  from  $p$  until one has accepted  $q$ . If we say that the occurrence which is one's inference begins when one accepts  $p$  we face the above difficulty. We might thus claim instead that when one infers  $q$  from  $p$  one's inference is the conscious act of acceptance and conscious event which is

---

processes nor on how the event/process distinction ought to be understood (see e.g. Crowther 2011; O'Shaughnessy 2000, pp. 42–9; Steward 2013; Stout 1997).



one's accepting  $q$ . This new suggestion, however, itself appears to be problematic. It is hard to see how it could be correct given that, as seen, the occurrence of a single inference requires the occurrence of *two* events. To infer  $q$  from  $p$  two events must occur—one must first accept  $p$  and then accept  $q$ . We might then ask 'How could one's inference from  $p$  to  $q$  be identical to the single event which is one's acceptance of  $q$  and yet *two* events of acceptance be required in order for this inference to occur?' If one's inference *just is* the single occurrence which is one's acceptance of  $q$  then surely the occurrence of this single event *suffices* for the occurrence of an inference. Yet, as seen, the occurrence of this inference requires the occurrence of two events of acceptance.

It seems that we are left in a quandary if we ask 'What is an inference?' where an answer to this question requires saying which occurrence one's inference is identical to when one infers  $q$  from  $p$ . Each potential answer appears problematic. Again, to be clear, none of the above problems are intended to be decisive. The aim is only to reveal that if we aim to directly answer the above ontological question then there is no answer which is obviously unproblematic. I thus suggest that we should begin with a different question. Our guiding question should concern what it is for an inference to occur. We should, that is, ask 'What is it for an inference to occur?' Crucially, this question can be answered whilst remaining neutral when it comes to the question of which occurrence one's inference ought to be identified with when one infers. Furthermore, an answer to the latter question and a resolution of our ontological difficulty will fall out of a correct answer to the former question—out of an answer to the question 'What is it for an inference to occur?'

### **1.9. Theoretical and practical inference**

Because my ultimate concern is inference's role in doxastic agency—i.e. its role in the exercise of agency over *belief*—my focus has been, and will continue to be, *theoretical* inference. It is worth noting, if only to set it aside, that there are also *practical* inferences. I might, for example, intend to go to Oxford tonight. On reflection I might realise that the only affordable way to do so is to go by bus and thus form the intention to go by bus. Doing so may involve making a

practical inference. The contrast here can be drawn by saying that theoretical inferences are those which can have beliefs as their conclusions, whilst practical inferences are those which can have intentions as their conclusions (Brandom 1994, p. 245).

A thorough discussion of practical as well as theoretical inference would require engagement in too many controversies. To take one example, there is controversy regarding what form practical inferences actually take and the role of desire in practical inference. Brandom claims I could infer as follows.

Only opening my umbrella will keep me dry,  
so I shall open my umbrella. (1994, p. 245)

In doing so I manifest my desire to stay dry. Davidson (2001a) disagrees. According to him the desire to stay dry would have to actually feature in such an inference. There are other controversies here too, such as that regarding whether only mere decisions or intentions are the conclusions of practical inferences (e.g. Broome 2013, p. 250) or whether actions themselves can be the conclusions of practical inferences (e.g. Dancy 2014). I will continue to focus on theoretical inference. The implications when it comes to giving an account of practical inference will be set aside.

### **1.10. Conclusion**

The aim of this chapter has been to make clear what my target phenomenon is and how I aim to approach the task of giving an account of it. Of concern is conscious theoretical inference. Inference of this sort, I suggest, is a distinctive kind of self-conscious occurrence. An adequate account of inference thus understood, I will argue, will ultimately pave the way for an explanation of how we can exercise doxastic agency by actively reasoning. It will likewise pave the way for explanation of why self-consciousness is required in order to be capable of doxastic agency and of how to exercise such agency is to exploit one's self-consciousness. I have also suggested that if we try to answer the question 'What is an inference?' directly we face an ontological quandary. When it comes to saying which occurrence one's inference is identical to when one makes an inference there is no answer which is clearly unproblematic. Instead I thus

suggest that we ask 'What is it for an inference to occur?', this question's being neutral on which occurrence one's inference is identical to when one infers.

## 2. Causalist accounts of inference

**Chapter abstract** My aim in this chapter is to assess the orthodox conception of inference. On what I will label the ‘causalist’ conception of inference for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  in some specifiable “right way”. Advocates of the “Taking Condition” have it that this requires that one’s taking it that  $q$  follows from  $p$  play an appropriate causal role in one’s coming to accept  $q$ . The Taking Condition’s having been the focus of much of the recent literature on inference, I consider its motivation and the objections its advocates face, finding the objections wanting. A simple and seemingly adequate reply, for instance, is available to the common suggestion that endorsing the Taking Condition leads to a regress of the sort to be found in Carroll (1895). I go on to argue that the causalist conception of inference allows that inference can fail to be self-conscious and thus fails to capture inference of the sort of concern. In response the self-consciousness of inference might be built into the causalist conception. But to do so, I argue, is to undermine the causalist conception’s motivation. The conception is either inadequate or unmotivated. A distinct response would be to suggest that the causalist conception is motivated insofar as the Taking Condition is, the Taking Condition’s yielding a version of the causalist conception. I thus go on to reveal why the most promising route to motivating the condition fails. It might be thought that we must have it that whenever one infers  $q$  from  $p$  one’s taking it that  $q$  follows from  $p$  must play an appropriate causal role in bringing about one’s acceptance of  $q$  such that we can capture the respect in which one’s making the inference must be intelligible from one’s point of view. On reflection it can be seen that the Taking Condition cannot be motivated by the apparent need to capture the way in which one’s inferences are intelligible to oneself in this way. If we are to arrive at an adequate account of inference of the sort apt to help explain how we are capable of doxastic agency then an alternative to the causalist conception is required.

## 2.1. The causalist conception of inference

The question of concern is ‘What is it for an inference to occur?’ If we ask this instead of asking ‘What is inference?’ then we can remain neutral on the ontology of inference. Our answer to the latter question can then be informed by our answer to the former. So far, the following has become clear when it comes to saying what it is for an inference from  $p$  to  $q$  to occur. One must first consciously accept  $p$  and then consciously accept  $q$ , where  $p$  is that which one infers from and where  $q$  is that which one infers. To do so one need not believe or come to believe either  $p$  or  $q$ . In accordance with this the standard reaction to the question at hand will be to endorse what I will label the ‘causalist conception of inference’. According to advocates of this conception to infer  $q$  from  $p$  one must accept  $p$  and then accept  $q$  where one’s acceptance of  $p$  causes one’s acceptance of  $q$ . But one’s acceptance of  $p$  must cause one’s acceptance of  $q$  in a specifiable “right way” if an inference is to qualify as having occurred. For an inference from  $p$  to  $q$  to occur, on this conception, is just for acceptance of  $p$  to cause acceptance of  $q$  in this right way. The task of saying what it is for an inference to occur then becomes that of specifying what the right kind of causation is. To say what it is for an inference from  $p$  to  $q$  to occur one just needs to specify the right way in which acceptance of  $p$  would need to cause acceptance of  $q$  such that an inference qualifies as having occurred, where in doing so one may, but equally may not have reductive ambitions.

The causalist conception is by a large margin the most commonly accepted conception of inference. It is accepted, if not always explicitly, by Boghossian (2014), Brewer (1995), Broome (2013, 2014), Chudnoff (2014), Fumerton (2015), Huemer (2016), Kornblith (2012), Koziol (2017) and Siegel (forthcoming), to name a few. To be precise, it is worth noting that some of those labelled causalists here in fact operate with a more permissive conception of inference than that in use here, allowing that inference can occur and conscious acceptance not. Brewer (1995) and Kornblith (2012), for instance, allow that an inference from  $p$  to  $q$  can occur and conscious acceptance not if belief in  $p$  appropriately causes belief in  $q$ . As has been noted, the focus here being *conscious* inference of the sort apt to occur as a constituent of conscious

reasoning, inferences in this more permissive sense are not of concern. However, advocates of causalist accounts of inference of the latter sort often do take it that their accounts are apt to shed light on inference of the sort of concern here. They can be read as taking their accounts to be such that they can be naturally extended to yield accounts of conscious inference of the sort which concerns us. On the accounts in question inference in the permissive sense can occur when belief in  $p$  causes belief in  $q$  in the right way, where this need not involve conscious acceptance. But inference in the sense of concern here can simply then be said to occur when conscious acceptance of  $p$  causes conscious acceptance of  $q$  in the right way, where what this right way is is specified in the same way as before. Inference of the sort of concern here, on the accounts which result, is just the conscious form of a broader phenomenon. And an account of inference of the sort of concern falls out of the account of the broader phenomenon.

Given this, my criticism of the causalist conception of inference to follow needs to be read as restricted in the following way. Causalist accounts of inference are either inadequate or unmotivated *insofar as they are intended to just be, or to be apt to yield, accounts of inference of the sort of concern here*. Our aim is an account of conscious inference of the sort apt to occur as a constituent of conscious reasoning. Inferences of the relevant kind are not merely the members of a subset of the occurrences which qualify as inferences in a more permissive sense and need not be self-conscious. They are not just the inferences in a more permissive sense which happen to be self-conscious. Inferences of the relevant kind are rather self-conscious occurrence of a distinctive kind. An adequate account of inference thus understood will be apt to help explain how it is that we are capable of exercising doxastic agency by reasoning. My aim in this chapter is to argue that if we want such an account then we should seek an alternative to the causalist conception of inference.

## **2.2. The causalist conception's motivation**

With the causalist conception of inference clear we can consider the conception's motivation. As Snowdon (1980, p. 175) notes, causalist accounts of various phenomena are given. Grice (1961), for instance, endorses a causalist

account of *seeing*. For one to see an object *o*, on such an account, is for one to be in a state which has been appropriately caused by *o* such that things look to one some way. Davidson (2001a) meanwhile endorses a causalist account of *acting for a reason*. To  $\varphi$  where in doing so one is acting for a reason, on such an account, is for one to  $\varphi$  where one's  $\varphi$ -ing is appropriately caused by one's reason (one's reason being conceived of as a mental state or complex thereof).

As Snowdon also notes, the arguments offered in favour of causalist accounts also tend to have a common form. They begin with consideration of cases in which we have much of what is required for the phenomenon of concern to be exemplified and yet in which the phenomenon is not exemplified. It is then noted that a causal connection is missing in the relevant case and concluded that part of what it is for the phenomenon of concern to be exemplified is for such a causal connection to obtain. Grice (1961, p. 142), for instance, discusses a case in which it looks to you as if there is a pillar in front of you. There in fact is a pillar in front of you, but this appears to be so only because a mirror is blocking your view of it and reflecting a different pillar. The case is one in which it looks to you as if there is a pillar in front of you and there is in fact is one, but not a case in which you see the pillar. It is also, we can suppose, a case in which your perceptual state is not caused by the pillar before you. Grice concludes that part of what it is to see the pillar before you is for it to be causing your perceptual state. In general, part of what it is to see *o* is for one to be in a perceptual state caused by *o*.<sup>19</sup> Snowdon (1980, pp. 181–2) observes that it does not follow from the fact that a relevant causal connection is absent in cases in which one fails to see *o* that part of what it is to see *o* is for such a causal connection to obtain. The advocate of the causalist account rather trades on an inference to the best explanation. We might ask *why* cases such as Grice's mirror case fail to amount to cases of seeing the object before one. The best explanation, it is alleged, is that in the relevant cases an appropriate causal connection between the object and oneself is missing and that part of what it is to see such an object is for such a causal connection to obtain.

---

<sup>19</sup> Putting the debate in these terms amounts to something of a reconstruction. Snowdon and Grice are concerned with the *conceptual analysis* of seeing, rather than with *what it is* to see.

A precisely analogous move is implicitly made in the inference case (e.g. Boghossian 2014, p. 3; Huemer 2016, p. 151). As seen, to infer  $q$  from  $p$  one must accept  $p$  and then accept  $q$ . But doing so does not suffice for inferring  $q$  from  $p$ . One might, for instance, judge that  $p$  and then incidentally later judge that  $q$ . To do so is not to infer  $q$  from  $p$ . Furthermore, it might be thought, a causal connection between one's acceptance of  $p$  and one's acceptance of  $q$  will likely be missing in such a case. On the basis of such observations it is concluded that part of what it is to infer  $q$  from  $p$  is for acceptance of  $p$  to cause acceptance of  $q$ . Again, it does not follow from the fact that a causal connection is typically missing in cases in which one fails to infer  $q$  from  $p$  that such a connection's obtaining is part of what it is for one to do so. It is rather assumed that the best explanation of such cases' failure to constitute cases in which one infers is that a causal connection is missing and that such a connection is required in order for there to be an inference, since the obtaining of such a causal connection is part of what it is for an inference to occur.<sup>20</sup>

To arrive at the conclusion that part of what it is to infer  $q$  from  $p$  is for acceptance of  $p$  to cause acceptance of  $q$  is not yet to arrive at the causalist

---

<sup>20</sup> In fact, I am being concessive to the causalist here in granting them the claim that in cases in which one accepts  $p$  and then incidentally later accepts  $q$  it will typically be false to say that one's acceptance of  $p$  causes one's acceptance of  $q$ . Many will have it that on reflection it is clear that most and perhaps even all cases in which one accepts  $p$  and later  $q$  are cases in which one's acceptance of  $p$  causes one's acceptance of  $q$ . Effects, after all, tend to have many and various causes. Suppose, to borrow an example used by Steward (1997, p. 190), that Mary dies because she allows her aluminium ladder to touch some power lines. We can correctly say that Mary's death was caused by an electrical discharge. But we can likewise correctly say that it was caused by the power line, by the ladder, by the ladder's being a good conductor and by Mary's decision this morning to take her ladder. Her decision to take her ladder will likewise have prior causes which in turn will qualify as causes of her death. And those prior causes will themselves have had prior causes which thereby qualify as causes of Mary's death in turn, and so on. To return to the case at hand, when one consciously accepts  $p$  and then later consciously accepts  $q$  it seems, on reflection, that it will at least likely be the case that one's acceptance of  $p$  is a cause of one's acceptance of  $q$ , albeit not a "proximate cause", simply because the conscious occurrences which make up one's stream of consciousness are part of what causally determine how one's future turns out and thus are part of what cause one's future conscious acts. It is perhaps worth noting here that one could in fact accept the claim that in order to infer  $q$  from  $p$  one must accept  $p$  and then  $q$  where one's acceptance of  $q$  is caused by one's acceptance of  $p$  without accepting the causal conception of inference. One might accept such a claim on the basis of the likes of the preceding reflections and reject anything which would amount to a causal *conception* of inference. Analogously, one might think that it is necessary for  $a$ 's being married to  $b$  that  $b$  have been causally involved in  $a$ 's coming to be married. But that would not force one to accept a causal conception of marriage—to endorse an account of what it is to be married on which it is characterised in terms of the causal history of one's current state (Snowdon 1980, pp. 183–4).



conception of inference. In doing so causalists begin with the observation that, just as acceptance of  $p$  and then  $q$  does not suffice for the occurrence of inference from  $p$  to  $q$ , nor does acceptance of  $p$ 's causing acceptance of  $q$  (Boghossian 2014, pp. 3–4; see also Brewer 1995, pp. 241–3; Huemer 2016, p. 151). Suppose, for instance, that I judge that  $p$ . Doing so might cause me to become nervous,  $p$ 's being the case being a worry for me. And becoming nervous in this way might happen to lead to my judging that  $q$ . Acceptance of  $p$  has caused acceptance of  $q$ , but clearly an inference from  $p$  to  $q$  need not thereby have occurred (c.f. Davidson 2001b, pp. 78–9). It is concluded, on consideration of such cases, that for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  in a to be specified “right way” (e.g. Boghossian 2014, pp. 3–4; see also Brewer 1995, p. 242; Huemer 2016, pp. 151–2). Cases such as that just considered are not cases in which an inference occurs because in such cases one's acceptance of  $q$  is caused by one's acceptance of  $p$  via some “deviant causal chain” rather than in “the right way”.

Once again, it does not follow from the fact that acceptance of  $p$  can cause acceptance of  $q$  and one not infer that for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  in some specifiable right way. The causalist rather trades on an inference to the best explanation here. We might ask why cases such as the above fail to constitute cases of inference. The best way to explain this is taken to be to posit a distinction between the right/non-deviant and wrong/deviant ways in which acceptance of  $p$  might cause acceptance of  $q$ , to claim that for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  in the right/non-deviant way, and to insist that in the relevant cases the causal connections are instead of the wrong/deviant kind. The distinction between the right and wrong way in which acceptance of  $p$  might cause acceptance of  $q$ , rather than being mandatory, is thus a distinction imposed by the causalist in order to make way for what they take to be the best explanation of the relevant cases' failure to constitute cases of inference. And the causal conception of inference is motivated by its making available the explanation in question. The causalist has it that for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  in the yet to be specified right way, where this conception of inference is motivated by its

being apt to explain why cases such as the above fail to qualify as cases of inference.

### 2.3. The Taking Condition

As emphasised, inference in the sense of concern is conscious inference of the sort apt to occur as a constituent of conscious reasoning. With this focus it becomes intuitive to think that making an inference requires that one “take it” that an appropriate inferential relation obtains between that which one infers and that which one infers from. The thought is that one would not infer  $q$  from  $p$  unless one took it that an appropriate inferential relation obtained between  $p$  and  $q$ . To infer  $q$  from  $p$  one must take it that such a relation obtains. The intuitiveness of this thought is acknowledged by Boghossian (2014), Fumerton (2015), Nes (2016), Neta (2013) and Valaris (2014). In fact, its intuitiveness even tends to be accepted by those who go on to deny that inferring in fact requires taking (e.g Wright 2014, pp. 28–30).

As well as inference’s requiring taking it is typical to have it that taking must be *causally* involved in inference. The motivating thought here is as follows. Intuitively, to infer  $q$  from  $p$  one must take it that an appropriate inferential relation obtains between  $p$  and  $q$ . One would not infer  $q$  from  $p$  if one did not take it that an appropriate inferential relation obtains between  $p$  and  $q$ . Furthermore, one would not qualify as doing so if one merely *incidentally* took it that such an inferential relation obtains—if one happened to take it that, say,  $q$  follows from  $p$ , but where one’s doing so played no role in one’s coming to accept  $q$ . The best way to capture this, the thought then is, is to say that when one infers  $q$  from  $p$  one’s taking must play an appropriate *causal* role in one’s doing so. When one infers  $q$  from  $p$  it will not then just be the case that one happens to take it that an appropriate inferential relation obtains between  $p$  and  $q$ . One’s taking will play a causal role in one’s inferring and thus in one’s coming to accept  $q$ . One will in that way qualify as accepting  $q$  in part *because* one takes it that an appropriate inferential relation obtains between  $p$  and  $q$ .

The above requirement on what it takes for an inference to occur is typically taken to be captured by Boghossian’s “Taking Condition”:

Inferring necessarily involves the thinker taking his premises to support his conclusion and drawing his conclusion because of that fact. (2014, p. 5)

The condition is in need of some refinement. Dogramaci (2016), for instance, objects that the Taking Condition must be false since we do not always infer from *premises*. We do not always do so, for example, in reasoning by conditional proof and by *reductio ad absurdum*. Consider the conditional proof case. One might suppose that  $p$  and then infer  $q$  from  $p$ . One might then discharge the supposition that  $p$  and conclude that if  $p$  then  $q$ . In doing so one infers that if  $p$  then  $q$ . But one does not infer if  $p$  then  $q$  from one's premise  $p$ , nor from any other premises. Inferring thus cannot require taking it that one's premises support one's conclusion, the objection goes, because not all inference involves inferring from premises.

A simple modification of the Taking Condition deals with the above issue. As already noted, and as Valaris (2016a, pp. 897–8) observes in his response to Dogramaci, whenever one infers there is always *something* which one infers from, it is just that what one infers from is not always one's premises. When, for instance, one reasons in accordance with conditional proof as above one supposes that  $p$  and then infers  $q$ . In doing so one can come to believe that  $q$  is derivable from  $p$  (and/or that  $q$  follows from  $p$ <sup>21</sup>). One can then infer if  $p$  then  $q$  from that  $q$  is derivable from  $p$  (or that  $q$  follows from  $p$ ). Similarly, when one reasons by *reductio ad absurdum* one might suppose that  $p$  and infer  $q$ . Again, in doing so one can come to believe that  $q$  is derivable from  $p$  (and/or that  $q$  follows from  $p$ ). One might then infer not- $p$  from that  $q$  is derivable from  $p$  (or that  $q$  follows from  $p$ ). The Taking Condition can be reformulated to accommodate such cases.<sup>22</sup>

---

<sup>21</sup> See sec. 4.11 below for discussion of the suggestion that one can acquire such beliefs *by* inferring.

<sup>22</sup> Alternatively the condition, as it is stated by Boghossian, could be given a more charitable reading than Dogramaci permits. The condition could be read as saying that whenever one infers one must take it that one's premises support one's conclusion where 'premises' is not to be understood narrowly as meaning premises in a formally representable logical argument. One's premises could rather be understood to just be whatever one infers from.

Just as we do not always infer from premises, as the Taking Condition implies, we do not always infer conclusions. In the above conditional proof reasoning, for instance, one infers  $q$  from  $p$  *without* concluding that  $q$ . One infers  $q$  from  $p$ , but  $q$  is not one's conclusion. Again, a simple modification of the Taking Condition deals with this issue. As was noted, whenever one infers there is always *something* which one infers, it is just that what one infers is not always one's conclusion. The Taking Condition can easily be reformulated to accommodate such cases.<sup>23</sup>

The above reveals that Boghossian's Taking Condition needs reformulating such that it demands that whenever one infers one takes it that an appropriate inferential relation obtains between *that which one infers* and *that which one infers from*, not necessarily between one's premises and conclusion. In stating the condition as he does Boghossian also takes a stance on what sort of inferential relation one must take to obtain in order to infer. To infer  $q$  from  $p$ , according to Boghossian, one must take it that  $p$  *supports*  $q$ . But there are several other options here. Thompson (1965) has it that to infer  $q$  from  $p$  one must take it that  $p$  is a *reason* to conclude that  $q$ . Broome (2013) and Valaris (2014, 2016b) have it that one must take it that  $q$  *follows from*  $p$ . Nes (2016), meanwhile, has it that one must take it that  $p$  *means*  $q$ .<sup>24</sup> A debate here will presuppose settled a matter which has not yet been engaged with: whether it is in fact correct to say that whenever one infers one must take it that *some* inferential relation obtains between that which one infers and that which one infers from. The Taking Condition tells us that this is correct because inference must causally involve taking. But it remains a possibility, given what has been said so far, that inference does not require taking whatsoever.

I will thus take it for granted that *if* one must take it that some inferential relation obtains between  $p$  and  $q$  in order to infer  $q$  from  $p$  then one must take it that  $q$  *follows from*  $p$ . Those who think that inference does require taking but

---

<sup>23</sup> Again, the alternative would be to allow a more charitable reading of the Taking Condition as stated by Boghossian. The condition could be read as saying that whenever one infers one must take it that one's premises support one's conclusion where 'conclusion' is just taken to mean that which one infers.

<sup>24</sup> Where the kind of meaning in question is Grice's (1957) "natural meaning".

who also think that this demand is too strong can read ‘follows’ as a placeholder for whatever inferential relation one in fact must take to obtain in order to infer.

Taking into account the above, the Taking Condition should thus be formulated as follows:

(Taking Condition) To make an inference one must accept that which one infers because one takes it that it follows from that which one infers from.

Thus stated the Taking Condition allows that when one infers one does not necessarily infer from premises nor infer a conclusion. It also builds in the assumption that to infer one must take it that what one infers follows from what one infers from, rather than potentially taking it that some other inferential relation obtains. Thus stated the Taking Condition should also be read as demanding that when one infers one accepts that which one infers *because* one takes, where this requires that one’s taking play a *causal* role in bringing about one’s acceptance of that which one infers. The Taking Condition in this way demands that taking plays a causal role in inference. Some will have it that it is a consequence of the Taking Condition that whenever one infers one’s taking causes the occurrence which is one’s inference. This will be correct if one’s inference is identical to the event which is one’s accepting that which one infers. But it will not be correct if one’s inference is identical to the process which begins with one’s acceptance of that which one infers from and terminates with one’s acceptance of that which one infers (see sec. 1.8 above). If the latter is correct then it is only a consequence of the Taking Condition that whenever one infers one’s taking is causally involved in one’s inference. Since at this stage we are remaining neutral when it comes to the ontology of inference it is not yet clear whether it is the former or just the latter which should be taken to be consequence of the Taking Condition (again, see sec. 1.8 above).

To advocate the Taking Condition, as it is now being understood, is to claim that when one infers  $q$  from  $p$  one’s taking it that  $q$  follows from  $p$  must play a causal role in one’s doing so such that one will qualify as accepting  $q$  *because* one takes it that  $q$  follows from  $p$ . But what, it might be asked, is it to take it that  $q$  follows from  $p$ ? The most natural suggestion is that to do so is simply to *believe* that  $q$  follows from  $p$ . Crane notes, for instance, that “[t]he essence of belief ...

is that it is taking something to be the case” (2013, p. 164). Saying that inferring  $q$  from  $p$  requires that one take it that  $q$  follows from  $p$  just sounds like another way of saying that inferring  $q$  from  $p$  requires that one believe that  $q$  follows from  $p$ . Despite this some have it that inference requires taking where taking it that  $q$  follows from  $p$  does not require believing that  $q$  follows from  $p$ . Chudnoff (2014) and Huemer (2016), for instance, both have it that to infer  $q$  from  $p$  one must have an *intuition* that  $q$  follows from  $p$ , but that one need not believe that  $q$  follows from  $p$ .

Whatever taking it that  $q$  follows from  $p$  amounts to it needs to be a kind of contentful state. Otherwise it is left unclear what it is to take it that  $q$  follows from  $p$  and appeal to such taking will no longer be illuminating when it comes to saying what it is for one to infer  $q$  from  $p$ . Wright expresses a similar sentiment when he confesses

to seeing no alternative interpretation of this “taking that” than to say that it requires an information-bearing state ... a state that, as we may say, registers the obtaining of an appropriate support relation or—to allow us to speak to the case where the inference is bad—that it represents such a relation as obtaining. (2014, p. 30)

To endorse the Taking Condition is to say that inference from  $p$  to  $q$  must causally involve a contentful state which represents  $q$  as following from  $p$ —most plausibly a belief—such that when one makes this inference one qualifies as accepting  $q$  *because* one takes it that  $q$  follows from  $p$ . Clear advocates of the Taking condition, thus understood include Chudnoff (2014), Huemer (2016), Leite (2008), and Fumerton (2015).<sup>25</sup>

---

<sup>25</sup> As before, setting aside the question of what inferential relation one must take to obtain in order to infer. Boghossian (2014) is not an advocate of the Taking Condition as it is being understood. On Boghossian’s view when one infers one does so because one takes it that what one infers follows from what one infers from. But this is not understood as a matter of the appropriate involvement of a contentful mental state. Rather, that inference involves taking is ensured by that to infer is to follow a rule of inference where

we can have no expectation that we will be able to give a non-circular analysis of what following a rule of inference amounts to. (2014, p. 17)

To infer, according to Boghossian, is to follow a rule of inference where no analysis of what it is to do so can be given and where doing so is assumed to suffice for satisfaction of the Taking Condition. As Boghossian sees, to say this is to give up on the project of giving an informative account of inference. What follows will reveal that there is no need to do so.

## 2.4. Two versions of causalism

Causalist accounts of inference can now be divided into two kinds: those which entail the Taking Condition and those which do not. On the first kind of causalist account for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  in the right way where this requires the appropriate causal involvement of one's taking it that  $q$  follows from  $p$ . One's taking it that  $q$  follows from  $p$  must be causally involved in bringing about one's acceptance of  $q$  such that one qualifies as accepting  $q$  *because* one takes it that  $q$  follows from  $p$ . Fumerton (2015, sec. 2) suggests an account of this form without attempting to specify what the relevant right way is. Chudnoff (2014), meanwhile, gives an account of the same form, but aims to be more informative when it comes to saying what the relevant right way is. According to Chudnoff, as seen, to infer  $q$  from  $p$  one must take it that  $q$  follows from  $p$  where this is a matter of having an intuition that  $q$  follows from  $p$ . Such an intuition, Chudnoff claims, can function as a mental imperative directing one to accept  $q$ . Chudnoff's view is then that for acceptance of  $p$  and taking it that  $q$  follows from  $p$  to cause acceptance of  $q$  in the right way they must do so such that one qualifies as following the mental imperative directing one to accept  $q$ , that mental imperative's being constituted by one's intuition that  $q$  follows from  $p$ .

On the second kind of causalist account for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  in the right way where this *does not* require the appropriate causal involvement of one's taking it that  $q$  follows from  $p$ . Advocates of causalist accounts of this form tend to insist that inference can be "blind" in the sense that it does not need to involve taking. Inference can be blind such that one might infer  $q$  from  $p$  without this involving one's taking it that  $q$  follows from  $p$  (e.g. Boghossian 2003; Dogramaci 2016; Kornblith 2012; Rosa forthcoming; Setiya 2013; Siegel forthcoming; Wright 2014).<sup>26</sup> The suggestion that inferring requires the involvement of taking, and is in that sense not blind, is taken to face insurmountable problems.

---

<sup>26</sup> In describing inference as potentially blind in this way authors take as their leave Wittgenstein's remark that

[w]hen I follow a rule, I do not choose. I follow the rule *blindly* (2009, para. 209).

Compare also his claim that

Various ways of characterising the right way in which acceptance of  $p$  must cause acceptance of  $q$  in order for an inference from  $p$  to  $q$  to occur are offered by advocates of the sort of causalist account now in question. Broome (2013), for instance, suggests that one's acceptance of  $p$  must yield one's acceptance of  $q$  in such a way that one qualifies as following an inference rule. He offers a dispositionalist account of what it is to follow such a rule. Brewer (1995) and Koziol (2017) suggest that one's acceptance of  $p$  must cause one's acceptance of  $q$  where one's acceptance of  $p$  causes one's acceptance of  $q$  in virtue of rationalizing it. One's acceptance of  $p$  must cause one's acceptance of  $q$  where the causation in question is a form of "causation in virtue of rationalization" (Brewer 1995, p. 246). Siegel (forthcoming), meanwhile, suggests that one's acceptance of  $p$  must cause one's acceptance of  $q$  such that in accepting  $q$  one is responding to  $p$  and such that in doing so one is responding to "information that admits of predicative structuring" (forthcoming, p. 14).

I will not go into detail into the various above attempts to characterise the right way in which acceptance of  $p$  needs to cause acceptance of  $q$  in order for an inference from  $p$  to  $q$  to occur according to causalists. There is no need to do so as, I will argue, the causalist conception of inference is misguided. For an inference from  $p$  to occur is not for acceptance of  $p$  to cause acceptance of  $q$  in some appropriate way. To say what it takes for an inference to occur is thus not to say what this appropriate way is. It is, however, worth looking at in some detail at the dispute over whether we should accept the Taking Condition. The latter dispute has been central in the majority of the recent literature on the nature of inference.

## 2.5. The case for the Taking Condition

All that has been said in favour of the Taking Condition so far is that it tends to be taken to be intuitive. Intuitively if one is to infer  $q$  from  $p$  one must take it that a relevant inferential relation obtains between  $p$  and  $q$ . Furthermore, one's taking it that this inferential relation obtains must be non-incidental. One must

---

it is not a kind of seeing on our part; it is our acting, which lies at the bottom of the language-game (1969, para. 204).



not just happen to take it that a relevant inferential relation obtains between  $p$  and  $q$  in order to infer  $p$  from  $q$ . One must rather take it that such a relation obtains where one's doing so is not incidental to one's inferring. This is taken to be best captured by the Taking Condition—by saying that whenever one infers one's taking it that a relevant inferential relation obtains must play a causal role in bringing about one's acceptance of that which one infers.

There tends to be a lack of clarity when it comes to offering more substantive considerations in favour of the Taking Condition. Boghossian, for instance, suggests that the Taking Condition should be endorsed because

no causal process counts as inference, unless it consists in an attempt to arrive at a belief by figuring out what, in some suitably broad sense, is supported by other things one believes. (2014, p. 5)

But it is in fact far from clear that an inference must be an attempt to do anything. I might, for instance, overhear someone playing Sudoku out loud. I might hear them say that the answer they are after is either 3 or 5 and cannot be 3 and infer that the answer is 5. In inferring that the answer is 5 it is far from clear that I must be attempting to do anything. I might even be trying my best to ignore the game. Furthermore, even supposing that inference always is an attempt to do something it seems far too restrictive to suggest that inference is always an attempt to arrive at belief by figuring out what is supported by what one believes. After all, as seen, we can infer from suppositions and from assumptions made for practical purposes. It is unclear why our aim in doing so must always be to arrive at belief by figuring out what is supported by things we believe.

A more concrete suggestion is that the Taking Condition should be endorsed in the light of considerations much like those offered in favour of the causalist conception above. As seen, acceptance of  $p$  can cause acceptance of  $q$  and one not infer  $q$  from  $p$ . The Taking Condition might be endorsed in order to explain why such cases fail to amount to cases of inference. Such cases are not cases of inference, the thought would be, because in such cases one fails to satisfy the Taking Condition, as one must in order to infer (Boghossian 2014, pp. 3–4; see also Korsgaard 2008, pp. 33–4; c.f. Neta 2013, pp. 388–9; Rödl 2013, pp. 213–4; Valaris 2016b, p. 4, 2016c, p. 4). However, as Boghossian

(2014, p. 5 note 2) sees, accepting the Taking Condition does not yield the resources to explain why all cases in which one's acceptance of something yields further acceptance without one's inferring fail to amount to cases of inference (see also McHugh and Way 2016, sec. 4.3). If, for instance, one's acceptance of  $p$  and one's taking it that  $q$  follows from  $p$  were to make one nervous and were that to lead to one's accepting  $q$ , one would not thereby have inferred  $q$  from  $p$ . Yet were this to occur one would have satisfied the Taking Condition.<sup>27</sup>

A further way to motivate the Taking Condition is gestured at by Boghossian when he suggests that we must endorse “something like” the Taking Condition if we are to continue “to think of ourselves as rational agents” (2014, p. 16). But as Boghossian admits the motivation gestured at here is “elusive” at best. What the connection between our being rational agents and the Taking Condition is supposed to be is left completely opaque.<sup>28</sup>

Yet another common way to motivate the Taking Condition is gestured at by Nes when he suggests that it must be endorsed in order to capture the sense in which inferences “bear their status as inferences on their conscious sleeves” (2016, p. 98; see also Boghossian 2014, p. 16; and for discussion Siegel forthcoming). The obvious way to do so would be to insist, as we have it here, that inference is not just conscious but *self-conscious*. However, the suggestion would go, merely insisting that inference is self-conscious is not enough to capture the *way* in which inference is self-conscious. The Taking Condition must be endorsed too in order to do so. The suggestion seems to be that the Taking Condition needs to be endorsed such that we can capture the respect in which our inferences are intelligible to us—such that we can capture that when one infers one does not merely know that one is inferring, but knows that one is doing so in such a way that one's doing so is intelligible to oneself.<sup>29</sup>

---

<sup>27</sup> It will not do to motivate the Taking Condition, and the version of the causalist conception which results, via appeal to its aptness to explain why *some but not all* cases of non-inference fail to qualify as cases of inference (see sec. 2.9 below).

<sup>28</sup> McHugh and Way (2016, sec. 4) discuss one unsuccessful way of making the connection, which they take to be Boghossian's, as well as several other unsuccessful attempts to motivate the Taking Condition.

<sup>29</sup> I ultimately suggest that something close to this claim is correct despite rejecting the Taking Condition. I ultimately claim that one must believe that  $q$  follows from  $p$  *when* one infers  $q$  from  $p$  in order for one's inference to be intelligible to oneself (see sec. 4.8).

Several authors note how when we infer we do not simply accept that which we infer without any real or apparent appreciation of why we are doing so or of the rationality of our doing so. Peacocke, for instance suggests that the inferences we make are “rational transitions” rather than mere “blind leaps into the dark, inclinations to make transitions in thought that just grip and take over the thinker’s rational self” (2003, p. 127). Dummett (1991, Chapter 9) makes what I take to amount to the same point. He describes circumstances in which we are conditioned to abide by alleged logical laws. And he takes abiding by the laws as not sufficient for inferring in accordance with them. We must also at least apparently appreciate why we are doing so. If we were to abide by alleged laws without even apparent appreciation of why we were doing so, Dummett says, although “[t]he rules of the language-game would be clear enough [the] point would now escape us” (1991, p. 207). When it comes to making inferences, Dummett claims, “[i]t is not enough [...] to be master of a practice” (1991, p. 207; see also Brewer 1995, sec. 3).

It remains to be seen how the above thought leads some to endorsement of the Taking Condition. The reasoning, I take it, goes as follows. Inference of the sort of concern, as seen, is self-conscious. When one makes an inference of the sort of concern here one does so self-consciously. Furthermore, when one does so one does not merely know that one is inferring. If one is to know that one is inferring  $q$  from  $p$ , the thought is, one must be aware of why one is doing so or have at least apparent awareness the rationality of one’s doing so. One’s inference must be in this way intelligible if it is to be self-conscious. The Taking Condition is endorsed in order to capture the respect in which our inferences must be intelligible to us in this way.

To see how the Taking Condition might be thought to be apt to play the above role suppose that it is true and that one infers  $q$  from  $p$ . In doing so one accepts  $q$  as a result of one’s accepting  $p$  and one’s taking it that  $q$  follows from  $p$ . One thereby at least seems to appreciate that  $q$  follows from  $p$  and accepts  $q$  as a result. One’s inference from  $p$  to  $q$  is self-conscious *and* in inferring  $q$  from  $p$  one accepts  $q$  because one takes it that  $q$  follows from  $p$ . One might thereby be thought to appreciate why one infers  $q$ . Were one asked why one did so it might be thought that one would be in a position to respond by saying ‘I

inferred  $q$  from  $p$  because I accepted  $p$  and took it that  $q$  follows'. One similarly might be thought to at least seem to appreciate why one's inferring is rational. Were one asked why one's inference was rational one would seem to be in a position to respond by pointing out that  $q$  follows from  $p$ . In this way, endorsing the Taking Condition might be thought to best capture the way in which inferences are not just self-conscious but intelligible from the point of view of the inferrer.<sup>30</sup>

## 2.6. The case against the Taking Condition: overintellectualization

Conclusive grounds to accept the Taking Condition and the resultant version of the causalist conception of inference have not been identified. Furthermore, it is common to have it that there are compelling grounds to reject the condition. Inference *cannot* be such that it must involve taking in the way the Taking Condition tells us.

As seen (sec. 1.7), it is common to object to accounts of inference on the grounds that they give us an “overintellectualized” picture of inference. Some will have it that it cannot be that in order to infer one must know that one is inferring. Some will likewise have it that it cannot be that in order to infer one must take it that an appropriate inferential relation obtains, and thus that the Taking Condition cannot be true. Animals and young children lack the relevant concepts, the objection goes, and yet are perfectly capable of making inferences. As was also seen, however, to object to the claim that inference is self-conscious and to the account of inference that results in the above way is to equivocate on the meaning of ‘inference’. Inference of the sort of concern here is conscious inference of the sort which is apt to occur as a constituent of self-conscious and active reasoning. The claim that animals and young children can make inferences of this sort, rather than inferences in a more permissive sense, is not well supported. That animals and young children can succeed in various problem-solving tasks, for instance, does not show that they are capable of conscious reasoning of the kind in question. Similarly, the question of whether one must possess concepts of relevant inferential relations in order to be able

---

<sup>30</sup> I will contest the suggestions below (see sec. 2.10).

to make inferences of the kind in question is still up in the air. It is yet to be seen whether self-conscious reasoners must possess such concepts. If there is a good argument to that conclusion then the claim that animals and young children can make inferences in a more permissive sense than that of concern will not cast doubt on it.

## 2.7. The case against the Taking Condition: Carroll's regress

Endorsement of the Taking Condition is also often taken to lead to the difficulty at the heart of Carroll's (1895) fable (e.g. Boghossian 2014; McHugh and Way 2016; Mole 2018; Rosa forthcoming; Valaris 2014, 2016b; Wright 2014). In it the Tortoise tells Achilles that he knows of

a race-course, that most people fancy they can get to the end of in two or three steps, while it *really* consists of an infinite number of distances, each one longer than the previous one (1895, p. 278).

The Tortoise draws Achilles' attention to the following argument:

- (A) Things that are equal to the same are equal to each other.
- (B) The two sides of this Triangle are things that are equal to the same.

Therefore,

- (Z) The two sides of this Triangle are equal to each other.

Both Achilles and the Tortoise accept that it is possible to accept (A) and (B) and not accept (Z). Achilles' challenge is then to take the Tortoise to be such a subject and to force him to accept (Z).

The Tortoise accepts any proposition which Achilles asks him to, the first being

- (C) (Z) follows from (A) and (B).<sup>31</sup>

At this stage the Tortoise accepts (A), (B) and (C), yet he still does not accept (Z). As Achilles admits, this seems possible in the same way that it is possible

---

<sup>31</sup> In fact, in the fable Achilles asks the Tortoise to accept that if (A) and (B) are true, (Z) *must* be true. Achilles takes this to be equivalent to saying that '(A), (B), therefore (Z)' is a valid argument. So Achilles seems to think that when we infer we need to take inferential relations to obtain where these relations are understood modally. Having set aside the matter of what inferential relations we must take to obtain to infer, I formulate (C) in terms of what follows from what for continuity.

to accept (A) and (B) and not (Z). In the hope that it will lead to his accepting (Z) Achilles thus asks the Tortoise to accept

(D) (Z) follows from (A), (B) and (C).

But the Tortoise still does not accept (Z). Again, this seems possible in the same way that it is possible for him to accept (A), (B) and (C) but not (Z). Achilles thus asks the Tortoise to accept

(E) (Z) follows from (A), (B), (C) and (D).

But the Tortoise still does not accept (Z), which on the same grounds again seems possible. The Tortoise cannot be forced to accept (Z) no matter how many propositions other than (Z) he accepts. We are left to draw our own moral.

It is noteworthy that despite that Achilles' challenge is to get the Tortoise to accept (Z), and despite that the Tortoise will accept any proposition he is asked to, Achilles does not simply ask him to accept (Z). This suggests that Achilles' real challenge is not just to get the Tortoise to accept (Z) but to get him to accept (Z) *by inferring it*. And Achilles aims to do so by getting the Tortoise to accept that there are relevant inferential connections between (Z) and propositions he is willing to accept without inference.

It is typical to extract the following problem from Carroll's fable—one which concerns the very possibility of inference given acceptance the Taking Condition. Of concern is what it is for an inference to occur. What is it, for instance, for an inference from  $p$  to  $q$  to occur? First, as seen, the causalist has it that to infer  $q$  from  $p$  one must accept  $p$  and as a result accept  $q$ . The Taking Condition tells us that one must also do so because one takes it that that  $q$  follows from  $p$  where this is a matter of the causal involvement of one's doing so in one's coming to accept  $q$ . So when one infers  $q$  from  $p$  acceptance of  $p$  along with taking it that  $q$  follows from  $p$  causes acceptance of  $q$ .

One could accept  $p$  and take it that  $q$  follows from  $p$  and yet not accept  $q$ . Achilles and the Tortoise both acknowledge that such a situation is possible. Similarly, as seen, one could accept  $p$ , take it that  $q$  follows from  $p$ , and later accept  $q$  as a result, where this does not suffice for one's having inferred  $q$  from  $p$ . In the causalists terms, one's acceptance of  $q$  might be the result of some "deviant causal chain". As was also seen, if an advocate of the Taking Condition

is to say what it takes for an inference from  $p$  to  $q$  to occur they thus must specify “the right way” in which acceptance of  $p$  and taking it that  $q$  follows from  $p$  would need to cause acceptance of  $q$  such that one qualifies as having inferred  $q$ . The alleged problem is that the only way for them to do so would be to say that this would happen if one were to infer  $q$  from  $p$  and from that  $q$  follows from  $p$ . Acceptance of  $p$  and taking it that  $q$  follows from  $p$  would yield acceptance of  $q$  such that one qualifies as having inferred  $q$  if one were to infer  $q$  from that  $p$  and that  $q$  follows from  $p$ .<sup>32</sup> Furthermore, there is no other way for the advocate of the Taking Condition to say that acceptance of  $p$  and taking it that  $q$  follows from  $p$  could play this role. McHugh and Way claim that in this way

[s]atisfying the Taking Condition would seem to require a further inference (2016, p. 319; see also Boghossian 2014, p. 9; Fumerton 2015, p. 212).

The above regress argument depends on a crucial assumption. It is claimed that the only way to say that acceptance of  $p$  and taking it that  $q$  follows from  $p$  could bring about acceptance of  $q$  such that one qualifies as having inferred  $q$  would be to say that one could infer  $q$  from  $p$  and ( $q$  follows from  $p$ ). We can label what is assumed here the ‘Core Assumption’ of the regress argument:

(Core Assumption) If acceptance of  $p$  and taking it that  $\Phi$  obtains is to bring about acceptance of  $q$  such that  $q$  qualifies as having been inferred then one must infer  $q$  from  $p$  and  $\Phi$ .

If acceptance of  $p$  and taking it that  $q$  follows from  $p$  brings about acceptance of  $q$  such that  $q$  qualifies as having been inferred, the Core Assumption tells us, then  $q$  must have been inferred from that  $p$  and that  $q$  follows from  $p$ . There is alleged to be no other way in which acceptance of  $p$  and taking it that  $q$  follows from  $p$  could play this role.

---

<sup>32</sup> Things might seem to have already gone wrong at this point. In question is what it takes for an inference from  $p$  to  $q$  to occur. Yet an appeal is being made to inference from  $p$  and ( $q$  follows from  $p$ ) to  $q$ —a distinct inference. One cannot really infer  $q$  from  $p$ . What we might call an inference from  $p$  to  $q$  must really be an “enthymematic” inference—an inference with a suppressed premise (namely, that  $q$  follows from  $p$ ) (Fumerton 2015, p. 213). The regress argument leads us to the conclusion that what looks like a simple inference from  $p$  to  $q$  must really be an enthymematic argument with *infinitely many* suppressed premises.

It is clear that the combination of the Taking Condition *and* the Core Assumption leads to a vicious regress. To see this, suppose that both are true and that one infers  $q$  from  $p$ . To do so, according to the Taking Condition, one must take it that  $q$  follows from  $p$ . Furthermore, given the Core Assumption, if acceptance of  $p$  and taking it that  $q$  follows from  $p$  are to yield acceptance of  $q$  such that one qualifies as having inferred  $q$  then one must infer  $q$  from  $p$  and ( $q$  follows from  $p$ ). But to make the latter inference, according to the Taking Condition, one must take it that  $q$  follows from ( $p$  and ( $q$  follows from  $p$ )). Furthermore, given the Core Assumption, if acceptance of  $p$  and taking it that  $q$  follows from  $p$  and taking it that  $q$  follows from ( $p$  and ( $q$  follows from  $p$ )) are to bring about acceptance of  $q$  such that one qualifies as having inferred  $q$  then one must make yet another inference—one must infer  $q$  from  $p$  and ( $q$  follows from  $p$ ) and ( $q$  follows from ( $p$  and ( $q$  follows from  $p$ ))). Accepting the Taking Condition *and* the Core Assumption leads to the conclusion that to make any inference one would need to make an infinitely complex inference.

As noted, the threat of Carrollian regress is often taken to be sufficient grounds to reject the Taking Condition. Advocates of the condition, meanwhile, typically make complex manoeuvres in order to attempt to avoid the regress (e.g. Chudnoff 2014). In fact, there is a simple and seemingly adequate response to the Carrollian worry available for the advocates of the Taking Condition. Advocates of the Taking Condition can simply reject the further assumption which leads to regress: they can reject the Core Assumption. Fumerton (2015, sec. 2), for instance, suggests that the lesson from Carroll's regress is that the assumption must be false. That the Taking Condition along with the Core Assumption leads to regress, he suggests, can simply be taken to show that the Core Assumption must be mistaken. There must be some way for acceptances and taking it that some relevant inferential relation obtains to bring about further acceptance such that an inference qualifies as having occurred which is not via inference. There must, for instance, be a way for acceptance of  $p$  and taking it that  $q$  follows from  $p$  to bring about acceptance of  $q$  such that an inference qualifies as having occurred other than via inference from  $p$  and ( $q$  follows from  $p$ ) to  $q$ . To respond in this way is to insist that the “right way” in which acceptance and taking must bring about further acceptance such that an



inference qualifies as having occurred is not simply via further inference. Furthermore, if the Taking Condition really is a well-motivated then it seems there is nothing wrong with this response. Supposing there are compelling grounds to accept the Taking Condition, and given the Taking Condition along with the Core Assumption leads to Carrollian regress, the Core Assumption must be rejected. And it can be rejected on the basis of our grounds for the Taking Condition. As Fumerton sees, if the Taking Condition really is well motivated then this response is not ad hoc, at least not until more compelling grounds have been offered on behalf of the Core Assumption.

## 2.8. Trouble for the causalist: inferences beyond one's ken

On the causalist conception of inference for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  in some specifiable right way. Causalists who accept the Taking Condition have it that this requires that one's taking it that  $q$  follows from  $p$  be causally involved in bringing about one's acceptance of  $q$ , such that one qualifies as accepting  $q$  *because* one takes it that  $q$  follows from  $p$ . Whether the Taking Condition should be accepted has not yet become clear. Regardless, however, all forms of causalism can be seen to be problematic. The causalist conception of inference, I will argue, allows that inference can occur without doing so self-consciously. If the occurrence of an inference from  $p$  to  $q$  is just a matter of acceptance of  $p$  causing acceptance of  $q$  in an appropriate way then such an inference can occur and one fail to know that one is inferring. The causalist conception of inference, in sum, fails to accommodate the self-consciousness of inference.

As will be seen, a natural response to the above worry on the behalf of the causalist would be to build the self-consciousness of inference into the causalist conception of inference. It could be insisted that all inferences are self-conscious on the basis of the very same grounds on which we have it that inference must be self-conscious here. I will go on to consider and object to this modified causalism after the present worry has been made clear (see sec. 2.9). To modify the causalist conception in the way in question, I will argue, is to render the conception unmotivated.

The present worry stems from the fact that the causalist conception of inference allows that an inference from  $p$  to  $q$  could occur and there be no difference from one's point of view between one's case and one in which one merely accepts  $p$  and then  $q$  without inferring. It similarly allows that an inference from  $p$  to  $q$  could occur and there be no knowledge one has which one would lack in a case in which one merely accepts  $p$  and then  $q$  without inferring (except the required knowledge that one is inferring itself). To allow this much is to allow that one can infer without knowing that one is doing so.

To illustrate the point suppose that one accepts  $p$  and then accepts  $q$  as a result. Or, if one accepts the Taking Condition, suppose that one accepts  $p$  and takes it that  $q$  follows from  $p$  and then accepts  $q$  as a result. One's accepting  $p$  (and one's taking it that  $q$  follows from  $p$ ) causes one's acceptance of  $q$ . On the causalist conception of inference for an inference to occur here is for one's acceptance of  $p$  (along with one's taking it that  $q$  follows from  $p$ ) to cause one's acceptance of  $q$  in the right way. Supposing this is correct we get two possible cases. In the first case one's acceptance of  $p$  (along with one's taking it that  $q$  follows from  $p$ ) causes one to accept  $q$  in the right way and one qualifies as inferring  $q$  from  $p$  accordingly. In the second case one's acceptance of  $q$  is not caused in the right way and one fails to qualify as inferring. In the causalist's terms, one's acceptance of  $p$  (along with one's taking it that  $q$  follows from  $p$ ) causes one's acceptance of  $q$  via some "deviant causal chain". In the former case one infers  $q$  from  $p$ . In the latter case one does not, and instead merely accepts  $q$  where one happens to do so because of prior acceptance of  $p$  (along with one's taking it that  $q$  follows from  $p$ ). Furthermore, the difference between the two cases is a matter which may be wholly beyond one's ken. The only difference that there need be between the two cases is a matter of the way in which one's acceptance of  $q$  is in fact caused. This means that there need be no difference between the inference case and the non-inference case from one's point of view. Nor need there be any knowledge one has in the inference case which one lacks in the non-inference case (except the required knowledge that one is inferring itself). And the non-inference case, we can suppose, will not even be a case in which one *seems* to be inferring. One simply accepts that  $q$  outright and happens to do so because of prior acceptance of  $p$  (and one's taking it that  $q$  follows from

*p*). One need not seem to be inferring for this to occur. Given all this, it is perfectly possible that in the case in which one is inferring one fails to know that one is doing so. The only difference between this case and one in which one may not even seem to be inferring is how one's acceptance of *q* is in fact caused, where this is a matter which may be completely beyond one's ken. There need be no difference between the two cases from one's point of view.

The problem, in sum, is that if whether one infers when acceptance of *p* (along with taking it that *q* follows from *p*) causes acceptance of *q* is determined by whether one's acceptance of *q* is caused in the right way then there need not be any difference from one's point of view between a case in which one infers *q* from *p* and a case in which one merely accepts *q* and happens to do so because of acceptance of *p* (along with one's taking it that *q* follows from *p*). Nor need there be any knowledge one has in the inference case which one lacks in the non-inference case (except the required knowledge that one is inferring itself). And to allow this much is to allow that one could infer *q* from *p* without knowing that one is doing so. The causalist conception of inference fails to accommodate the self-consciousness of inference.

It would be no help, in response to the above worry, to appeal to the suggestion that inference has a distinctive phenomenology (c.f. James 1976). It would be no help, that is, to suggest that inference is a conscious occurrence with a distinctive phenomenal character apt to make inference distinct from non-inference from one's point of view such that one can always know that one is inferring when doing so. Even supposing, as is somewhat doubtful (see O'Brien 2005, pp. 583–4), that such an account of the self-consciousness of inference could be given, the account will not be available for the causalist. The causalist lacks grounds to insist that a case in which one infers *q* from *p* must differ phenomenologically from one in which acceptance of *p* (along with taking it that *q* follows from *p*) just happens to yield acceptance of *q*. For the causalist, after all, the only difference which there need be between two such cases is that in one case one's acceptance of *q* is caused in the right way. Why think that the differences in the causal connections which hold between such conscious occurrences *necessitate* a difference in the phenomenology?

Causalist who accept the Taking Condition are no better off here than those causalists who reject the condition. Suppose, for instance, that acceptance of  $p$  along with taking it that  $q$  follows from  $p$  yields acceptance of  $q$ . The only difference between a case in which this occurs and one infers and a case in which this occurs and one does not infer, for the causalist now in question, is a matter of how one's acceptance of  $q$  is in fact caused. This may be a matter which is entirely beyond one's ken and makes no difference from one's point of view. The causalist who endorses the Taking Condition is thus equally committed to allowing that one can infer and fail to know that one is doing so.

The objection is that the causalist fails to accommodate the self-consciousness of inference. If the matter of whether one is inferring  $q$  from  $p$  or merely accepting  $q$  as a result of acceptance of  $p$  (and taking it that  $q$  follows from  $p$ ) is settled entirely by whether one's acceptance of  $q$  is caused in the right way then one can infer without knowing that one is doing so. It might seem natural to respond to this objection on behalf of the causalist as follows. We have it here that inference is self-conscious. Conscious inference of the sort apt to occur as a constituent of conscious and active reasoning is self-conscious. Given this, it seems the causalist is entitled to simply insist that inference is self-conscious on the very same grounds on which we have it that this is so. When asked on what grounds they insist that inference is self-conscious it seems the causalist can simply reply that they have it that inference is self-conscious on the very same grounds on which we have it that that is so here. To succeed in undermining the causalist conception of inference I need to say what is wrong with the modified causalist conception of inference which results.

## **2.9. Causalism's motivation undermined**

The causalist conception of inference, as initially understood, tells us that for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to yield acceptance of  $q$  in the right way. The problem with the causalist conception thus understood, I have argued, is that it fails to accommodate the self-consciousness of inference. In response it can be insisted that inference is self-conscious on the same grounds on which we have it that that is so. If this response is endorsed then the causalist conception of inference needs modifying accordingly. This might

be done by building the claim that inference is self-conscious into one's account of the right kind of causation or by offering a conjunctive analysis of inference. The causalist, that is, can say that for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  (and taking it that  $q$  follows from  $p$ ) to yield acceptance of  $q$  in the right way where acceptance of  $p$  (and taking it that  $q$  follows from  $p$ ) only does so if one infers self-consciously. Acceptance of  $q$  is only yielded in the right way such that one qualifies as inferring  $q$  from  $p$  if one infers self-consciously. Or they can offer a conjunctive analysis and say that an inference from  $p$  to  $q$  occurs iff acceptance of  $p$  (along with taking it that  $q$  follows from  $p$ ) yields acceptance of  $q$  in the right way *and* one knows that one is inferring  $q$  from  $p$ .

The problem with the resultant conception of inference stems from the fact that if it is to be acceptable then appeal to knowledge of inference in one's account of inference must likewise be acceptable, not rendering it problematically circular. The causalist now in question appeals to knowledge of inference in their account of inference, either as a condition which must be satisfied for an inference to occur, or in their account of the right kind of causation. But if appeal to knowledge of inference in one's account of inference is both acceptable and required, and given that, since inference is self-conscious, one is inferring iff one knows that one is inferring, one's lack of knowledge that one is inferring can be appealed to in order to account for why *any* given case of non-inference fails to amount to a case of inference. And if that can be done for *any* case in which we might want to account for an inference's failure to occur then appeal to causation and to the right kind of causation in one's account of inference cannot be motivated by the need to account for cases of non-inference's failure to amount to cases of inference. That, however, as seen (sec. 2.2), is precisely how the causalist conception of inference is motivated.

The point here can be explained in terms of the search for necessary and sufficient conditions for the occurrence of an inference. I have suggested that the claim that part of what it is for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  is motivated by consideration of cases in which one accepts  $p$  and then  $q$  where one does not infer and where there is allegedly no causal connection between one's acceptance of  $p$  and  $q$  whatsoever (see sec. 2.2 above). The best explanation of the relevant cases' failure to amount

to cases of inference is taken to involve appeal to the fact that part of what it is for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$ . One might, for example, judge that  $p$  and then later incidentally judge that  $q$  without inferring  $q$  from  $p$ . The best explanation of why such a case fails to amount to a case of inference is alleged to be that part of what it is for an inference to occur is for acceptance of  $p$  to cause acceptance of  $q$ , and that acceptance of  $p$  does not cause acceptance of  $q$  in the case in question. I have likewise suggested that the causalist conception of inference is motivated by consideration of cases in which one accepts some  $p$  and then some  $q$  where one's acceptance of  $p$  causes one's acceptance of  $q$  and where one does not infer  $q$  from  $p$  nonetheless (again, see sec. 2.2 above). One might, for example, judge that  $p$  where one's doing so makes one nervous and where that in turn leads to one's judging that  $q$ . The best explanation of such a case's failure to amount to a case in which one infers is taken to involve appeal to the fact that what it is for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  in the right way. One does not infer  $q$  from  $p$  in the case, it is alleged, because to infer  $q$  from  $p$  one's acceptance of  $p$  must cause one's acceptance of  $q$  in the right way, where one's acceptance of  $p$  does not do so in the case in question.

In terms of necessary conditions, we begin with the observation that to infer  $q$  from  $p$  one must accept  $p$  and then accept  $q$ . We begin, that is, with the following condition on what it takes for an inference from  $p$  to  $q$  to occur:

(C1) One must accept  $p$  and then accept  $q$ .

It is then observed that one can satisfy (C1) and not infer  $q$  from  $p$ . In order to account for such cases' failure to amount to cases of inference the following further condition is imposed:

(C2) One's acceptance of  $p$  must cause one's acceptance of  $q$ .

It is then observed, however, that one can satisfy (C1) and (C2) and not infer  $q$  from  $p$ . The need to account for such cases' failure to amount to cases of inference motivated the imposing of the following further condition:

(C3) One's acceptance of  $p$  must cause one's acceptance of  $q$  in the right way.

Initially, it was suggested that the satisfaction of these three conditions suffices for the occurrence of an inference. The causalist conception of inference was proposed accordingly: for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  where one's acceptance of  $p$  does so in a to be specified right way. This conception's failure to ensure the self-consciousness of inference, however, reveals that (C1) – (C3) can be satisfied and one not infer  $q$  from  $p$ . That is so since one could satisfy (C1) – (C3) and not know that one is inferring (see sec. 2.8). In response, the modified causalist conception can be endorsed. In order to account for all cases in which inferences fail to occur the advocate of this conception can endorse the following further condition on what it takes for an inference from  $p$  to  $q$  to occur:

(C4) One must know that one is inferring.

Alternatively, they can insist that (C3) must be understood in terms of knowledge. They can, that is understand (C3) as:

(C3\*) One's acceptance of  $p$  must cause one's acceptance of  $q$  in the right way where one's acceptance of  $p$  does so iff one knows that one is inferring.

(C3\*), however, can be broken down into the condition that one's acceptance of  $p$  must cause one's acceptance of  $q$  and the further condition that one must know that one is inferring. Inference's being self-conscious, they must have it that one's acceptance of  $p$  causes one's acceptance of  $q$  in the right way iff one knows that one is inferring. (C3\*), that is, is equivalent to (C2) plus (C4).

The aim above is to arrive at a conception of inference which is apt to yield the best explanation of why all cases of non-inference fail to amount to cases of inference. Reflection on cases in which one accepts  $p$  and then  $q$  without inferring motivates imposing (C2) as a condition on what it takes for an inference from  $p$  to  $q$  to occur. However, it quickly becomes apparent that appeal to (C2) will not suffice to explain why all cases in which one accepts  $p$  and then  $q$  without inferring fail to amount to cases of inference. What can now be seen to be an odd move is then made. Rather than abandoning (C2), (C2) is *maintained* and *added to*—(C3) is added. One might, at this stage, wonder why an alternative condition to (C2) is not sought. Why *keep* and *add to* (C2) when it comes to explaining why all of the relevant cases of non-inference fail to amount

to cases of inference once it becomes evident that appeal to (C2) is not apt to do so? Once (C3) is added, furthermore, it quickly becomes apparent that appeal to (C2) and (C3) will not suffice to explain why all cases in which one accepts  $p$  and then  $q$  which do not amount to cases of inference fail to do so. (C4) is added accordingly (or (C3) is read as (C3\*), which amounts to (C2) plus (C4)). But appeal to (C4) will suffice *alone* to account for why *all* cases in which one accepts  $p$  and then  $q$  without inferring fail to amount to cases of inference, since one infers iff one satisfies (C4). It is because of this that adding condition (C4) (either explicitly, or implicitly by understanding (C3) as (C3\*)) undermines the causalist in question's motivation for (C2) and (C3). If one concedes that we must accept (C4) as part of our characterisation of what it is for an inference from  $p$  to  $q$  to occur then (C4) can be appealed to rather than (C2) and (C3) in order to explain why *any* given case in which one accept  $p$  and then  $q$  without inferring  $q$  from  $p$  fails to amount to a case of inference. The advocate of the modified causalist conception of inference cannot motivate (C2) and (C3) via appeal to the need to explain why cases of non-inference fail to amount to cases of inference, as they in fact do, since they concede that (C4) must also be endorsed and since appeal to (C4) suffices to explain why all such cases fail to amount to cases of inference alone.

The advocate of the modified causalist conception of inference now in question concedes that one both can and must appeal to knowledge of inference in one's account of inference. Once such appeal is conceded to be legitimate and required it can be made in order to explain why any given case of non-inference fails to amount to a case of inference. Given this, insisting in addition that appeal to causation and to the right kind of causation must be made in one's account of what it is for an inference to occur is not motivated. The latter appeal, for the causalist, is motivated by the need to explain why cases of non-inference fail to amount to cases of inference. But appeal to knowledge of inference suffices to do that alone. To endorse the modified causalist conception of inference is thus to render one's conception of inference unmotivated.

The causalist conception of inference, as initially understood, fails to accommodate the self-consciousness of inference. The conception allows that



one can infer without doing so self-consciously. To accept the modified causalist conception, where the self-consciousness of inference is built into the conception, meanwhile, is to undermine one's motivation for accepting a causalist conception in the first place. A final possible response to consider on behalf of the causalist is as follows. One might attempt to modify the causalist conception of inference in order to ensure that inference is self-conscious without simply building the self-consciousness of inference into the conception. The reason the causalist conception as initially understood fails to accommodate the self-consciousness of inference is that it allows that one can infer and there be no difference from one's point of view between one's case and one in which one does not even seem to be inferring. To rule this out one might build the claim that to be inferring one must *believe* that one is inferring into the causal conception.<sup>33</sup> If one must believe that one is inferring to be doing so then there will thereby be a difference from one's point of view in the relevant sense between a case in which one infers and a case in which one does not even seem to be inferring. As before, this might be done by building the claim that inference requires belief into one's account of the right kind of causation or by offering a conjunctive analysis of inference. The causalist, that is, might say that for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  (and taking it that  $q$  follows from  $p$ ) to yield acceptance of  $q$  in the right where acceptance of  $p$  (and taking it that  $q$  follows from  $p$ ) only does so if one believes that one is inferring. Or they might offer a conjunctive analysis and say that an inference from  $p$  to  $q$  occurs iff acceptance of  $p$  (along with taking it that  $q$  follows from  $p$ ) yields acceptance of  $q$  in the right way *and* one believes that one is inferring  $q$  from  $p$ .

On the modified causalist conception that results to be inferring one must believe that one is doing so. Whenever one infers one thereby truly believes that one is doing so. Some might worry about whether this conception is really apt to capture that inference is self-conscious—i.e. that whenever one infers one *knows* that one is inferring—rather than just that whenever one infers one will

---

<sup>33</sup> Alternatively, the claim that one must *seem* to be inferring to be doing so might be built into the conception. Or one might build the claim that a conscious occurrence must have a distinctive phenomenology in order to be an inference into one's conception of inference. These suggestions face difficulties exactly analogous to those faced by the belief proposal.

truly believe that one is doing so. On the conception in question to infer  $q$  from  $p$  one must accept  $p$  and then  $q$  where one's acceptance of  $p$  causes one's acceptance of  $q$  in the right way *and* where one believes that one is inferring (or where one's acceptance of  $p$  only causes one's acceptance of  $q$  in the right way if one believes that one is inferring). If that suffices for the occurrence of an inference then whenever one infers one truly believes that one is inferring. But for most  $p$  it is possible to truly believe that  $p$  without knowing that  $p$ .<sup>34</sup> The question thus arises of whether the conception in question accommodates that inference is self-conscious, or whether further conditions must be sought and built into the conception in order to ensure that it does.<sup>35</sup> Rather than pursue this difficulty, it can be brought out how the suggestion in question faces a problem much like that raised by Snowdon (1980) for the causal account of seeing. Snowdon argues that once the direct realist account of seeing is made available the causal account of seeing can be seen to be unmotivated, since the latter account is motivated by the thought that it *must* be appealed to in order to explain why cases of failing to see do not amount to cases of seeing (and since appeal to the direct realist account of seeing is apt to explain why any given case in which one fails to see does not amount to a case of seeing). Similarly, the modified causalist conception of inference in question can be seen to be unmotivated one it is seen that an account of inference where inference is characterised directly in terms of its self-consciousness is available (and apt to account for why any given case of non-inference fails to amount to a case of inference).

The suggestion in question is that we must endorse a conception of inference on which for an inference from  $p$  to  $q$  to occur one must accept  $p$  and then  $q$  where one's acceptance of  $p$  causes one's acceptance of  $q$  and on which for one to infer  $q$  from  $p$  one must believe that one is inferring. The conception

---

<sup>34</sup> I consider possible exceptions in my (forthcoming).

<sup>35</sup> If such further conditions are needed, and if in stating them appeal to knowledge of inference must be made such that the resultant conception of inference can be appealed to in order to explain why any given case of non-inference fails to amount to a case of inference, then the conception that results will be rendered unmotivated as per above. As seen, if appeal to knowledge is ultimately needed in one's conception of inference then appeal to causation and to the right kind of causation is rendered unmotivated.

is motivated by the need to account for why cases in which one accepts  $p$  and then  $q$  without inferring fail to amount to cases of inference. An alternative conception of inference is available on which for an inference from  $p$  to  $q$  to occur is just for a distinctive kind of self-conscious occurrence to occur (see ch. 4 below). On this conception, for an inference from  $p$  to  $q$  to occur is not for acceptance of  $p$  to cause acceptance of  $q$  in some specifiable right way, nor need inference in some way involve belief that one is inferring. For an inference to occur is rather just for a distinctive kind of self-conscious occurrence to occur. The very availability of this alternative conception of inference, I suggest, undermines the motivation for the modified causalist conception of inference now in question. The alternative's availability reveals that the modified causalist conception is unmotivated, since its very availability reveals that the modified causalist conception cannot be motivated by the *need* to appeal to it in order to explain why cases of non-inference fail to amount to cases of inference. The alternative conception of inference on which inference is characterised directly in terms of its self-consciousness can account for all such cases.<sup>36</sup>

## 2.10. The Taking Condition's motivation undermined

With my case against the causalist conception of inference on the table it is worth returning to the question of whether the Taking Condition is well motivated. After all, I have argued that the causalist cannot accommodate the self-consciousness of inference without undermining the causalist conception's motivation. It might seem like the Taking Condition thereby loses its motivation too. The most promising suggestion in favour of the Taking Condition was that it needs to be endorsed if we are to capture the way in which one's inferences are intelligible to oneself. When one infers, the suggestion was, one appreciates why one is doing so and/or has at least apparent appreciation of the rationality of one's doing so. It was suggested that the Taking Condition needs to be endorsed in order to capture this. But if one's inferences are to be intelligible to

---

<sup>36</sup> It thus must be admitted that the objection in question to the modified causalist conception turns on the adequacy the alternative conception of inference which is yet to be developed (see ch. 4). The objection is also bolstered by the fact that the alternative's is, unlike the causalist conception of inference, apt to help explain how we are capable of doxastic agency via reasoning (see ch. 5).

oneself they must also be self-conscious. One could not appreciate why one is inferring when doing so, nor seem to appreciate the rationality of one's inferring when doing so, if one did not even know that one was inferring. It may thus seem like the Taking Condition cannot be motivated by the need to capture the intelligibility of inference since it appears to yield a conception of inference on which this cannot be done. Only a conception of inference apt to accommodate its self-consciousness could likewise accommodate its intelligibility.

A causalist could respond to the above line of thought as follows. I have argued that the modified causalist conception of inference is unmotivated. However, the advocate of the Taking Condition has a distinct means to motivate a version of the modified causalist conception to that discussed above. The Taking Condition, they may claim, is motivated by the need to capture the intelligibility of one's inferences. And the Taking Condition yields a version of the causalist conception of inference—a conception which is itself motivated by the need to capture the intelligibility of inference. To complete my case against the causalist conception of inference I need to explain why the causalist conception cannot be motivated in this way.

Consider first the suggestion that the Taking Condition should be endorsed to capture that when one infers one does not merely do so self-consciously but does so in awareness of *why* one is doing so. The Taking Condition might be thought to capture this since it entails that when one, say, infers  $q$  from  $p$  one accepts  $q$  *because* one takes it that  $q$  follows from  $p$ . But the suggestion in question in fact faces a problem much like the first problem raised for the causalist above (sec. 2.8). On the view now in question when one infers  $q$  from  $p$  one does so *because* one takes it that  $q$  follows from  $p$  where this is in turn true because one's taking it that  $q$  follows from  $p$  plays an appropriate causal role in bringing about one's acceptance of  $q$ . The problem is that supposing that one's taking plays this role its doing so is a matter which may be entirely beyond one's ken. The causalist lacks grounds to insist that the de facto appropriate causal involvement of one's taking it that  $q$  follows from  $p$  will necessarily result in one's having knowledge apt to secure the intelligibility of one's inference of a sort which one would lack if one did not come to accept  $q$  in the way in question. The causalist in question has it that when one infers  $q$  from  $p$  one's acceptance of  $q$  will be

appropriately caused by one's taking it that  $q$  follows from  $p$ . But they lack grounds to insist that one's taking it that  $q$  follows from  $p$  playing this role will guarantee awareness of its doing so such that one will thereby know that one accepted  $q$  because (one has it) it follows from  $p$  where this is a matter of one's taking having played an appropriate causal role.

The other way in which it was suggested that one's inferences might need to be intelligible to oneself was that they might need to be at least apparently rational. When one infers  $q$  from  $p$ , the thought was, one does not merely do so self-consciously. One's doing so will be at least apparently rational. As seen, it might be thought that endorsement of the Taking Condition is needed to accommodate this. In fact, however, the Taking Condition cannot be motivated by the alleged need to capture the apparent rationality of one's inferences. On the view in question inferring  $q$  from  $p$  requires taking it that  $q$  follows from  $p$  such that one who infers  $q$  from  $p$  might thereby seem to be in a position to appeal to the alleged fact that  $q$  follows from  $p$  in order to explain why it is rational to infer  $q$  from  $p$ . But even supposing that this is correct it does not speak in favour of the Taking Condition. To see this, consider a distinct view on which to infer  $q$  from  $p$  one must believe that  $q$  follows from  $p$  but on which one's taking it that  $q$  follows from  $p$  need not play a causal role in one's coming to accept  $q$  (see sec. 4.8). An advocate of this view can explain why one's inferences must be at least apparently rational in the very same way in which the advocate of the Taking Condition does—via appeal to the required taking—and can do so without insisting that when one infers the one's taking must play a causal role in one's doing so. The need to capture that one's inferences are not just self-conscious but apparently rational does not suffice to motivate the claim that to infer  $q$  from  $p$  one's taking it that  $q$  follows from  $p$  must play a *causal role* in one's inference as advocates of the Taking Condition have it. At most it motivates the claim that to infer  $q$  from  $p$  one must in fact take it that  $q$  follows from  $p$ .

## 2.11. Conclusion

On the causalist conception of inference for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to yield acceptance of  $q$  in “the right way”. Advocates of

the Taking Condition have it that the obtaining of such a causal connection requires the appropriate causal involvement of one's taking it that  $q$  follows from  $p$ . The causalist conception is the orthodox conception of inference. On reflection, however, it can be seen that it cannot be the right conception to endorse when it comes to giving an account of inference of the sort of concern here—when it comes to giving an account of self-conscious inference of the sort apt to occur as a constituent of active and conscious reasoning. This is because the causalist cannot accommodate the self-consciousness of inference without undermining the causalist conception's motivation. If one's acceptance of the causalist conception of inference is motivated by its allegedly yielding the best explanation of why relevant cases fail to be cases of inference then one must also appeal to the self-consciousness of inference in one's account. Doing so, however, renders one's appeal to causation theoretically redundant. One's acceptance of the causalist conception might instead be motivated via one's motivation of the Taking Condition, the latter being seemingly motivated by the need to capture inference's intelligibility. On reflection, however, it can be seen that the Taking Condition is no help when it comes to accommodating one's awareness of why one is inferring when doing so and not needed when it comes to accommodating one's apparent awareness of the rationality of one's inferences. The Taking Condition cannot be motivated by the need to accommodate the intelligibility of inference, and thus nor can the causalist conception itself.

### 3. Constitutive accounts of inference

**Chapter abstract** The causalist conception of inference has been seen to be either unable to accommodate the self-consciousness of inference or unmotivated. We can thus move on to consideration of accounts of inference which reject the causalist conception. On what I label ‘constitutive’ accounts what it takes for an inference to occur is specified in terms of what inferring constitutively involves rather than in terms of how inferences, or their constituent occurrences, are caused. One such account is offered by Valaris (2014, 2016a, 2016b, 2016c). According to Valaris there can be constitutive relations between beliefs. Believing  $p$  and that  $q$  follows from  $p$  can constitute believing  $q$ . Similarly, believing  $q$  on the basis of believed  $p$  can constitute believing that  $q$  follows from  $p$ . I argue that what Valaris offers fails to yield an account of inference of the kind of concern: he at most characterises a kind of potentially non-conscious belief fixation. On Rödl’s (2013) account, meanwhile, to infer  $q$  from  $p$  one must judge  $p$  and then judge  $q$  on the basis of  $p$ . When one does so one’s inference is identical to one’s judgment of  $q$ . And the latter judgment is also *token identical* to one’s knowledge of one’s judging and to one’s (at least apparently) realising that  $p$  is sufficient grounds to judge that  $q$ . Rödl’s account is too narrow in scope, ruling out the possibility of inference without judgment, and involves unnecessary commitment to a revisionary ontology. Despite their shortcomings I will build on both Valaris’ and on Rödl’s accounts in offering my own.

#### 3.1. An alternative to the causalist conception

The task at hand is to say what it is for an inference to occur. What is it for an inference from  $p$  to  $q$  to occur? Inference of the sort of concern is conscious inference of the sort apt to occur as a constituent of active and self-conscious reasoning and apt to be appealed to in order to explain how reasoning is a means to exercise doxastic agency. Such inferences are not just the inferences in a more permissive sense which happen to be self-conscious. Inference of the sort of concern is a distinctive kind of conscious and self-conscious occurrence.

On the orthodox causalist conception of inference for an inference from  $p$  to  $q$  to occur is just for acceptance of  $p$  to cause acceptance of  $q$  in the right way. On this conception, to say what it is for an inference from  $p$  to  $q$  to occur what this right way is just needs specifying. It has become apparent that this conception is either unable to accommodate the self-consciousness of inference or is unmotivated. We thus need to move on to a consideration of what I will label ‘constitutive’ accounts of inference. On such accounts what it takes for an inference to occur is not to be specified in terms of the occurrence of the right kind of causation. Instead, what it takes for an inference to occur is specified in terms of what an occurrence must *constitutively* involve in order to be an inference. On such accounts, whether an occurrence is an inference is determined by whether it constitutively involves appropriate belief or knowledge.

### 3.2. Valaris’ account

Valaris is in agreement with me that when one accepts  $p$  and then  $q$  the matter of whether an inference occurs is not settled by whether one’s acceptance of  $q$  is appropriately caused by one’s acceptance of  $p$  or by one’s acceptance of  $p$  along with one’s taking it that  $q$  follows from  $p$ . He accordingly aims to offer what amounts to an account of inference on which beliefs about what follows from what, rather than playing a causal role in inference, as advocates of the Taking Condition have it, play a constitutive role (see Valaris 2014, 2016b, 2016a, 2016c).

It is clear that Valaris aims to give an account of inference, despite his being somewhat equivocal on the matter of what his target phenomenon is. He says that his concern is *reasoning*, where reasoning “is a personal-level, conscious activity” (2014, p. 105). But he also has it that an inference is just “a piece of reasoning” (2016c, p. 1). Furthermore, in his description of his examples it is clear that it is inference which is of concern. Valaris, for instance, asks us to consider a subject who knows that if Socrates is human then he is mortal and that Socrates is human. He suggests that



[w]e naturally think that there is a cognitive act—albeit a rather trivial one, in this particular example—that the subject can perform in order to get to know [that] Socrates is mortal (2016b, p. 1).

The act of the kind in question, it seems clear, is an inference. And it is this act which Valaris states is his target phenomenon. He even sums up his view by saying that

inferring a conclusion from a set of propositions may simply consist in taking it that the conclusion follows from these propositions” (2016a, p. 895).

He also quotes what he acknowledges is an explicit discussion of inference in Ryle (2009a, pp. 302–3) as an instance of someone being in agreement with him in one respect on the nature of the act of concern (2016c, p. 2). It is clear that Valaris’ concern is what it is to make an inference where inference is understood as a kind of conscious act, as it is here.

Despite inference’s being Valaris’ target phenomenon his focus is on the possibility of their being constitutive relations between beliefs. His account is as follows. If one believes both  $p$  and that  $q$  follows from  $p$  then “barring inattention or irrationality” one thereby believes  $q$  (2014, p. 110). If one believes  $p$  and believes that  $q$  follows from  $p$ , the suggestion is, then there need be nothing left for one to do in order for one to qualify as believing  $q$ . Believing that  $p$  and that  $q$  follows from  $p$  can *constitute* believing  $q$ . As Valaris (2014, p. 110 note 12) acknowledges, a similar suggestion is made by Hieronymi when she suggests that “if you take certain reasons to show that  $p$ , you therein believe that  $p$ ” (2006, p. 51). The same is so of Ryle when he claims that

[w]e do not first see an implication and then go on to draw a conclusion, any more than we first accept the solution of an anagram and then go on to solve it. (2009a, p. 279)

Valaris illustrates the point with the case of Fred:

If Fred believes that there is a bloody knife in the garden, and he comes to realize that it follows from this that the butler did it, his reasoning is done—no further process is required to get him to believe that the butler did it. (2014, p. 110)

If Fred is fully rational and attentive then his believing that there is a bloody knife in the garden and that it follows from this that the butler did it suffices for his believing that the butler did it. And this is in turn so because his believing that there is a bloody knife in the garden and that it follows from this that the butler did it can constitute his believing that the butler did it.

Valaris sees his account as incomplete at this stage. He has said that believing  $p$  and that  $q$  follows from  $p$  can constitute believing that  $q$ . But he notes that belief that  $q$  follows from  $p$  is belief of a sort which is itself typically the product of reasoning or inference (2014, pp. 109–10). Fred's belief that it follows from the fact that the bloody knife is in the garden that the butler did it, for instance, is likely the product of reasoning or inference. To have an account of inference on which an inference can occur without a prior inference needing to have occurred Valaris appeals to the suggestion that we have a capacity to come to believe things on the basis of other things we believe by conforming to general patterns—one which operates “largely unconsciously” and which is at least to some extent hard-wired, but which may also be to some extent the product of training (2014, p. 113). The suggestion that we have such a capacity, Valaris (2014, p. 113) notes, is not something which his opponents want to deny.

Valaris then considers a case in which one exercises the above capacity. Suppose that one does so and comes to believe  $q$  on the basis of believed  $p$ . Valaris claims that one will thereby believe that one believes  $q$  for reason  $p$ . In general, Valaris suggests, if one believes  $q$  on the basis of believed  $p$  then one must believe that one believes  $q$  for reason  $p$  (2014, p. 113). Valaris has it that

it is a necessary truth about us that we are often (though not always) in a position to know what we believe and why, without observation or inference. (2014, p. 114)

Just as one can know that one believes  $p$  without observation or inference one can know that one believes  $q$  for reason  $p$  without observation or inference. Valaris takes this to support his suggestion that if one believes  $q$  for reason  $p$  one must also believe that  $p$  is one's reason to believe  $q$ . The fact that we can know what we believe without observation or inference is typically taken to be reflected in the fact that in order to answer the question of whether one believes

$p$  all one need typically do is ask oneself whether  $p$  (G. Evans 1982, p. 225; see also Barnett 2016; Boyle 2011b; Moran 2001; Valaris 2011). Valaris suggests that the fact that to believe  $q$  on the basis of  $p$  one must believe that  $p$  is one's reason to believe  $q$  is similarly reflected in the fact that in order to answer the question of whether one believes  $q$  for reason  $p$  one typically need only attend to the matter of whether  $p$  is in fact sufficient reason to believe  $q$  (2014, p. 114).<sup>37</sup>

So Valaris has it that we have a capacity to form beliefs on the basis of reasons and that if one believes  $q$  for reason  $p$  one must believe that one believes  $q$  for reason  $p$ . He also has it that if one believes that  $p$  is one's reason for believing  $q$  then one must believe that  $q$  follows from  $p$  (2014, p. 116). One would not believe that  $p$  is one's reason to believe  $q$ , the thought is, if one did not believe that  $q$  follows from  $p$ . If one believes that one believes  $q$  for reason  $p$  then one must believe that  $p$  and  $q$  stand in an appropriate inferential relation—that  $q$  follows from  $p$ —otherwise one would not see  $p$  as sufficient reason to believe  $q$  and would not have it that  $p$  is one's reason to believe  $q$  accordingly.

Assuming all of the above is correct, believing  $q$  on the basis of  $p$  entails believing that  $q$  follows from  $p$ . To believe  $q$  on the basis of  $p$  one must believe that one believes  $q$  for reason  $p$ . And to believe that one believes  $q$  for reason  $p$  one must believe that  $q$  follows from  $p$ . Valaris has it that this is so because believing that one believes  $q$  on the basis of  $p$  constitutes believing that  $q$  follows from  $p$ . If one believes that one believes  $q$  on the basis of  $p$  then one's doing so will constitute one's believing that  $q$  follows from  $p$ . If one comes to believe  $q$  on the basis of  $p$  by exercising the above capacity to form beliefs on the basis of others then one will thereby believe that one believes  $q$  for reason  $p$ . And one's believing that one believes  $q$  for reason  $p$  will constitute one's believing that  $q$  follows from  $p$ .

---

<sup>37</sup> It is far from clear that the “transparency” considerations which Valaris cites here support anything as strong as his claim that *necessarily* if one believe  $q$  for reason  $p$  then one believes that one believes  $q$  for reason  $p$ . Such considerations support the claim that one can *typically* know whether one believes  $p$  without observation or inference. They likewise support the claim that one can *typically* know that one believes  $q$  for reason  $p$  when one does without observation or inference. But Valaris' claim is that to believe  $q$  for reason  $p$  one *must* believe that one does so.

### 3.3. Valaris' failure to characterise inference

In sum, Valaris' view is that if one believes that  $p$  and that  $q$  follows from  $p$  then, assuming one is rational and attentive, one will thereby believe  $q$ . In the rational and attentive believing  $p$  and that  $q$  follows from  $p$  constitutes believing  $q$ . In addition, if one believes  $q$  on the basis of  $p$  then one will thereby believe that one believes  $q$  for reason  $p$ . And believing that one believes  $q$  for reason  $p$  will constitute one's believing that  $q$  follows from  $p$ . There can in this way be constitutive relations between one's beliefs. Rosa (forthcoming) objects to the particular way in which this suggestion is fleshed out by Valaris (2016b). A more pressing issue for Valaris, however, is that of whether he in fact manages to give an account of his target phenomenon—whether what he says amounts to an account of inference.

Valaris (2016c) characterises his account as one of a kind of acceptance. But by 'acceptance' he does not mean a kind of conscious act and a kind of conscious occurrence as I have been using the term. 'Acceptance' as used by Valaris denotes a kind of state. What Valaris offers is an account of what kinds of constitutive relations there can be between states of acceptance. This is clear both in the account offered by Valaris, as seen, as well as in how he sums it up. At one point, for instance, he sums up by stating that his concern is reasoning where

[r]easoning is a sophisticated species of the broader phenomenon of believing for a reason. (2014, p. 105)

He characterises his central claim as being that inference "is not to be identified with a causal process" and is rather a kind of *state* (2014, p. 103). As Valaris puts it elsewhere, on his view

inferring a conclusion from a set of propositions may simply consist in taking it that the conclusion follows from these propositions (2016a, p. 895).

Valaris' account is an account of what constitutive relations can obtain between beliefs—between particular kinds of states of acceptance. The problem for Valaris is then clear. It seems that he simply fails to give an account of his target phenomenon and of the phenomenon which concerns us here. What Valaris

says does not amount to offering an account of conscious inference. Yet, as seen, his aim is to do just that.

In response to the above objection Valaris stresses that on his view “inferring a conclusion from a set of premises is a special way of accepting that conclusion” and that “accepting something can, in a perfectly colloquial sense, be a kind of change—if what you accept now is something you did not accept before” (2016c, p. 11). To object that Valaris fails to give an account of inference understood as a kind of conscious occurrence, Valaris insists, is to “equivocate on the notion of change” (2016c, p. 11). But in offering his account Valaris does not use ‘acceptance’ to denote such changes. He uses it to denote states of belief, not events of coming to believe. The only way to make sense of Valaris is then as follows. Valaris must intend for what he says regarding states of belief and the constitutive relations they can stand in to yield an account of inference understood as a kind of occurrence. What he says of states of acceptance and their constitutive relations is intended to yield an account of a distinctive kind of event of acceptance. And Valaris identifies events of the latter sort with inferences. On Valaris’ view inference is a kind of acceptance were ‘acceptance’ means “a kind of change”—that of the sort which occurs when “what you accept now is something you did not accept before”. On the view which results to infer  $q$  from  $p$  is to come to accept  $q$  (in the stative sense) in a particular way. It is to come to accept  $q$  either by coming to believe  $p$  and that  $q$  follows from  $p$  or by coming to believe  $q$  on the basis believed  $p$  (via an exercise of one’s capacity to form beliefs in accordance with general patterns). On Valaris’ view to believe  $p$  and that  $q$  follows from  $p$  can constitute believing that  $q$ . Coming to believe  $p$  and that  $q$  follows from  $p$  can thus constitute coming to believe  $q$ . To come to believe  $q$  in this way, on Valaris’ account, is to infer  $q$  from  $p$ . On Valaris’ view it is also the case that believing  $q$  on the basis of believed  $p$  constitutes believing that  $q$  follows from  $p$ . Coming to believe  $q$  on the basis of believed  $p$  will thus constitute coming to believe that  $q$  follows from  $p$ . And to do so, on Valaris’ account, is to infer  $q$  from  $p$ . Valaris’ view is that inference is a kind of acceptance in that to infer  $q$  from  $p$  is to come to believe  $q$  in one of these ways.

Given the above, it is clear that what Valaris says fails to yield an account of conscious inference of the sort of concern. On Valaris' account inference is just a kind of belief fixation—that which either just is or which constitutively involves the acquisition of beliefs about what follows from what. As has already been seen, the occurrence of such belief fixation is neither necessary nor sufficient for the occurrence of inference where inference is understood as a kind of conscious act. One can infer  $q$  from  $p$  without coming to believe  $q$  or that  $q$  follows from  $p$ , as one might when going through familiar reasoning, when inferring from suppositions, or when inferring from assumptions made for practical purposes (see sec. 1.3). This is precluded by Valaris' account, and nothing which he says suggests how it might be accommodated.<sup>38</sup> One can also come to believe  $q$  on the basis of believed  $p$  without inferring  $q$  from  $p$  (see sec. 1.4). Doing so suffices for accepting  $q$  in a way which amounts to inferring  $q$  from  $p$  on Valaris' account. The latter point is even exemplified by one of Valaris' own examples. Valaris asks us to

suppose that while watching a football game I come to believe that a goal was just scored. It is plausible that this belief of mine is based on more basic perceptual knowledge, as well as my knowledge of the rules of the game. Nevertheless, to the extent that I never explicitly considered the matter, this is not a case of reasoning in our sense. (2014, p. 105)

By Valaris' own admission the occurrence of acceptance of the sort which Valaris characterises and which on his account suffice for the occurrence of an inference does not suffice for the occurrence of conscious inference of the sort which concerns us here, and which is Valaris' target phenomenon. As Valaris himself puts it,

[r]easoning ... is a personal-level, conscious activity; and it seems clear that many of our beliefs are based on other beliefs without such explicit reasoning. (2014, p. 105)

---

<sup>38</sup> Valaris is thus mistaken in claiming that although he focuses entirely on the case in which inference from  $p$  to  $q$  yields belief in  $q$  consideration of further cases is “unlikely [to] introduce significant new issues” (2014, p. 103 note 3).

What Valaris offers neither amounts to nor yields an account of inference of the sort of concern and which is his target phenomenon. That is not to say, however, that what Valaris offers is not of value, nor that what he offers is not of value when it comes to the task of saying what it takes for an inference to occur. Valaris' central point is that we can see beliefs as constitutively related. Believing  $p$  and that  $q$  follows from  $p$ , for instance, can *constitute* believing  $q$ . Similarly, believing  $q$  and that one believes  $q$  for reason  $p$  can *constitute* believing that  $q$  follows from  $p$ . Beliefs need not be conceived of as "isolated atoms" (2016c, p. 2). This yields a way to see one's beliefs as apt to do more than just play a causal role in the acquisition of further beliefs and knowledge. Furthermore, if believing  $q$  in a certain way can constitute believing that  $q$  follows from  $p$  then coming to believe  $q$  in a certain way can likewise constitute coming to believe that  $q$  follows from  $p$ . Coming to believe  $q$  and coming to believe that  $q$  follows from  $p$  need not be seen as requiring *both* the occurrence of an event which is one's coming to believe  $q$  *and* the occurrence of a *distinct* event which is one's coming to believe that  $q$  follows from  $p$ . I will ultimately build on this suggestion by offering an account of inference on which inference from  $p$  to  $q$  can yield belief that  $q$  and can involve belief that  $q$  follows from  $p$  without the latter belief's playing a causal role in one's inference (see sec. 5.2). To infer  $q$  from  $p$  can *just be* to *both* come to believe  $q$  *and* to come to appreciate that  $q$  follows from  $p$ . I will thus be in complete agreement with Valaris when he suggests that

it is a mistake to think that beliefs about what follows from what can be involved in reasoning only as links in a causal chain. Indeed ... such beliefs play a constitutive, rather than causal, role in reasoning. (2014, p. 103)

Furthermore, I will suggest, seeing how beliefs about what follows from what can be constitutively involved in inference allows us to shed light on the way in which inference is self-conscious and to become clear on what kind of occurrences inferences are.

### 3.4. Rödl's account

Rödl is also in agreement with me on that when one accepts  $p$  and then  $q$  the matter of whether an inference occurs is not determined by how one's acceptance of  $q$  is caused. Unlike with Valaris, what Rödl (2013) offers also does amount to a characterisation of a conscious act of the sort of concern. According to Rödl to infer  $q$  from  $p$  one must first judge that  $p$  and then judge that  $q$  (2013, pp. 213–4). He is thus in agreement with me that to infer  $q$  from  $p$  one must first consciously accept  $p$  and then  $q$ . In doing the latter, on Rödl's account, one must also judge  $q$  on the basis of  $p$ . Rödl identifies one's inference with the latter act of acceptance—with the event which is one's judging  $q$  on the basis of  $p$ . When one infers  $q$  from  $p$ , he claims, one's inference is the unity of the following three elements:

(Element 1—the “Dependence of judgments”) One judges  $q$  on the basis of  $p$ .

(Element 2—the “Consciousness of dependence”) One's consciousness of one's judging  $q$  on the basis of  $p$ .

(Element 3—the “Consciousness of justification”) One's recognising  $p$ , the ground on which one judges  $q$ , as sufficient ground to do so. (2013, p. 215)

Rödl says that Element 1 *just is* Element 3. When one infers  $q$  from  $p$  one's judging  $q$  on the basis of  $p$ —which for Rödl is one's inferring  $q$  from  $p$ —just is one's recognising  $p$  as sufficient ground to judge that  $q$ . When we see this, Rödl claims, it becomes evident that Element 2 *just is* Element 3. One's inference from  $p$  to  $q$  is a self-conscious occurrence of a kind apt to yield knowledge only because Element 2 is Element 3—because one's being conscious of judging  $q$  on the basis of  $p$  just is one's recognising  $p$  as sufficient grounds to do so. One can come to know  $q$  by inferring it from  $p$ , the thought is, because to do so is to be conscious of one's judging  $q$  on the basis of sufficient grounds.

On the view that results to infer  $q$  from  $p$  one must judge  $p$  and then judge  $q$  on the basis of  $p$ , one's inference being the latter act of acceptance. And to do that is to recognise  $p$  as sufficient ground to judge  $q$ , which is in turn to be conscious of judging  $q$  on the basis of  $p$ . It is because inference just is the unity



of these three elements, in the sense that they are all literally *identical*, that inference is apt to yield knowledge.

Although along the right lines Rödl's account is too narrow in scope and involves commitment to an unmotivated revisionary ontology. It is too narrow in that it rules out that one might infer  $q$  from  $p$  without judging that  $q$ , as one might when inferring from a supposition or assumption made for practical purposes. It similarly rules out that one might infer  $q$  from  $p$  without first judging that  $p$ . On Rödl's account inferring  $q$  from  $p$  requires judging  $p$  and then  $q$  and the account thus rules out the possibility of inference without judgment. It might also be thought to be too narrow in applying only to successful inferences—to potentially knowledge transmitting inference. But it is in fact clear that Rödl takes himself to offer the means to say what it is for an inference to occur via reference to successful inference. He takes himself to provide the means to say what it takes for an inference to occur—successful or otherwise—where the correctness of the account which results is made clear via consideration of the case of successful inference.<sup>39</sup> Strictly speaking then, Rödl's claim is that to successfully infer  $q$  from  $p$  is to knowledgeably judge  $p$  and then judge  $q$  where one's inference is identical to one's judging that  $q$ , to one's consciousness of one's doing so, and to one's recognising  $p$  as sufficient grounds to do so. What it is to infer  $q$  from  $p$  can then be characterised similarly. To infer  $q$  from  $p$  is to judge  $p$  and then judge  $q$  where one's inference is identical to one's judging that  $q$ , to one's consciousness of one's doing so, and to one's *apparently* recognising  $p$  as sufficient grounds to do so. For such an inference to be successful one's judgment that  $p$  must be knowledgeable and one's apparent recognition of the sufficiency of one's grounds amount to genuine awareness.

Rödl's view is thus that when one infers  $q$  from  $p$  the following are *token identical*:

- One's judging  $q$  on basis  $p$ .

---

<sup>39</sup> This same strategy is employed in the inference case by Brewer (1995). What it is to make an inference is characterised via reference to what it is to infer successfully. The strategy is commonly applied elsewhere too, perhaps most familiarly in the case of perceptual experience. Brewer (2011), Martin (2002) and McDowell (1994), amongst many others, aim to say what it is to have a perceptual experience in terms of what it is to have a veridical perceptual experience.

- One's consciousness of judging  $q$  on basis  $p$ .
- One's recognising (or apparently recognising)  $p$  as sufficient grounds to judge that  $q$ .

This is naturally read as meaning that when one infers  $q$  from  $p$  the *event* which is one's judging  $q$  on basis  $p$  is token identical to one's *knowledge* of one's doing so and to one's *knowledge* or *belief* that  $p$  is sufficient grounds to judge that  $q$ . Rödl is committed to a token identification of judgments with beliefs and knowledge.

The obvious worry with Rödl's suggestion is that it runs afoul of a simple Leibniz's Law argument. Take a case in which one infers  $q$  from  $p$ . On Rödl's view one's inference is the event which is one's consciously judging that  $q$ . And this event is token identical to one's knowledge of one's inferring and to one's knowledge or one's mere belief that  $p$  is sufficient grounds to judge that  $q$ . It is both natural and standard, however, to have it that the latter are *states*. As Shah and Velleman put it

[a] judgment is a cognitive mental act of affirming a proposition ... It is an act because it involves occurrently presenting a proposition, or putting it forward in the mind .... A belief, by contrast, is a mental state of representing a proposition as true, a cognitive attitude rather than a cognitive act. (2005, p. 503)

On Rödl's view when one infers  $q$  from  $p$  one's inference is the event which is one's judging  $q$ . Belief, meanwhile, is standardly taken to be a kind of state, as is knowledge. It is even typical to cite knowledge as a paradigm state when introducing the state/occurrence distinction (e.g. Mourelatos 1978; Vendler 1957). Furthermore, events and states have distinct characteristics such that they cannot be identified. How the distinction between states and occurrences is to be drawn and how, within the latter, the distinction between events and processes is to be drawn is a matter of controversy (e.g. Crowther 2011; O'Shaughnessy 2000, pp. 42–9; Steward 2013; Stout 1997). But it is nonetheless clear that there is a distinction to be made here. As Soteriou explains, it is typical to appeal to the different ways in which events, processes and states "fill time" in order to draw an ontological distinction between them. States, for instance, "*obtain* over, and throughout, intervals of time and at times". Events and

processes, meanwhile, “occur/happen/unfold over time and/or at times” (2013, p. 27).

So events occur and do not obtain at times. States, meanwhile, obtain and do not occur at times. Events happen and do not obtain at times. States, meanwhile, obtain but do not happen at times. It similarly seems clear that whilst judgment is a kind of occurrence—specifically, a kind of event—belief and knowledge are kinds of states. Beliefs and knowledge obtain rather than occur or happen. It seems that such differences can be appealed to in order to establish that judgments cannot be token identical to states. Rödl’s apparent suggestion that judging  $q$  on basis  $p$  can be token identical to knowledge of one’s doing so and to knowledge or to mere belief that  $p$  is sufficient grounds to judge that  $q$  thus cannot be correct.

A different way to understand Rödl would be to read his talk of recognising  $p$  as sufficient grounds to judge that  $q$  as reference to what is naturally understood as a kind of *occurrence* rather than as a kind of *state*. Thus read Rödl’s suggestion is that when one infers  $q$  from  $p$  the event which is one’s judging  $q$  is identical to the event which is one’s coming to believe or to know that  $p$  is sufficient grounds to do so. As noted in the discussion of Valaris above, I will ultimately suggest that such events can be token identical and appeal to this in my account of inference. Nonetheless, understanding Rödl in the way in question fails to render unproblematic his suggestion that when one infers  $q$  from  $p$  the event which is one’s inference is also token identical to one’s knowledge of one’s inference. It seems clear that in suggesting this Rödl is committing to a problematic token identification of states of knowledge with the occurrences of which they are knowledge.

Rödl’s response to the above worry will simply be to deny that belief and knowledge are kinds of states rather than kinds of occurrences. Rödl (2007, pp. 75–9) even explicitly argues against the natural suggestion that belief is a kind of state. It cannot be, he claims, because beliefs are not such that “any limit of [their] duration is accidental” (2007, p. 78). The claim is that beliefs are such that they persist unless something happens to stop them from persisting and that this is an essential characteristic of beliefs rather than a contingent fact about them. This is claimed to be inconsistent with their being states. Believing

and likewise knowing, rather than being kinds of states, are on Rödl's view rather kinds of "temporally unlimited act[s]" (2007, p. 78).

It is unclear, however, why we should think that beliefs differ from states in that the former and not the latter persist non-accidentally. The air's smelling of jasmine, for example, is a paradigm state (Mourelatos 1978, p. 201).<sup>40</sup> If Rödl is to maintain that the latter is a state then he is committed to saying that any limit of its duration is accidental. But it is unclear why we should think that this is so—why the air's smelling of jasmine is not a state of a sort which will persist unless something happens to stop it from doing so. In sum, it is far from clear that there is really any contrast of the kind which Rödl aims to draw attention to in order to argue that beliefs are not states. He even cites as evidence for his view the fact that we do not say of beliefs that they happen at particular times (2007, p. 78). This would support Rödl's suggestion that beliefs are not states only if we are typically happy to say of other states that they happen at particular times. We do not, however, do any such thing. It seems clear that Rödl imposes demanding conditions on what would be required in order for belief and knowledge to be kinds of states where it is far from clear or just evidently false that these conditions are satisfied by other paradigm states.

At the very least then, Rödl's view is revisionary in requiring a rejection of the natural and orthodox view that belief and knowledge are kinds of states. Rödl does not offer a compelling case against the orthodoxy. Nor does he show that there is any need to accept his revisionary ontology. Rödl claims that events of inferring are identical to knowledge of their occurrence on the basis of insistence that inference needs to be essentially self-conscious in order to play its epistemological role. I am in agreement that inference is not merely contingently self-conscious. What I will deny is that we need to accept Rödl's revisionary ontology in order to accommodate this fact.

Rödl's account is presented only briefly amongst a critical discussion of Kitcher (2011). So it would not be fair to be overly critical. The account I will ultimately offer builds on Rödl's. But it also allows us to see how inference is

---

<sup>40</sup> Other typical examples include desiring, loving, hating and wanting (Vendler 1957). Rödl might deny that such alleged psychological examples are really examples of states on analogous grounds to those on which he denies that belief is a state.

essentially self-conscious without denying that knowledge and belief are kinds of states. I will offer an account on which the relation between inference, self-consciousness and awareness of inferential connections is neither causal nor one of identity. On the view to be offered, the relation between inference and knowledge of its occurrence is one of constitutive dependence rather than identity. The relation between inference and awareness of inferential connection is likewise not one of identity. Rather, when inference yields such awareness one's inference is identical to the event which is one's coming to be aware. This view preserves the intimate relation between inference and knowledge in such a way that it can be seen how inference plays its epistemological role and its role in doxastic agency whilst avoiding commitment to Rödl's revisionary ontology.

### **3.5. Conclusion**

Valaris offers an account of what constitutive relations can obtain between beliefs. Rödl, meanwhile, offers an account of inference understood as a kind of conscious occurrence, but one which excludes the possibility of inference without judgment. His view also comes with unnecessary revisionary commitments. In offering my own account I will aim to build on the insights from both Valaris and Rödl. On the account I will offer an inference from  $p$  to  $q$  can be *token identical* to an event of coming to appreciate that  $q$  and that  $q$  follows from  $p$ . As both Rödl and Valaris suggests, inference can involve appreciation of inferential relations without such appreciation having to play a causal role in inference. How this can be true of conscious occurrences of the kind of concern will become clear.

## 4. Inference and self-consciousness as its “silent partner”

**Chapter abstract**     So far it has become clear that to make an inference one must accept *that which one infers from* and then *that which one infers*, but need not believe or come to believe either. Neither causalist accounts nor the constitutive accounts considered can accommodate this adequately. In this chapter I construct an alternative account of inference via consideration of Koziolok’s (2017) causalist account. On Koziolok’s view inference is self-conscious because when one infers one’s inference causes awareness of one’s inference. Much like the causalist conception of inference, I argue, Koziolok’s account of inference’s *self-consciousness* is either inadequate or unmotivated. The inadequacy of Koziolok’s account of the self-consciousness of inference motivates an alternative on which the self-consciousness of inference is not a causal product of one’s inferences but a cotemporaneous “silent partner” (O’Shaughnessy 2000, p. 106). When one infers one knows that one is inferring where this knowledge of one’s inference constitutively depends on the occurrence which is one’s inference. This suggestion yields a way to say what it is for an inference to occur without appeal to what causes or is causally involved in inference. For an inference to occur one must simply accept what one infers from and then what one infers where in doing so one knows that one is inferring, the latter knowledge being a silent partner to one’s inference. This in turn tells us that inference is a kind of conscious *event*—one’s acceptance of *that which one is inferring*. I aim to shed light on what is involved in knowing that one is inferring via appeal to the need for appropriate taking beliefs in order for an inference one makes to be intelligible to one. I also explain how this account avoids the difficulties faced by those considered so far, how the account can be developed, and what it suggests concerning the place of taking beliefs in inference.

### 4.1. What we have so far

So far, the following has become clear. Inference is a kind of conscious and self-conscious occurrence. It is also a kind of conscious and self-conscious act—something which we consciously and self-consciously *do*. Furthermore, when

one infers there is something which one *infers from*, and consciously accepts, and also something which one *infers*, and consciously accepts. That is not to say that to infer one must believe what one infers from, nor that one must believe or come to believe what one infers. Self-conscious reasoners can infer from assumptions made for practical purposes and from mere suppositions. In such cases we infer from propositions which we do not believe and which we may even take to be false.

No account considered so far has been able to adequately capture all of the above. Causalist accounts either fail to accommodate the self-consciousness of inference or are unmotivated. If inferring  $q$  from  $p$  is merely a matter of acceptance of  $q$  being appropriately caused by acceptance of  $p$  then whether one is inferring when one does so can be completely beyond one's ken. A causalist might respond by building the self-consciousness of inference into their conception of inference. For them to do so, however, is for them to undermine their motivation for endorsement of a causalist conception (see ch. 2). The constitutive accounts considered, meanwhile, fail to accommodate that inference need not yield belief in that which one infers, that inference need not yield belief in the obtaining of inferential relations, and that inference need not involve judgment (see ch. 3). With the insights and shortcomings of these views now clear, it is time to develop an alternative. I will do so via consideration of what I take to be the shortcomings of the causalist view presented by Koziolok (2017).

#### **4.2. Koziolok's causalist account of inference**

Koziolok (2017) begins by offering an account of what it is to make an “epistemically successful inference”, by which he means a knowledge transmitting inference—an inference which yields inferential knowledge of what one infers on the basis of knowledge of what one infers from. According to Koziolok, for an epistemically successful inference from  $p$  to  $q$  to occur is for *knowledge of  $p$*  to cause *inferential knowledge of  $q$*  where the causation in question is a kind of “causation in virtue of rationalization”. The kind of causation in question when one makes a knowledge transmitting inference is also a species of “theoretical rational causation”, where this is a kind of causation appeal to

which is apt to “explain knowledge as knowledge” (2017, p. 10).<sup>41</sup> As Koziolak explains

[f]or causation to explain knowledge as knowledge is for it to explain that knowledge in such a way as to reveal how it is that the subject knows what she knows. (2017, p. 10)

Koziolak thus presents a developed account of the form suggested by Brewer (1995). Like Koziolak, Brewer claims that in order to give an adequate account of inference we must make sense of “causation in virtue of rationalization”. Brewer aims to shed light on this distinctive kind of causation via reflection on basic perceptual knowledge, whilst Koziolak aims to do the same via reflection on paradigm cases of knowledge transmitting inference.

Koziolak’s suggestion here can be summed up as follows. When one makes an epistemically successful inference one’s knowledge of what one infers from causes one’s knowledge of what one infers in a distinctive way: such that the knowledge one acquires qualifies as having been caused by prior knowledge in a way which makes it *inferential knowledge* and such that it was caused in this way *in virtue of* the fact that one’s prior knowledge was apt to play this causal role.<sup>42</sup>

Koziolak has no reductive ambitions in offering the above account of epistemically successful inference. His approach is rather “knowledge-first” in the following way. His aim is to say what it is to make an inference in general via appeal to a non-reductive account of knowledge transmitting inference. Another thing to note is that there appears to be a respect in which Koziolak’s concerns are broader than mine, and that as a result his target phenomenon is not the same as mine. Koziolak counts an inference as having occurred in all cases in which one comes to believe something on the basis of something else one believes (2017, p. 2). As seen, however, one can come to believe something

---

<sup>41</sup> Causation in virtue of rationalization is a broader category than theoretical rational causation, since the former is alleged to be apt to yield both knowledge and *action*, whilst the latter is alleged to be apt to yield only knowledge (see Antony 1989; Brewer 1995).

<sup>42</sup> Koziolak, like Brewer, is thus committed to an anti-Davidsonian, non-extensionalist conception of causation on which causes can have their effects in virtue of some of their properties and not others. Crane (1995) argues that this conception of causation has anti-physicalist consequences. I am inclined to say that it is the conception of causation which is at fault. Part of what I see myself as doing in what follows is revealing how we can make do just fine without it.



on the basis of something else without consciously inferring (see sec. 1.3 above). As Moran puts it, such belief fixation can occur “below the threshold of consciousness” (2001, p. 110). It appears that Koziolok’s concern, unlike mine, thus cannot be inference where by ‘inference’ we mean a kind of conscious occurrence (see sec. 1.3 above).<sup>43</sup> Despite this, Koziolok has insisted in correspondence that his concern is inference understood as a conscious phenomenon. He is mistaken in thinking that he can have it both that inference is necessarily conscious and that one’s coming to believe for a reason suffices for the occurrence of an inference. We can come to believe for reasons without consciously inferring.<sup>44</sup>

Koziolok is in agreement with me that inference is self-conscious. This is something which he both aims to capture with his account of inference and which his account makes essential appeal to (see 2017, pp. 13–19). He also claims to share Rödl’s (2007) account of self-consciousness (Koziolok 2017, p. 13 note 19). On the account which he endorses

for a state to be self-conscious is for you to know that you’re in it, where you know that you’re in it in a particular, and special, way, namely, by being in it. Similarly, for an act to be self-conscious is for its performance to involve coming to know that you’ve performed it, where you come to know that you’ve performed it precisely by performing it. (2017, p. 13)

When one infers one knows that one is inferring *by inferring*, just as, Koziolok suggests, one might know that one is in pain when one is simply *by being in pain*. When one makes an inference one’s doing so suffices for one’s knowing that one is doing so—one need do nothing other than infer itself in order to know that one is inferring. To put it another way, when one infers one does so self-consciously, where one’s knowledge of one’s doing so is a manifestation of

---

<sup>43</sup> There is also a respect in which Koziolok’s focus appears to be narrower than mine. He is unable to count cases in which one infers from mere suppositions or from assumptions for practical purposes as inference, a fact which he recognises (Koziolok 2017, p. 2 note 2).

<sup>44</sup> Koziolok could easily ensure that inference is a conscious phenomenon on his account by stating it terms of conscious *acceptance* rather than in terms of *knowledge* and *belief*. He could have it that for an epistemically successful inference from *p* to *q* to occur is for knowledgeable *acceptance* of *p* to cause knowledgeable *acceptance* of *q* where the latter knowledge is inferential and where the causation in question is a kind of “causation in virtue of rationalization”.

one's capacity for self-consciousness, and where all one need do in order to manifest this capacity in this way is infer.<sup>45</sup>

In further developing his account of the self-consciousness of inference Koziolk uses as a constraint a generalisation of Sellar's suggestion, itself from Kant, that

[e]ven our consciousness of what is going on in our own mind is a conceptual response which must be distinguished from that which evokes the response (Sellars 1967, p. 280; in Koziolk 2017, p. 6).

As Koziolk (2017, p. 6 note 10) notes, Sellars tentatively follows Kant in restricting this point. Sellars sees the point as limited to the case of sensory experience and holds back from applying it to the likes of acts like judgment and inference. Koziolk, meanwhile, has it that no such restriction of the claim is warranted. *All knowledge*, he claims, must be a conceptual response to that which it is knowledge of. This is understood as requiring that all knowledge must be distinguished from that which it is knowledge of. It is likewise understood as requiring that all knowledge is such that its subject matter's being as it is known to be does not depend upon its being known to be so. From this Koziolk concludes that, even when it comes to knowledge of our own minds, what is known must always be causally prior to one's knowledge. As he himself puts it

even when it comes to our knowledge of our own minds, what is known is both causally and epistemically prior to our knowledge of it. (2017, p. 6)

Koziolk concludes that when one infers one's knowledge of one's inferring—that is, the self-consciousness of one's inference—must be a *causal product* of one's inference. He has it that when one infers one's inference causes one's self-consciousness “in a special way”, namely, “the way distinctive of self-consciousness” (2017, p. 14). One's inference “self-consciously rationally causes” one's self-consciousness where such causation is a further species of theoretical rational causation (2017, p. 15).

---

<sup>45</sup> Koziolk is thus in agreement with O'Shaughnessy (2000, p. 105) in having it that, in a sense, there is no *way* in which we acquire such knowledge.

Here it becomes clear that Koziolak is mistaken in claiming to share Rödl's account of self-consciousness. According to Rödl (2013, see also 2007, p. ix), when one infers one's inference is self-conscious because the act which is one's inferring is *identical* to one's knowledge that one is inferring (see sec. 3.4 above). On Koziolak's account, meanwhile, one's knowledge of one's inference and one's inference itself are distinct. The former is a causal product of the latter.<sup>46</sup>

For completeness, we can now present Koziolak's account of inference in full. For a knowledge transmitting inference from  $p$  to  $q$  to occur is for knowledge of  $p$  to cause inferential knowledge of  $q$ . When this occurs the causation in question is a kind of theoretical rational causation. When such an inference occurs it also self-consciously rationally causes knowledge of one's inference where the causation in question here is a further kind of theoretical rational causation.

The above makes it possible to say what a "good inference" is where, for Koziolak, although a good inference must yield belief it may not yield knowledge, as one may make a good inference and yet infer from a belief which fails to constitute knowledge. For a good inference from  $p$  to  $q$  to occur is *either* for a knowledge transmitting inference to occur *or* for belief in  $p$  to cause belief in  $q$  in such a way that had one known  $p$  one would have acquired inferential knowledge of  $q$ . Such inferences involve "inferential rational causation"—yet another species of theoretical rational causation—where this is causation that is *either* actually productive of inferential knowledge *or* potentially productive of inferential knowledge. Again, when such an inference occurs it also self-consciously rationally causes knowledge of one's inference where the causation in question is a further kind of theoretical rational causation. However, this only amounts to a partially successful exercise of one's capacity for self-consciousness. One's inference self-consciously rationally causes knowledge of one's inference and of the goodness of one's inference, but mere belief that one's inference is knowledge transmitting (Koziolak 2017, sec. 6).

---

<sup>46</sup> Correspondence with Koziolak has made it evident that we disagree about how Rödl intends to be understood. I stand by my interpretation. When one infers, on Rödl's view, the occurrence which is one's inference is literally *token identical* to one's knowledge of one's inference. It is because of this that Rödl must deny that one's knowledge of one's inference is a state (see sec. 3.4 above).

It is also now possible to say what a bad or fallacious inference is, where such inferences are not even potentially productive of knowledge. For a bad inference from  $p$  to  $q$  to occur is for belief in  $p$  to cause belief in  $q$  where in question is “non-rational causation” rather than inferential rational causation or any other species of causation in virtue of rationalization where this nonetheless self-consciously rationally causes knowledge of one’s inference. And again, this only amounts to a partially successful exercise of one’s capacity for self-consciousness. One’s inference self-consciously rationally causes knowledge of one’s inference, but mere belief in the goodness of one’s inference and that one’s inference is knowledge transmitting (Koziolk 2017, sec. 8). In such a case, although one’s belief in  $q$  is not the product of any kind of causation in virtue of rationalization, it is still correct to conceive of oneself as inferring. This is because even though no rational causation occurs one still conceives of oneself as inferring. As Koziolk (2017, p. 19) puts it, in the case of bad inferences one’s conception of oneself and of what one is doing makes all the difference.<sup>47</sup>

As seen (sec. 2.8-2.10), causalist accounts either fail to accommodate the self-consciousness of inference in the self-conscious or are unmotivated. Koziolk’s account is no exception. However, considering the particular issues faced by Koziolk’s account will help to make evident a new way to go about constructing an account of inference.

### **4.3. How Koziolk fails to accommodate the self-consciousness of inference**

Koziolk faces a serious difficulty when it comes to his account of the *self-consciousness* of inference. The difficulty is much like that raised by Kitcher (2011) and Rödl (2013) for perceptual models of self-consciousness. The picture attacked by Kitcher and Rödl is one on which the self-consciousness of a given self-conscious mental act is a matter of perception of the act. It is also a picture

---

<sup>47</sup> Koziolk (2017, sec. 8) is equally happy to be understood as denying that such bad inferences are really inferences whatsoever. On this way of understanding Koziolk his view is a revisionary one. Real inferences involve inferential rational causation. One may simply want to label bad inferences inferences in virtue of the fact that we mistakenly take them to be real inferences.

on which the very same perceived act could have occurred unperceived and still have been the kind of mental act that it in fact is. The self-consciousness of any given self-conscious mental act, on this picture, is the contingent result of what O'Shaughnessy calls a "quasi-perceptual faculty of 'inner sense'" (2000, p. 105).

In objecting to the above picture Rödl states that self-consciousness is  
a consciousness that is of her who is so conscious, not *per accidens*, but  
in virtue of being the kind of consciousness it is. (2013, p. 217)

Perceptual models of the self-consciousness of inference cannot capture this, Rödl claims, because

[a]n act of receptivity, as such, does not satisfy this definition. The consciousness of unity in inference is no act of inner sense, but a synthesis. (2013, p. 217)

Kitcher similarly objects that when a self-conscious mental act occurs one's knowledge of one's act

is not a perception of an act that could be separate from it. It is an indissoluble component of the self-conscious act itself. (2011, p. 130)

Applied to the case of inference, the objection appears to be as follows. To infer self-consciously is not to infer *and* to perceive one's inference via a quasi-perceptual faculty of inner sense. If that were right then the self-consciousness of any given inference would either be something which contingently accompanies one's inference or something which is a contingent product of one's inference. This cannot be right, the objection goes, because inferences are self-conscious in virtue of being the kinds of acts they are and are *necessarily* self-conscious.<sup>48</sup>

Koziol's account faces a difficulty of the same form. Inference is necessarily self-conscious. As seen, this is a claim which an adequate account of inference needs to accommodate. But it is unclear how Koziol could be entitled to this claim given that on his account the self-consciousness of any

---

<sup>48</sup> Rödl (2013, p. 214) suggests that if we deny this then we deny ourselves the resources needed to distinguish inference from non-inference. Ultimately, I will agree with Rödl that appeal to the self-consciousness of inference is needed in order to *correctly* distinguish inference from non-inference. But to say this now would not be to offer a dialectically effective objection to Koziol's suggestion that we can instead do so via appeal to an appropriate kind of causation in virtue of rationalization.

given inference is a causal product of the inference's occurrence.<sup>49</sup> Why could an inference not occur and yet *fail* to yield knowledge of its occurrence? Suppose, for instance that knowledge of  $p$  causes knowledge of  $q$  where the causation in question here is inferential rational causation and thus that one qualifies as gaining inferential knowledge of  $q$ . On Koziol's account of self-consciousness this inference will be self-conscious if it self-consciously rationally causes knowledge of its occurrence. But why could knowledge of  $p$  not cause knowledge of  $q$  via inferential rational causation and this *fail* to yield knowledge of one's inference via self-conscious rational causation? Why is it *impossible* for knowledge of  $p$  to cause knowledge of  $q$  via inferential rational causation and then fail to yield knowledge of one's inference via self-conscious rational causation? Causes, after all, have their effects contingently. Koziol's account seems to leave us without the resources to accommodate his claim that inference *must* be self-conscious, rather than the weaker claim that inference is self-conscious all being well—so long as one's inferences self-consciously rationally cause awareness of themselves.

A natural way to respond on Koziol's behalf at this point would be say that he can simply use the grounds on which we have it that inference in the self-conscious is necessarily self-conscious as grounds to deny that it is possible for an inference to occur and yet fail to yield knowledge of its occurrence via self-conscious rational causation. This is not an ad hoc reply, the response can go, because we have independent grounds to deny that non-self-conscious inference is possible (c.f. sec. 2.9 above). Koziol could insist that whether an inference from  $p$  to  $q$  occurs can depend *both* on whether belief in  $p$  yields belief in  $q$  via inferential rational causation *and* on whether this yields knowledge of one's inference via self-conscious rational causation.<sup>50</sup> Alternatively, he could

---

<sup>49</sup> In correspondence Koziol has suggested that he would be happy to drop the claim in light of the objection to follow. With Koziol thus understood the objection to him can be read as restricted in the same way in which my objections to the other views considered have been restricted. Koziol fails to give an adequate account of inference insofar as by 'inference' we mean the distinctive kind of conscious and self-conscious occurrence. That is the phenomenon which we need an account of if we are to understand how doxastic agency is possible and how to exercise such agency is to exploit one's self-consciousness (see sec. 1.6 above).

<sup>50</sup> In the case of bad inferences whether one is inferring will depend *only* on whether knowledge of one's inference is produced via self-conscious rational causation.

insist that whether inferential rational causation occurs in a case depends in part on whether what occurs will yield knowledge of one's inference via self-conscious rational causation.<sup>51</sup> Whether one is presently inferring can depend on what *will* happen.<sup>52</sup>

It has already been seen why a response of the above form on Koziolok's behalf will not work. For Koziolok to respond in the above way is for him to appeal to knowledge of inference in his account of inference. Supposing such a move is legitimate, as it must be if Koziolok is to make it, knowledge of inference can be appealed to in order to explain why any case of non-inference fails to amount to a case of inference. But if that can be done then Koziolok's endorsement of a causalist conception of inference cannot be motivated by the need to explain why cases of non-inference fail to amount to cases of inference (see sec. 2.9 above).

What is important for present purposes, however, is that Koziolok cannot accommodate the claim that inference is self-conscious without undermining the motivation for his causalist account of inference's self-consciousness. As seen, Koziolok could modify his view in response to the observation that it fails to accommodate the self-consciousness of inference. He could say that when belief in  $p$  causes belief in  $q$  via inferential rational causation this only suffices for the occurrence of an inference if it yields knowledge of one's inference via self-conscious rational causation. The result is a conjunctive analysis of inference. For an inference to occur belief in  $p$  must cause belief in  $q$  via inferential rational causation *and* this occurring must yield knowledge of the occurrence of one's inference via self-conscious rational causation.<sup>53</sup> Alternatively, Koziolok could say that when belief in  $p$  causes belief in  $q$  whether inferential rational causation is in question and in turn whether an inference is in question depends on whether what occurs causes knowledge of inference via

---

<sup>51</sup> Again, in the case of bad inferences there will be no inferential rational causation and whether one is inferring will just depend on whether knowledge of one's inference is produced via self-conscious rational causation.

<sup>52</sup> A claim which Koziolok is already committed to in the case of bad inference. Whether one is presently making a bad inference depends on whether what is now occurring will cause knowledge of one's inferring.

<sup>53</sup> Unless a bad inference is in question, in which case all that is required is that knowledge of one's inference be caused via self-conscious rational causation.

self-conscious rational causation.<sup>54</sup> On both of these suggestions whether an inference occurs can depend on whether knowledge of one's inference *will* be produced via self-conscious rational causation.

On the account Koziolok is forced to accept when some inference causes knowledge of its occurrence via self-conscious rational causation the occurrence will qualify as an inference only given that it plays this causal role. A dilemma for Koziolok then arises when we ask '*When* such an inference occurs, is it *already* an inference?' If Koziolok's answer is 'Yes' then the above problem is not evaded. If when one infers one's inference *already* constitutes an inference and *then* causes knowledge of its occurrence, then when one makes the inference everything which needs to be the case in order for one to qualify as inferring is already the case. But if that is right then it should be possible for one's inference to fail to yield knowledge of one's inference. There is no need for it to do so if one is to qualify as having inferred. Answering 'No', meanwhile, is problematic by Koziolok's own standards. To respond in this way is to say that when one infers it is not already settled that one is doing so. Whether one will qualify as having inferred depends on whether what is occurring will cause knowledge of inference. Koziolok cannot maintain this without undermining the motivation for his account of inference.

Koziolok cannot maintain both that inference is self-conscious and his account of inference's self-consciousness unless he has it that whenever one infers whether one is inferring depends upon whether the occurrence in question will cause knowledge of inference. Analogously, we can maintain that necessarily whenever a lethal poisoning occurs a death will occur because we can maintain that whether any given occurrence qualifies as a lethal poisoning depends upon whether it will cause a death (c.f. Davidson 2001c, pp. 177–8). Crucially, for Koziolok to take this line would be for him to abandon his claim that when one infers one's knowledge of one's inference must be a "response"

---

<sup>54</sup> Again, unless a bad inference is in question, in which case whether an inference occurs depends upon whether knowledge of one's inference be caused via self-conscious rational causation.



to one's inferences as he understands that claim.<sup>55</sup> It would be for him to accept that when one infers one's inference's status as an inference depends upon one's knowledge of one's inference. As Koziolk understands it, the claim that when one infers one's knowledge of one's inference must be a response to one's inference entails that one's inference must be distinct from and must not depend for its status as an inference upon one's knowledge of one's inference.

Koziolk motivates his claim that we must see the self-consciousness of one's inferences as a causal product of one's inferences by suggesting that we must do so in order to accommodate that whenever one infers one's knowledge of one's inference is a response to one's inference. The resultant view, we have seen, fails to accommodate that inference is necessarily self-conscious. In response, we have seen, Koziolk could maintain that inference is necessarily self-conscious by insisting that any given inference depends for its status as an inference on whether it causes knowledge of inference. To say that, however, would be to abandon the claim that when one infers one's knowledge of one's inference is a conceptual response to one's inference. It would instead be to have it that when one infers one's inference depends for its status as an inference upon one's awareness of one's inference. And for Koziolk to accept that would be for him to abandon the claim the need to capture which motivates his account.

Some may suspect that it is not just the fact that inference is necessarily self-conscious which Koziolk is unable to adequately accommodate. Many will see the self-consciousness of one's acts and one's acts themselves as being more intimately connected than Koziolk allows. When a self-conscious act occurs, the thought is, one is not simply acting and, in addition, aware that one is doing so. One's awareness of what one is doing is part of what makes the act in question the kind of act and the kind of exercise of agency which it is (e.g. Anscombe 1957; Kitcher 2011; O'Shaughnessy 2000; Rödl 2007; Soteriou

---

<sup>55</sup> Strictly speaking, Koziolk is already committed to abandoning this claim given his account of bad inferences. It is because of this that Koziolk is happy to be understood as denying that bad inferences are real inferences (see note 47 above).

2013).<sup>56</sup> Koziolak's account of self-consciousness does not allow for such an intimate relation between one's acts and one's awareness of them.<sup>57</sup>

#### 4.4. Self-consciousness as a “silent partner”

As seen, Koziolak has it that the self-consciousness of inference must be a causal product of one's inferences. This is itself taken to be correct on the basis of Koziolak's understanding of Sellars' remark that

[e]ven our consciousness of what is going on in our own mind is a conceptual response which must be distinguished from that which evokes the response (Sellars 1967, p. 280; in Koziolak 2017, p. 6).

Koziolak has it that one's knowledge of one's inferences must be a “response” to one's inferences where this is taken to entail that one's inferences must be causally prior to one's knowledge of them. In general, he assumes, what is known must always be causally prior to one's knowledge. The account which results is either ill-equipped to accommodate the self-consciousness of inference or unmotivated. We should seek an alternative account of inference's self-consciousness accordingly. We should seek an account of inference's self-consciousness on which it is not a matter of inference's being such that it always causes knowledge of its occurrence. Inference, we can instead say, is self-conscious, which is to say that one must know that one is inferring *when* one infers and thus that the self-consciousness of one's inferences is not a causal product of one's inferences. In developing this account we need not be constrained by Sellars' claim as it is understood by Koziolak.

What is now needed is a way of understanding how inference is self-conscious without seeing the self-consciousness of any given inference as a causal product of that inference. Seeing Kitcher's (2011, p. 130) suggestion that the self-consciousness of an inference is “an indissoluble component of the self-conscious act itself” as insufficiently precise, Rödl (2013) aims to provide what

---

<sup>56</sup> We will see a manifestation of this suggestion in the account to follow.

<sup>57</sup> Koziolak (2017, p. 7 note 12) implies that as he sees it there is no need to accommodate the point in question. In fact, his use of Sellars' remark amounts to a rejection of it. Koziolak cites Strawson (2003) and Setiya (2013)—sceptics about the possibility of doxastic agency—in support of this stance in the case of inference. Such scepticism will itself be revealed to rest on a mistake and to rest on a mistake that can be corrected via appeal to an improved account of self-conscious inference (see chapter 5 below).

is needed here. As seen, according to Rödl inference is self-conscious because when one makes an inference one's inference is *identical* to one's awareness of that very inference (see sec. 3.4 above). However, although it amounts to a more precise suggestion than Kitcher's, Rödl's account faces difficulties. According to Rödl we ought to see inferences, understood as conscious occurrences, as being token identical to *knowledge of* conscious occurrences. It is recommended that we endorse an identification of conscious *occurrences* with what it is both typical and natural to understand as a kind of *state*. As seen, this suggestion either runs afoul of a simple Leibniz's Law argument or forces us to accept a revisionary ontology (again, see sec. 3.4 above). The suggestion to be developed reveals that we need do no such thing and thus reveals Rödl's revisionary ontology to be unmotivated.

A better way of understanding how inference is self-conscious which avoids the difficulties faced by both Koziolak's and Rödl's suggestions is available. On this suggestion inference is self-conscious where one's awareness of one's inference when one infers is neither identical to one's inference, nor a causal product of one's inference, but a *cotemporaneous state*. One knows that one is inferring *when* and *because* one infers whenever one infers. This is not because one's inference causes or is identical to knowledge of one's inference, but because when one infers a cotemporaneous state which is one's knowledge of one's inferring *constitutively depends* on the conscious occurrence which is one's inference. On this suggestion, when a self-conscious reasoner makes an inference this suffices for knowledge of the inference where this is understood as a matter of the obtaining of a cotemporaneous state of awareness which constitutively depends on the conscious occurrence which is the inference.<sup>58</sup> This view captures what appears right in Koziolak's suggestion that when one infers one knows that one is inferring *by inferring*. When one makes an inference one's doing so suffices for one's knowing that one is doing so. One need do nothing other than infer in order to know that one is inferring. This is because

---

<sup>58</sup> Someone sympathetic to Steward's (1997) concerns with the very idea of a *token state* might worry that this suggestion commits me to the existence of such states. In fact though, the suggestion need not be read that way. Someone convinced by Steward could understand my claim to be that whenever one infers one's doing so is constitutive of one's being such that one knows that one is inferring.

to be capable of inferring is to be such that when one makes an inference a cotemporaneous state of awareness of one's inference will constitutively depend on the conscious occurrence which is one's inference. To be capable of inferring is to be such that for one to make an inference is for one to do so self-consciously.

Some may worry that this suggestion faces a vicious regress.<sup>59</sup> On the suggestion inference is self-conscious, but an inference and the self-consciousness of the inference are distinct. Whenever one infers a state of knowledge of one's inference which is distinct from one's inference obtains. It might be thought that this will require that whenever one infers one must *judge* that one is inferring. One must judge that one is inferring whenever one infers in order to know that one is inferring when one does so. However, it might be thought that judgment must be self-conscious on the same grounds that inference has been taken to be. If this is right then self-consciously inferring requires self-consciously judging that one is inferring when one infers. Now the regress arises. It begins with the suggestion here that inference is self-conscious and that to infer self-consciously a state of knowledge of one's inference must obtain. Furthermore, it is suggested, for a state of knowledge of one's inference to obtain one must judge that one is inferring when one infers. However, as inference is self-conscious, judgment is self-conscious by the same token. To judge self-consciously a state of knowledge of one's judgment must obtain. Furthermore, for a state of knowledge of one's judgment to obtain one must judge that one is judging when one judges. In order to make a self-conscious inference one must infer *and* simultaneously judge that one is inferring *and* judge that one is judging that one is inferring. And given judgment is self-conscious this latter judgment must itself be accompanied by a further judgment. The suggestion that inference and judgment are self-conscious *and* that knowing that one is inferring or judging requires judging that one is doing so clearly generates a vicious regress.

---

<sup>59</sup> The regress worry was raised by an anonymous reviewer for my (forthcoming) and is responded to there as here. It is also briefly discussed and set aside by O'Shaughnessy (2000, p. 107), whilst a regress of the form in question moves Sartre to a characterisation of self-consciousness in terms of "an immediate, non-cognitive relation of the self to itself" (2003, pp. 8–9).

Coming to know does not always require conscious judgment in a mediating role (O’Shaughnessy 2000, p. 107; Soteriou 2013, p. 231) (see also sec. 1.3 above). We can and do come to know things without doing so by making conscious judgments. Given this, the lesson from the above regress is a simple one. We should deny that in order to come to be aware of what one is doing when one makes an inference one must make a judgment about what one is doing. One’s knowledge of what one is doing when one is inferring must instead be understood as “a silent or non-conscious partner” to one’s inferring (O’Shaughnessy 2000, p. 106). One knows that one is inferring when one does so, that is, where being aware of what one is doing in this way does not require the occurrence of a further event of judging—a further occurrence and a further occupant of the stream of consciousness. When one infers self-consciously, on this proposal, no occurrence other than that which is one’s conscious inference is required in order for one to know that one is inferring. One’s knowledge that one is inferring constitutively depends on the conscious occurrence which is one’s inference, obviating the need for a further occurrence which yields one’s self-knowledge.

#### **4.5. My account of inference**

We can now, at long last, begin to answer the question posed in chapter 1: ‘What is it for an inference from  $p$  to  $q$  to occur?’ This much is already clear. One must first consciously accept  $p$ . But one need not do so in a distinctive way. One might simply judge that  $p$  and  $q$  occur to one in a way such that one qualifies as having inferred  $q$  from  $p$  without having accepted  $p$  in a way any different to how one would have done in a case in which one did not infer. In order for a step to qualify as a constituent of a process of waltzing, for instance, one might have it that the step must be taken in a distinctive way—in order to waltz. One need not analogously accept  $p$  in order to do anything, or in any distinctive way whatsoever, in order to infer  $q$  from  $p$ . As well as consciously accepting  $p$  one must consciously accept  $q$  if an inference from  $p$  to  $q$  is to occur. One must first accept that which one infers from and then accept that which one infers. But accepting  $p$  and then  $q$  does not suffice for having inferred  $q$  from  $p$ . Nor does doing so where one’s acceptance of  $p$  causes one’s acceptance of  $q$ . Acceptance

of  $p$  could cause acceptance of  $q$  and one fail to qualify as having inferred  $q$  from  $p$ . A natural response to this point, as seen, is to appeal to the “right kind of causation”. That is, it is natural to respond that for an inference from  $p$  to  $q$  to occur is for acceptance of  $p$  to cause acceptance of  $q$  in the right way, where what this right way is is specified either reductively or otherwise. But as has also been seen this approach cannot yield an adequate and motivated account of inference. To take the approach in question is to endorse the causalist conception of inference. And that conception is either unable to accommodate the self-consciousness of inference or unmotivated (see sec. 2.8 - 2.10 above).

For a self-conscious reasoner to infer  $q$  from  $p$ , then, one must first accept  $p$  and then accept  $q$ . But appeal to whether or not an appropriate causal connection obtains between one’s acceptance of  $p$  and  $q$  is not the way to distinguish inference from non-inference. Furthermore, we have now seen that the claim that inference is self-conscious ought to be understood as a matter of the obtaining of a cotemporaneous state of awareness of one’s inference whenever one infers—a “silent partner” to one’s inferring which constitutively depends on the occurrence of one’s inference. This last observation makes available a way to specify what it is for an inference to occur without any sort of appeal to the right kind of causation. We can instead appeal to the self-consciousness of inference *directly* in order to say what it takes for an inference to occur.<sup>60</sup> We can, that is, say what it is for an inference from  $p$  to  $q$  to occur as follows. For an inference from  $p$  to  $q$  to occur is for one to accept  $p$  and then accept  $q$  where in doing so one is aware that one is inferring  $q$  from  $p$ , one’s knowledge that one is inferring  $q$  from  $p$  being a “silent partner” to the conscious occurrence which is one’s inference as that is understood above. To put it another way, for an inference from  $p$  to  $q$  to occur is for one to accept  $p$  and then  $q$  where in doing so one manifests one’s awareness that one is inferring  $q$  from  $p$ . When one infers, on this suggestion, one does so self-consciously where we can appeal to this “silent partner” to one’s inference directly, rather than to the occurrence of the right kind of causation, in order to say what it is for an inference to occur.

---

<sup>60</sup> A suggestion made by Rödl (2013, p. 214) in passing (see also sec. 2.9 above).

The above strategy renders the account of inference offered circular. In saying what it takes for an inference to occur an appeal to knowledge is made. But it is knowledge *of inference* which is appealed to. One might worry about the substantiveness and informativeness of an account which moves in such a small circle. The account offered, however, does have substantive consequences and is apt to be put to significant theoretical work accordingly. The account can also be developed further. Light can also be shone on what is involved in knowing that one is inferring rather than, say, guessing or supposing, even if what it is to have such knowledge cannot be non-circularly specified. We have failed to alight upon a reductive account of inference. The account to be offered here, although articulated in terms of knowledge of inference and thus circular, allows us to see how inference is self-conscious, how inference plays its role in reasoning, and how reasoning is a means to exercise doxastic agency. We should, I suggest, endorse the non-reductive conception offered accordingly.

What has been said already yields the resources to resolve the ontological issue from chapter 1. We began by asking the question ‘What is an inference?’ Every available answer to this question appeared to be problematic in its own way and we were left in a quandary (see sec. 1.8). We thus instead turned to the question ‘What is it for an inference to occur?’ We now have an answer to the latter question. Furthermore, this answer yields an answer to the ontological question.

As seen, to infer  $q$  from  $p$  one must accept  $p$  and then accept  $q$ . But one need not accept  $p$  in a distinctive way. And when one infers  $q$  from  $p$  one must know that one is doing so where this knowledge is a “silent partner” to one’s inference as understood above. Furthermore, one could not know that one is inferring  $q$  from  $p$  when accepting  $p$ . That would require that when one accepts  $p$  in a non-distinctive way and will happen to go on to accept  $q$  and qualify as having inferred  $q$  from  $p$  one must know already that one is inferring  $q$  from  $p$ . It is mysterious how one could have such knowledge when all one is doing is accepting  $p$  in a non-distinctive way. After all, when one accepts  $p$  it may not yet be settled whether one will go on to accept  $q$ , or whether one will get interrupted or otherwise fail to go on to accept  $q$  and qualify as having inferred  $q$  from  $p$ . One would have to know that one is inferring  $q$  from  $p$  when accepting  $p$  in a

non-distinctive way where this knowledge depends upon whether one will happen to go on to accept  $q$  in an appropriate way.

It would not be similarly mysterious to suggest that in a case in which one infers  $q$  from  $p$  one knows that one is inferring *when one accepts  $q$* . Such knowledge would not need to depend on what one will happen to go on to do. Suppose that one accepts  $p$ , then accepts  $q$  and knows that one is inferring  $q$  from  $p$  when one does the latter. Nothing further needs to occur for the matter of whether an inference from  $p$  to  $q$  has occurred to be settled. It seems unproblematic to say that in a case in which one infers  $q$  from  $p$  one knows that one is inferring *when one accepts  $q$*  unlike how it would be mysterious to suggest that one already knows that one is inferring when one accepts  $p$ . We can thus say the following. To infer  $q$  from  $p$  one must accept  $p$  *in any old way* and then accept  $q$  where one must know that one is inferring  $q$  from  $p$  *when one accepts  $q$*  and not before. The latter knowledge is a “silent partner” to one’s inferring where this is in turn a matter of its constitutively depending on the conscious occurrence which is one’s inference.

It can now be seen how this yields an answer to the question ‘What is an inference?’ When one infers  $q$  from  $p$  one knows that one is doing so *when one accepts  $q$*  and *not before*. Furthermore, this knowledge of what one is doing constitutively depends on the conscious occurrence which is one’s inference. And the conscious occurrence which is contemporaneous with this knowledge of one’s inference is the *conscious event* which is one’s acceptance of  $q$ . It is the latter conscious occurrence which thus ought to be identified with one’s inference. When one infers  $q$  from  $p$  it is the event which is one’s acceptance of  $q$  which ought to be identified with one’s inference. Generalising, when one makes an inference one’s inference is identical to the *conscious event* which is one’s acceptance of *that which one infers*.

It is now possible to state my account of inference. To infer  $q$  from  $p$  one must first consciously accept  $p$  and then consciously accept  $q$  where one’s knowledge that one is inferring  $q$  from  $p$  is a silent partner to one’s inference. This is in turn a matter of a contemporaneous state of awareness that one is inferring  $q$  from  $p$  constitutively depending on the occurrence which is one’s inference. And the occurrence which is one’s inference is one’s accepting  $q$ —



the *conscious event* which is one's acceptance of *that which one infers*. My answer to the question 'What is it for an inference from  $p$  to  $q$  to occur?' is thus the following:

**(What it is for an inference to occur)** For one to infer  $q$  from  $p$  is for one to consciously accept  $p$  and then consciously accept  $q$  where a cotemporaneous state of knowledge that one is inferring  $q$  from  $p$  constitutively depends on the conscious occurrence which is one's acceptance of  $q$ .

As seen, this yields the following answer to the question 'What is an inference?':

**(What an inference is)** An inference is a kind of conscious *event* and conscious act. Namely, it is the conscious event and act of acceptance which is one's acceptance of *that which one infers*.

#### 4.6. Our ontological quandary resolved

When we first considered the above answer to the question 'What is an inference?' the suggestion appeared to be problematic (see sec. 1.8 above). If we say that when one infers  $q$  from  $p$  one's inference is the conscious event which is one's acceptance of  $q$  then we appear to be committing ourselves to a view on which the occurrence of a *single event* suffices for the occurrence of an inference. However, we also have it that *two events* must occur in order for an inference to occur (see sec. 1.4 above). One must first accept that which one infers from and then accept that which one infers. How could the occurrence of a single event suffice for the occurrence of an inference and yet the occurrence of two events be required for an inference to occur? It has now become clear that when one infers one's inference really is identical to the single event which is one's accepting that which one infers. The above difficulty for this suggestion thus must be resolved.

In response to the above difficulty we can say the following. When an inference from  $p$  to  $q$  occurs one's inference is identical to the event which is one's acceptance of  $q$ . In addition, for one to qualify as inferring  $q$  from  $p$  in accepting  $q$  one must *first* accept  $p$  and *then* accept  $q$ . An event of acceptance's qualifying as an inference requires that it be preceded by acceptance of that which one is inferring from. Although an inference from  $p$  to  $q$  is a single event

it is an event of a kind which can only occur after a prior event. When an inference occurs one's inference is thus *identical to a single event* despite the fact that the occurrence of two distinct consecutive events is required for an inference to qualify as having occurred. The mistake was to think that an inference's being identical to a single event entails that the occurrence of a single event suffices for the occurrence of an inference.

This point is one which we can see exemplified in various other cases (see Ryle 2009a, pp. 149–53). For instance, an event which is one's arriving might be identified with the event which is one's stepping into the room. But for this event to count as an arrival a prior journey must have taken place and thus prior occurrences have occurred. That one's arriving is a single event does not entail that the occurrence of a single event suffices for one's having arrived. Similarly, an event which is one's checkmating might be identified with the event which is one's placing one's bishop. But for this event to count as checkmating a game must have started and prior moves have been made. That one's checkmating is a single event does not entail that the occurrence of single event suffices for one's having checkmated. It becomes particularly plausible to see the inference case as having this structure when we consider a case in which one *concludes reasoning* by inferring. One might conclude one's reasoning about whether  $q$  by inferring  $q$ . But for the event which is one's concluding one's reasoning to be an event of this kind prior reasoning must have taken place. That one's concluding is a single event does not entail that the occurrence of single event suffices for one's having concluded. And in the case in which one concludes *by inferring* it seems that the single event which is one's concluding just is the single event which is one's inferring (see sec. 5.2 below).

Because inferring  $q$  from  $p$  requires first accepting  $p$  and then accepting  $q$  we can see that the occurrence of an inference entails the occurrence of reasoning. To infer  $q$  from  $p$  one must accept  $p$  and then infer  $q$  and be aware that one is inferring  $q$  from  $p$  in doing so. And to do that is to engage in a minimal bit of reasoning. It is to engage in reasoning which may begin with one's acceptance of  $p$  and terminate with one's inferring  $q$ , although it may just as well be to engage in a minimal bit of reasoning which is itself part of a longer chain of reasoning. A natural way to express the point is as follows. Inference

is a kind of conscious and self-conscious *event* of a kind which always occurs as a constituent of a *process* of conscious reasoning. The occurrence of an inference entails the occurrence of at least a minimal bit of reasoning and thus entails the occurrence of a process of reasoning.

#### **4.7. The problem of deviant causation resolved**

The account of inference offered only gives us a circular answer to the question ‘What is it for an inference to occur?’ Despite this, it does yield a resolution of the ontological quandary with which we began. A further virtue of the account offered is that it does not face the problem of deviant causal chains. As seen, an account of inference must be able to distinguish cases in which one infers  $q$  from  $p$  from those in which one merely accepts  $p$  and then accepts  $q$ . A natural way to do so is to appeal to the need for an appropriate causal connection to obtain between one’s acceptance of  $p$  and acceptance of  $q$ . But as has also been seen the causal conception of inference that results either fails to accommodate the self-consciousness of inference or is unmotivated (see chapter 2 above). On the account proposed no such problematic appeal to “the right kind of causation” is needed in order to distinguish cases in which one infers from cases in which one does not. We can instead appeal to one’s knowledge of what one is doing when one infers directly. An act of acceptance which is an inferring is not to be distinguished from non-inferences by its having been caused in some specifiable appropriate way and rather by its amounting to a manifestation of one’s awareness that one is inferring.

#### **4.8. Taking beliefs and the intelligibility of inference**

What it is to infer has been specified via reference to inference’s self-consciousness. I do not intend to give a non-circular account of what it is to know that one is inferring and thus of what it is to infer. We can, nonetheless, shed light on what is involved in knowing oneself to be inferring. In doing so our guiding question can be the following: ‘What must be the case for one to know that one is inferring?’

As was noted earlier (sec. 2.5), when one infers one does not merely know that one is inferring. When one infers one knows that one is doing so in such a

way that one's making the inference in question is intelligible to oneself. As Peacocke puts it, the inferences we make are "rational transitions" rather than mere "blind leaps into the dark, inclinations to make transitions in thought that just grip and take over the thinker's rational self" (2003, p. 127). The need to capture this does not suffice to motivate acceptance of the Taking Condition. As seen, advocates of the condition have it that to infer  $q$  from  $p$  one must accept  $p$  and then  $q$  where one's acceptance of  $q$  is caused by one's taking it that an appropriate inferential relation obtains between  $p$  and  $q$  (see sec. 2.3 above). Acceptance of the condition does not suffice to explain how it is that our inferences are intelligible in the way in question (see sec. 2.10 above). Nor is acceptance of the condition required in order to accommodate this intelligibility (again, see sec. 2.10 above, a point which will come out further in what follows). But although reflection on what is needed for inference to be intelligible does not suffice to motivate the Taking Condition, it does suffice to motivate the claim that *when one infers  $q$  from  $p$  one must take it that an appropriate inferential relation obtains between  $p$  and  $q$* . The need to capture the respect in which one's inferences are intelligible to oneself does suffice to motivate the claim *when one infers  $q$  from  $p$  one must take it that  $q$  follows from  $p$* , where this is understood, as is natural, as a matter of inferring  $q$  from  $p$  requiring that one *believe* that  $q$  follows from  $p$  when one infers.<sup>61</sup>

Suppose, for instance, that one infers  $q$  from  $p$ . To do so is to accept  $q$  in a way such that one manifests awareness that one is inferring  $q$  from  $p$ , rather than, say, merely judging that  $q$  outright, supposing that  $q$ , or guessing that  $q$ . A simple and natural way to say what is required in order for one to know that one is inferring  $q$  here, rather performing some other act of acceptance, is as follows. To do so one must be aware, not just that one is accepting  $q$ , but that one is accepting  $q$  *because  $p$*  (one believes, or is supposing, or is assuming). Mere awareness that one is accepting  $q$  because  $p$  (one believes, or is supposing, or is assuming), however, would not suffice alone to render one's acceptance of  $q$

---

<sup>61</sup> As before, I assume that *if* one must have an appropriate taking belief in order to infer  $q$  from  $p$  then one must believe that  *$q$  follows from  $p$* . It remains true that the issue of *whether* one must have an appropriate taking belief in order to infer is prior to that of *what* taking beliefs one must have in order to infer.

intelligible. Knowing only that one is accepting  $q$  because  $p$  (one believes, or is supposing, or is assuming), assuming such a thing is possible, might perfectly well leave it completely opaque to one why it is  $q$  that one is accepting. If one is inferring  $q$  from  $p$  where one's doing so is intelligible to oneself then one must be aware that one is accepting  $q$  and aware that one is doing so because  $p$  (one believes, or is supposing, or is assuming) where  $q$  follows from  $p$  (one believes). For one's self-conscious act of acceptance of  $q$  to be intelligible to oneself as an inference, that is, one must take oneself to be accepting  $q$  because  $p$  (one believes, or is supposing, or is assuming) where one believes that  $q$  follows from  $p$ . One's belief that  $q$  follows from  $p$  is required such that it is intelligible from one's perspective why one would accept  $q$  in the way in question given that one believes that  $p$ , is supposing that  $p$ , or is assuming that  $p$ . Without belief that  $q$  follows from  $p$  knowledge that one is accepting  $q$  because  $p$  would not render an inference from  $p$  to  $q$  intelligible from the point of view of the inferrer.

We have arrived at the following suggestion. To infer  $q$  from  $p$  one must accept  $p$  and then  $q$  where in accepting  $q$  one manifests one's awareness that one is inferring  $q$  from  $p$ . To manifest such awareness one's act of accepting  $q$  must be intelligible to oneself as an inference from  $p$  to  $q$ , where this requires that one be aware, not just that one is accepting  $q$ , but that one is doing so because  $p$  (one believes, or is supposing, or is assuming). And one's accepting  $q$  in this way because  $p$  (one believes, or is supposing, or is assuming) will be intelligible to oneself whenever one makes such an inference because to infer  $q$  from  $p$  one must believe that  $q$  follows from  $p$ . To infer  $q$  from  $p$ , on the view which results, one must do so self-consciously and where one's doing so requires that one believe that  $q$  follows from  $p$  when one infers. Generalising, in order to make an inference one must believe that *what one is inferring* follows from *what one is inferring from* and must believe this *when one infers*.

On the view suggested it is a mistake to think that inference is "blind" in the following sense (c.f. sec. 2.4 above). Whenever one infers one does so self-consciously, manifesting awareness that one is inferring. Likewise, whenever one infers one is aware of one's grounds for doing so and is at least apparently aware that what one is inferring follows from what one is inferring from. It is

because of this that whenever one infers one's doing so is intelligible to oneself.<sup>62</sup>

The above does not demand that in order to infer  $q$  from  $p$  one's acceptance of  $q$  must be caused by one's acceptance of  $p$ , nor by one's belief that  $q$  follows from  $p$ . When one infers  $q$  from  $p$  one self-consciously accepts  $q$  where in doing so one is at least apparently aware that  $p$  is one's grounds for doing so. This, I suggest, is a matter of one's manifesting appropriate awareness in accepting  $q$  itself, rather than a matter of one's acceptance of  $q$  having been appropriately caused. It is not inconsistent with one's accepting  $q$  *because*  $p$  (one believes, or is supposing, or is assuming) in the sense in question that one's acceptance of  $q$  is in fact caused by one's prior acceptance of  $p$  and/or by prior belief that  $q$  follows from  $p$  (c.f. Williams 1970, pp. 141–2). The point here is just that we need not say that to infer  $q$  from  $p$  one's acceptance of  $p$  or taking it that  $q$  follows from  $p$  must cause one's acceptance of  $q$  in order to capture how it is that one's inferences are intelligible to oneself. One's inferences are intelligible to oneself because to infer  $q$  from  $p$  is to manifest awareness that one is accepting  $q$  because  $p$  (one believes, or is supposing, or is assuming) where one's accepting  $q$  *because*  $p$  in the relevant sense will be intelligible to oneself so long as one believes/is supposing/is assuming that  $p$  and so long as one believes that  $q$  follows from  $p$ . I am in fact happy to concede to the orthodoxy that whenever one infers  $q$  from  $p$  one's acceptance of  $p$  is a cause of one's acceptance of  $q$  (see note 20 above). What I deny is that one's awareness of one's accepting  $q$  because  $p$  (one believes, or is supposing, or is assuming) when one infers  $q$  from  $p$  is matter of awareness of one's acceptance of  $q$ 's being caused by one's acceptance of  $p$ . It is instead a matter of one's acceptance of  $q$  manifesting one's at least apparent awareness of one's grounds for accepting  $q$ .<sup>63</sup>

---

<sup>62</sup> I am happy to admit that inference is blind in distinct sense, one which might even be closer to Wittgenstein's (see note 26 above). To infer one need not do so in the light of prior (real or apparent) awareness of the correctness of the inference in question (see sec. 5.8 below).

<sup>63</sup> I *will* reject the claim that when one infers  $q$  from  $p$  one's belief that  $q$  follows from  $p$  must be causally prior to one's acceptance of  $q$ . I will thus reject the Taking Condition and only endorse what McHugh and Way (2016, p. 316) label the "Consequence Condition". I claim only that to infer  $q$  from  $p$  one must believe that an appropriate inferential relation obtains between  $p$  and  $q$  (see sec. 4.11 below).

The above reveals what kind of conceptual sophistication is required in order to be capable of inference. To be capable of inference one must be capable believing that appropriate inferential relations obtain and thus possess the relevant concepts. To infer, that is, one must possess the concept *follows* or *entails* (or of whatever inferential relation one must take to obtain in order to infer). To insist this is not to “overintellectualize” (see sec. 1.7 above). Inference of the kind in question is that which occurs as a constituent of conscious and self-conscious reasoning of the sort which is our means to exercise doxastic agency. The capacity to engage in such reasoning is a sophisticated one. I have offered what amounts to an argument to the effect that to be capable of such reasoning one must be capable of appreciating that inferential relations obtain. That claim is not an implausible one in itself.

The picture of the self-consciousness of inference presented may also ease a further kind of overintellectualization worry. I have suggested that to be inferring one must know that one is doing so. Some may object that the conscious lives of inferrers need not be as conceptually rich as this seems to demand. To say that whenever one infers one must know *that one is inferring*, the thought is, is to demand that one must have a conceptually richer understanding of one’s own conscious acts than one in fact need have in order to be able to infer. When people inquire into what I mean by ‘inference’, for instance, it often takes a while for me to explain. Does this not suggest that they do not know *that they are inferring* when they do so and perhaps even lack the concept *inference* in use here?

One thing to note in response here is that my claim that to infer  $q$  from  $p$  one must know that one is inferring  $q$  from  $p$  is merely a theorist’s articulation of the claim that inference is self-conscious. Furthermore, as has been seen (sec. 4.4), inferring  $q$  from  $p$  does not require that one judge that one is inferring  $q$  from  $p$  when doing so. One’s inferences’ being self-conscious does not require that they be accompanied by further conscious occurrences and occupants of one’s stream of consciousness with content ‘I am inferring ...’. When one infers one’s knowledge of one’s doing so is rather a “silent partner” to one’s inference where this is a matter of a state of knowledge of one’s inference constitutively depending on the conscious occurrence which is one’s inference. In fact, the

account of inference's self-consciousness offered could even be maintained and it be accepted that in order to be able to infer self-consciously one *need not* possess the concept *inference*. To be able to infer self-consciously, on the account offered, there is no demand that one be such that one can judge that one is inferring. There is no demand that one must judge that one is inferring when one infers, nor is there even a demand that inferrers be able to make judgments where *inference* is a constituent of their contents. Inference is rather self-conscious because to infer just is to manifest awareness that one is inferring. If possession of the concept *inference* is understood as requiring the capacity to make judgments with *inference* as a constituent of their contents (c.f. G. Evans 1982, pp. 100–5) then on the view in question there is no demand that one must possess the concept *inference* in order to be capable of inferring. Nor does the account offered demand that one in fact does or even can conceptualise one's acts as acts of inferring. Again, it might then not be taken to demand that in order to be able to infer one must possess the concept *inference*. One rather just needs to be such that when one accepts something and one's act is an inference in so acting one manifests one's awareness that one is inferring, rather than merely judging, guessing, supposing, ... This does not require that one does or even can conceive of one's act as an inference—that in inferring one does or that one can manifest one's competence with the concept *inference* in use here. We can imagine a subject, for instance, who operated with the concept *inference\** where to infer\* is to infer where one's inference amounts to a judgment of that which one infers.<sup>64</sup> Such a subject might have a further concept *inference\*\** where to infer\*\* is to infer where one's inference does not amount to judgment of that which one infers. Such a subject does not operate with the concept of inference in question here. But this does not preclude them from inferring self-consciously as that is now being understood. Such a subject can infer and in doing so manifest their awareness that they are inferring, rather than performing a distinct kind of act of acceptance. That, I claim, is what is required in order for them to be capable of inferring self-consciously, rather than grasp of the concept of inference in use here. The account of the self-consciousness of

---

<sup>64</sup> Compare Koziol's (2017) idiosyncratic use of 'inference' discussed above (see sec 1.4)



inference offered thus allows us to maintain that all inference is self-conscious without attributing a too specific and conceptually rich understanding of one's acts to all who are capable of inferring. We arrive at a perhaps surprising result. Maintaining a conception of inference on which inferring requires having an appropriate taking belief, rather than being too conceptually demanding of those who are capable of inferring, in fact paves the way for an understanding of self-conscious inference which dispels with concerns to the effect that to see inference as necessarily self-conscious is to commit to an overintellectualized account of inference.

#### 4.9. The account in sum

In sum then, the account offered is as follows. For an inference from  $p$  to  $q$  to occur one must first consciously accept  $p$ , where one need neither believe  $p$  nor accept it in a distinctive way. One must then accept  $q$ , where again one need not believe  $q$  or come to believe  $q$ . Nor need one's acceptance of  $q$  be caused in a distinctive way. One's acceptance of  $q$  does not need to be the product of "the right kind of causation" or of "causation in virtue of rationalization". One must, however, know that one is inferring  $q$  from  $p$  *when one accepts  $q$* . This tells us that inference is a conscious and self-conscious act of acceptance—the event which is one's accepting that which one infers. To infer  $q$  from  $p$  is to accept  $q$  where in doing so one knows that one is inferring from  $p$ . To have such knowledge, one must know oneself to be accepting  $q$  because  $p$  (one believes, is supposing, or is assuming) where doing so requires that one believe that  $q$  follows from  $p$  such that one's act of acceptance is intelligible to oneself. This knowledge that one is inferring  $q$  from  $p$  is to be understood as a "silent partner" to one's act of acceptance, where this is in turn a matter of its constitutively depending on the conscious occurrence which is one's accepting  $q$ .

This account applies equally to inferences from beliefs, to inferences from suppositions and to inferences from assumptions made for practical purposes. Furthermore, the account does not face a problem of deviant causal chains. The account does not demand that in order for an inference from  $p$  to  $q$  to occur one's acceptance of  $p$  must cause one's acceptance of  $q$  in some right way. It instead requires that one's acceptance of  $q$  constitutively involve awareness of

one's inferring. Although non-reductive and circular, the account offered thus resolves the serious issues faced by causalist and by constitutive accounts of the sort considered above. Furthermore, the account paves the way for an understanding of how self-conscious reasoners are capable of exercising doxastic agency despite being unable to believe at will, and for an explanation of why to exercise such agency is to exploit one's self-consciousness.

#### **4.10. Developing the account further**

Inference, on the account offered, is a conscious event and self-conscious act of acceptance. One's knowledge that one is inferring when one infers is not a causal product of one's inference and rather constitutively depends on the occurrence which is one's inference. We can appeal to this "silent partner" to one's inferences to distinguish inference from non-inference.

As seen, saying this much allows us to answer the questions 'What is it for an inference to occur?' and 'What is an inference?', albeit circularly. There is, however, a further question which we might want to ask if our aim is a full understanding of the nature of inference. We might ask: What determines that one is inferring when one infers? The question is to be understood as asking what, metaphysically speaking, settles that one is inferring when one infers. One thing which has become clear is that appeal to the how inferences are caused is not the way to answer the question.

The account above yields an alternative way to answer the present question. On the account offered whenever one infers one knows that one is doing so because a state of awareness of one's inference constitutively depends on the act of acceptance which is one's inference. Furthermore, an act of acceptance's being an inference just is a matter of its being self-conscious in this way. In answer to the question 'What determines that one is inferring when one infers?' we can thus say that it is one's knowledge of one's inference itself which plays this role. When one infers it is one's knowledge that one is inferring, rather than the causal history of one's act, which settles that the conscious act of acceptance in question is an inference. One's inferences thus depend for their status as inferences upon one's knowledge of them. In this way, when one infers what is known does not occur independently of one's knowledge of what is known.

This amounts to a rejection of Koziol's (2017) claim that whenever one infers one's knowledge that one is inferring must be a "response" to one's inference. One's knowledge that one is inferring when one infers is not a mere response to one's inference and is rather what settles that the act of acceptance in question is an inference. This is a way to make more precise Kitcher's (2011, p. 130) suggestion that when one infers one's knowledge of what one is doing "is an indissoluble component of the self-conscious act itself." It is likewise a way to capture the above suggestion that when one infers one is not simply acting and, in addition, aware that one is doing so (see sec. 4.3). One's awareness of what one is doing is rather part of what makes the act in question the kind of act which it is. And it is a way to do this without saying that one's knowledge of one's inference (naturally understood as a kind of *state*) is simply *identical* to one's inference (a conscious *occurrence*) as Rödl (2013) suggests.

A move of the above form is made explicitly by Soteriou (2013, Chapter 10) in the case of judgment. According to Soteriou, as I have suggested is so in the case of inference, when one judges one does so self-consciously where this knowledge that one is judging constitutively depends on the conscious occurrence which is one's judgment. As Soteriou puts it, when one judges one knows that one is doing so where

[t]his distinctive form of self-knowledge involves the obtaining of an occurrent mental state—a mental state whose obtaining is constitutively dependent on the occurrence of a conscious mental event with temporal extension (2013, p. 251).

Soteriou notes that this means that when one judges one thereby has

a form of self-knowledge that is only available to the subject whose mental life includes a stream of consciousness. (2013, p. 251)

When one judges one knows that one is doing so where this knowledge constitutively depends on the conscious occurrence in question such that one has knowledge of a kind only available to one whose mental life involves conscious occurrences. I have argued that the very same is true of inference. But Soteriou also argues that when one judges it is one's knowledge that one is judging which settles that the conscious occurrence in question is a judgment. As he puts it

the occurrence of an event of the kind conscious judging that  $p$  depends upon the concurrent obtaining of some mental state of the subject—namely her belief that she is judging that  $p$ . (2013, p. 248)

When one judges that  $p$  one does so self-consciously where the obtaining of one's state of knowledge depends on the occurrence of the conscious occurrence which is one's judgment. However, one's knowledgeable belief that one is judging is what determines that the conscious occurrence in question is a judgment. In this way, one's knowledge that one is judging when one judges is not a "response" to one's judging as Koziolok understands this. One's knowledge of one's judgments is what determines that the conscious occurrences are judgments, despite that this knowledge itself constitutively depends on the occurrences in question.

Soteriou thus offers an account of judgment with the following structure. When one judges it is one's knowledge that one is judging—a silent partner to one's act—which determines that one is judging. Nonetheless, one only knows that one is judging *because* one is doing so. One's awareness of one's judgment constitutively depends on the conscious occurrence which is one's judgment. In this way one knows that one is judging whenever one judges simply *by judging* where it is this awareness of what one is doing which settles that the occurrence in question is a judgment. I can offer a similar answer to the question 'What determines that one is inferring when one infers?' When one infers it is one's knowledge that one is doing so which determines that one is inferring. Nonetheless, one only knows that one is inferring *because* one is doing so. One's awareness of one's inference constitutively depends on the conscious occurrence which is one's inference. In this way, one knows that one is inferring whenever one does so simply *by inferring* where it is this awareness of what one is doing which settles that the occurrence in question is an inference.

I have argued that when one infers one necessarily knows that one is doing so, where this knowledge of what one is doing constitutively depends on the occurrence which is one's inference.<sup>65</sup> I have also argued that whether one is

---

<sup>65</sup> My defence of this view will not be complete until my account of doxastic agency is on the table. The account is motivated by the need to capture inference's self-consciousness, where

inferring when one infers is not settled by the causal history of one's act. We should instead have it, I suggest, that when one infers it is one's knowledge that one is doing so which settles that the occurrence in question is an inference. To infer, on the view I suggest, just is to manifest one's awareness that one is inferring where it is this awareness of what one is up to, rather than the causal history of one's act of acceptance, which settles that the conscious act of acceptance in question is an inference. To infer is to manifest such awareness, I have suggested, because whenever one infers a state of awareness of one's inference constitutively depends on the act of acceptance which is one's inference. On the suggested account the obtaining of a state can constitutively depend on an occurrence whilst the nature of that occurrence is determined by the obtaining of that very state. A state of awareness of one's inference can constitutively depend upon the occurrence of one's inference whilst the nature of one's act of acceptance is determined by that very awareness. This account can be seen to be tenable, I have argued, once we see that we can reject Sellars' claim, as understood by Koziolk, that one's knowledge must always be a "conceptual response" to that which it is knowledge of. A promising account of inference is available on which one's knowledge of one's inferences constitutively depends upon and determines the nature of the conscious acts of acceptance which are one's inferences. Accepting this account not only resolves the difficulties which we began with, but paves the way for an explanation of how we are capable of exercising doxastic agency despite being unable to believe at will.

#### **4.11. Taking beliefs as neither epistemically nor causally prior**

I have argued that inference is self-conscious and that believing that one is inferring requires believing that *what one is inferring* follows from *what one is inferring from*. As seen above (sec. 2.3), it is common to agree. However, it seems that the place of such beliefs in inference has been widely misunderstood. One potential mistake is to muddle taking beliefs and knowledge of one's inferences.

---

the claim that inference is self-conscious is in turn motivated by the explanatory work the account of inference which it yields can do (see sec. 1.6 above and chapter 5).

Some speak as if having an appropriate taking belief and knowing that one is inferring is one and the same thing. For instance, it has been seen that Koziolak argues that the self-consciousness of one's inferences is a causal product of one's inferences. But Koziolak (2017, pp. 6–7) also takes this to be equivalent to the suggestion that *taking beliefs* are a causal product of one's inferences. The above account reveals why this is a natural error to make. To infer  $q$  from  $p$  one must do so self-consciously. Doing so requires that one's accepting  $q$  is intelligible to oneself where this in turn requires that one believe that  $q$  follows from  $p$ . It is because there is an intimate relation of this sort between knowing that one is inferring and having a corresponding taking belief, I suggest, which makes it natural to talk as if one's knowledge that one is inferring and one's having an appropriate corresponding taking belief are one and the same.

There are two further common claims concerning the *role* of taking beliefs in inference which I want to question. As seen, advocates of the Taking Condition claim that one's taking beliefs are *causally* prior to one's acceptance of that which one infers. We have also seen that when one infers one's accepting that which one infers *just is* one's inference. The Taking Condition thus commits us to the claim that when one infers one's taking belief is causally prior to one's inference. That is, when one infers  $q$  from  $p$  one does so *because* one takes it that an appropriate inferential relation obtains between  $p$  and  $q$ , where this 'because' is understood causally. One's inference is caused by one's taking belief. Relatedly, it is common to think that when one infers one's taking belief is *epistemically* prior to one's inference. That is, for one's inference from  $p$  to  $q$  to be warranted one's corresponding taking belief must be warranted, and if the inference is indeed warranted then this is so, at least in part, *because* one's taking belief is warranted. One's inferences, the claim is, depend both causally and epistemically on one's taking beliefs.

On inspection, both of the above claims appear questionable. Consider the suggestion that taking beliefs are causally prior to inference first. It is in fact not clear that one must have an appropriate taking belief *before* one makes a given inference. It is thus unclear in turn that one's inferences must be caused by appropriate prior taking beliefs. Suppose, for instance, that I infer  $q$  from  $p$  or  $q$  and from not- $p$ . Can I not come to believe that  $q$  follows from  $p$  or  $q$  and not- $p$

*by* making this inference? That is, can I not acquire beliefs about the inferential relations which obtain between what I infer from and what I infer *by* making inferences? This suggestion becomes plausible when we consider how we acquire our beliefs in the obtaining of inferential connections. We can, of course, do so via testimony. But being able to do so presupposes that we can acquire such beliefs other than via testimony. And a plausible answer to the question ‘How do we acquire our beliefs in particular inferential relations?’ is that, in many cases at least, we do so *by* making inferences. One might, for instance, suppose that  $p$  and then infer  $q$  and come to believe that  $q$  follows from  $p$  *by* making this inference.

As well as its being plausible to deny that one’s taking beliefs must be causally prior to one’s inferences the suggestion that they must be seems to be problematic. The suggestion appears to render much of our suppositional reasoning redundant. I might, for instance, suppose that  $p$  and then infer  $q$ . We are assuming that in order to do so I must believe that an appropriate inferential relation obtains between  $p$  and  $q$ . We can further assume that in question is a logical inference and thus that I must believe that  $q$  *follows* from  $p$ . Having supposed  $p$  and then inferred  $q$  I can then discharge the supposition that  $p$  and conclude that if  $p$  then  $q$ . The problem is then as follows. If I really did *already* need to believe that  $q$  follows from  $p$  in order to go through the above reasoning then there would be no need for me to do so. I could instead simply have inferred if  $p$  then  $q$  from that  $q$  follows from  $p$  outright. After all, if  $p$  then  $q$  does evidently follow from that  $q$  follows from  $p$ . It becomes clear why suppositional reasoning like the above need not be redundant if we have it that we can acquire taking beliefs *by* making inferences rather than having it that we already need appropriate taking beliefs in order to make inferences. In the above case, for instance, I suppose that  $p$  and then infer  $q$  and *by doing so* come to believe that  $q$  follows from  $p$ . Only then will I take myself to be in a position to conclude that if  $p$  then  $q$ .

The observation that one’s taking beliefs need not be *causally* prior to one’s inferences in turn casts doubt on the suggestion that they must be *epistemically* prior. If we can come to believe that inferential relations obtain by making inferences then plausibly we can come to *know* that they obtain by doing so as

well. It is, for instance, plausibly the case that one could come to know that  $q$  follows from  $p$  by going through the above suppositional reasoning. One might come to know that  $q$  follows from  $p$  by inferring  $q$  from  $p$  having supposed that  $p$ . When one makes an inference one's inferring can thus make one's taking belief warranted rather than one's taking belief being what makes one's inference warranted. That is, one's taking beliefs need not be epistemically prior to one's inferences.

Reflection on cases casts doubt on the suggestion that one's taking beliefs must be causally and epistemically prior to one's inferences. Further development of the account of inference offered—in particular when it comes to the relation between inference and belief—will further support the suggestion that inference can in fact be a *source* of taking beliefs and of knowledge of inferential relations (see in particular sec. 5.2 below).

#### **4.12. Conclusion**

Reflection on the shortcomings of Koziol's causalist account of inference has yielded an alternative proposal. Inference can be distinguished from non-inference via appeal to the knowledge which it constitutively involves, rather than via appeal to the causal history of the occurrence or to what it causally involves. On the account of inference which results inference is a kind of conscious event and act of acceptance the nature of which depends upon the self-knowledge which it constitutively involves. This account resolves the ontological difficulty with which we began and does not face the difficulties faced by causalist accounts, nor those faced by the constitutivist accounts considered. On it, whenever one infers one must have an appropriate taking belief such that one's inference is not just self-conscious but intelligible to oneself. The account suggests that the role of taking beliefs in inference has been widely misunderstood. In the next chapter I go on to further consider the relation between inference and belief and the role of inference in doxastic agency.



## 5. Self-conscious reasoning as doxastic agency and “making up our minds”

**Chapter abstract** In this chapter I apply the account of inference developed in order to yield an account of doxastic agency and an explanation of how we can exercise such agency despite the fact that we cannot believe or come to believe “at will”. We can do so, I suggest, by *reasoning* because reasoning is a kind of action, because reasoning involves inferences as constituents, and because to infer can *just be* to come to believe. Given this, to reason can *just be* to exercise doxastic agency. I explain how the temptation to respond to the challenge to the possibility of doxastic agency by moving to a view on which we are at best capable of actions of a sort which lead causally to belief revision is removed once the causalist conception of inference is abandoned. I likewise explain why, with the suggested alternative account of inference on the table and with the relevant reasoning conceived of appropriately, there should be no temptation to think that to exercise doxastic agency by reasoning one would need to be able to infer at will where one’s inference yields belief. The reasoning by which we exercise doxastic agency is best conceived of as primarily *epistemic* agency—as agency the aim of which is *knowledge* (rather than judgment or belief). The fact that a process with this aim fails to terminate with inference at will should not prevent us from seeing it as amounting to the active achievement of one’s aim. I conclude by showing how the account offered sheds light on the role of self-consciousness in doxastic agency. Self-consciousness is required for and exploited in the exercise of doxastic agency, not because we must know what we believe in order to be active with respect to our beliefs, but because the exercise of doxastic agency is *itself a self-conscious activity*.

### 5.1. The challenge to the possibility of doxastic agency again

I have offered an account of inference on which inference is understood as a distinctive kind of conscious and self-conscious occurrence. Inference is a kind of self-conscious occurrence, on the account offered, in that the inferences are not just the subset of the inferences in a more permissive sense which happen

to be both conscious and self-conscious. Inference is a kind of essentially conscious and self-conscious occurrence. Furthermore, when an inference occurs the act of acceptance's being an inference is not determined by its causal history. It is rather determined by what it constitutively involves, namely, knowledge of its occurrence. This account has been motivated by its aptness to yield an explanation of how we are capable of doxastic agency and of why to exercise such agency is to exploit one's self-consciousness. We thus need to move on to the task of showing how this is so.

The key thing to be explained is how it is that we are capable of doxastic agency. We take ourselves to be capable of being active with respect to our beliefs, and not merely by engaging in actions of a sort which can lead causally to belief revision. We are capable of "directly" exercising agency over our beliefs in a way which does not amount to mere self-manipulation—of "settling questions" for ourselves (see sec. 1.5 above). As seen, there is a distinctive challenge to the suggestion that we are capable of such agency which stems from the observation that belief in some sense aims at truth (again, see sec. 1.5 above). Belief's aiming at truth, it seems clear, precludes us from being able to believe or come to believe at will. To believe or come to believe at will, after all, would be to do so where one's doing so is up to oneself. It would thus be to do so whilst seeing the constraint of treating  $p$  as true as self-imposed. Yet to believe  $p$  one must treat  $p$  as true without seeing the constraint of treating  $p$  as true as self-imposed.

To believe or come to believe at will, the thought is, would be to do so whilst seeing the constraint of treating  $p$  as true as self-imposed. But an attitude's being held or formed in that way would undermine any claim to the effect that the attitude is a belief. It seems clear that we cannot believe or come to believe at will accordingly.<sup>66</sup> And given this we must ask how it is that we are capable of doxastic agency nonetheless. How are we capable of direct agency over belief of the form in question despite that we cannot believe or come to believe at will? As noted (sec. 1.5), this challenge is taken to force upon us a view on which

---

<sup>66</sup> As noted (sec. 1.5), if there are exception they are not apt to be appealed to in order to explain how we are capable of doxastic agency to the extent sought.

agency is connected to belief only in that beliefs can be the effects of something agential. We are best capable of actions of a sort which can lead causally to beliefs via “mental ballistics” (Strawson 2003, p. 241).

## 5.2. Reasoning as action and the token identity view

It might seem like appeal to inference will be of little help when it comes to explaining how we are capable of doxastic agency in the face of the above challenge. The cases of interest, after all, are primarily those in which inferences yield belief in that which is inferred. In such cases one judges *by* inferring. As seen, one’s inference just is the event of accepting that which one infers. And that is the same event which is one’s judging. In such cases one does not both judge and infer where this is a matter of the occurrence of two distinct events. One judges *by* inferring where the event which is one’s inference *just is* the event which is one’s judgment. Furthermore, it seems clear that whenever judgment yields belief one’s judgment cannot be performed at will. Judgment, just like belief, aims at truth. And this seems to preclude us from being able to judge at will when judgment yields belief, just as belief’s aiming at truth seems to prevent us from being able to believe or come to believe at will. To judge that *p* at will where one’s judgment yields belief, after all, would be to accept *p* where one’s doing so is up to oneself. It would thus be to accept that *p* where in doing so one sees the constraint of treating *p* as true as self-imposed. And doing that would seem to undermine the status of one’s act as one of judgment (c.f. sec. 1.5 above).

Such considerations do not cast doubt on the thought that one might be able to judge that *p* at will if one *already believes that p*. If one already believed that *p* then one could judge that *p* where one’s doing so is up to oneself without seeing the constraint of treating *p* as true as self-imposed. But it remains clear that we cannot judge at will when judgment yields belief any more than we can believe or come to believe at will. And when inference yields belief in what one infers one’s inference amounts to a judgment. If we attempt to appeal to inference in order to explain how doxastic agency is possible we thus seem to face a challenge much like that with which we began. It must be explained how we can exercise doxastic agency by inferring despite that whenever inference

yields belief in what one infers one does not infer at will. Appeal to inference, one might think, is not apt to help explain how doxastic agency is possible since we can no more infer at will when inference yields belief than we can believe or come to believe at will.

We can accept that inferences, when they yield beliefs, are not performed at will. Doing so does not prevent us from saying that inference plays a central role in the exercise of doxastic agency. Inference, as seen, is an event of a kind which occurs, whenever it does, as a constituent of a process of reasoning (see sec. 4.6 above). To infer  $q$  from  $p$ , for instance, one must accept  $p$  and then accept  $q$ , inferring  $q$  from  $p$  in doing the latter. To do so is to engage in a minimal bit of reasoning. Inference, when it occurs, at a minimum occurs as a constituent of such a minimal bit of reasoning and can also occur as a constituent of a more extended process of reasoning. Furthermore, it seems that we can unproblematically maintain that *reasoning* is action. Reasoning can be maintained to be a kind of action without its being insisted that all constituents of one's reasoning when one reasons must be performed at will. Reasoning, the thought is, is a conscious process where seeing constituent acts of the process as not being performed at will need not prevent us from seeing the process as a whole as amounting to action (see sec. 5.7 below for further discussion).<sup>67</sup> Reasoning is even the kind of thing which we can perform at will. One can reason about, say, whether  $p$  where one's doing so is up to oneself. No pressure to think otherwise should stem from the observation that some constituent acts which make up such reasoning cannot themselves be performed at will as a consequence of their aiming at truth (although see sec. 5.5 below).

On the suggested view reasoning is a kind of action where this *does not* require that all of the constituent events of the process be performed at will.<sup>68</sup>

---

<sup>67</sup> Reasoning, being a kind of process, might be more appropriately labelled a kind of "activity" (Hornsby 2012, 2013). What is crucial is that reasoning can be an exercise of agency of a sort apt to potentially amount to the exercise of doxastic agency.

<sup>68</sup> I am not the first to make a suggestion along these lines. Soteriou (2013, pt. 2) sees reasoning as a kind of exercise of agency without insisting that constituent judgments must themselves be actions, his focus being the special case in which we *suspend judgment* by reasoning. Hornsby (2012, 2013), meanwhile, suggests that we perform actions by engaging in activities where these activities are processes. She also denies that such actions "can always be resolved into indefinitely many different sub-actions" (2013, p. 12). My suggestion is that a point of this form is exemplified in the case of reasoning. We can engage in the activity of reasoning where doing

This structure is not unique to the case of reasoning. It seems it can even be found in cases of ordinary bodily action. Take, for instance, my typing right now. My doing so, it seems clear, is an action. Nonetheless, some finger movements which are constituents of the action are plausibly not performed at will. I may, for instance, just have pressed the ‘P’ key with my ring finger. But my doing so may not have been something I even knew I was doing distinct from my awareness of my typing each word or phrase as a whole. It may well have not been something I did at will accordingly. My typing was an action where this does not require that all constituent events of the action themselves have been performed at will.

This observation alone does not yield an explanation of how we are capable of doxastic agency. We can reason where such reasoning amounts to action, doing so can constitutively involve inference, and inference can yield belief. However, as noted, the relevant inferences are judgments of that which is inferred. Furthermore, it appears to be standard to assume that when judgments yield beliefs they do so by *causing* beliefs. More precisely, the idea seems to be this: when judgment yields belief, an event of judging that *p* causes formation of a belief that *p*, that is, an event of coming to believe that *p*. This suggestion is endorsed explicitly by Golob (2015, sec. 3), Lee (forthcoming), Owens (2000, p. 87) and O’Shaughnessy (2000, p. 106). Toribio also appears to be endorsing the view when she states that “judging is an event that very often leads to the formation of a belief?” (2011, p. 346) as does Cassam when he states that “judging that *p* normally leads one to believe that *p*” (2010, p. 81; see also Shah and Velleman 2005, p. 503). As Boyle (2011a, p. 5) notes, authors tend not to unambiguously endorse the suggestion in question, but nor do they unambiguously reject it and often seem implicitly committed.<sup>69</sup>

What results is a view on which we can reason where such reasoning is action and where doing so can *cause* beliefs. We have not yet identified an alternative to the suggestion that we are at best capable of agency with respect

---

so does not require that all constituent events of one’s reasoning themselves be performed at will.

<sup>69</sup> There are some who reject the view in question. McHugh, for instance, claims that when judgment yields belief one’s judgment “constitutes the acquisition” of the belief (2009, p. 246).

to our beliefs in that we are capable of actions of a sort which can cause beliefs. However, a different conception of the relation between inference and belief (and between judgment and belief) is available.<sup>70</sup> When inference yields belief in that which one infers one's inference is *token identical* to the event which is one's coming to believe that which one infers. And reasoning can conclude with such inference.<sup>71</sup> Given this, to reason can, in part, *just be* to come to believe that which one concludes and, in this way, can *just be* to exercise doxastic agency.

This alternative conception of the relation between inference and belief can be motivated by consideration of the apparent inadequacies of the standard view on which judgments and inferences yield beliefs by causing them. In drawing out these inadequacies we can begin with the observation that when one reasons one will, at least typically, do so with an aim. A typical aim will be to *figure out* or *find out whether* something is the case (see sec. 5.5 below). Crucially, one has not figured out or found out whether *p* unless one *knows* whether *p*. One has thus not achieved the aim of reasoning the aim of which is to find out or figure out whether *p* unless one knows whether *p* and thus either knows that *p* or knows that not-*p* (Soteriou 2013, pp. 350–1). Furthermore, when one concludes deliberation with inference one conceives of the act as amounting to the achievement of the aim of one's deliberation. Boyle expresses such a thought when he says that when I conclude reasoning with judgment “[t]he act of judgment is the completion of my project, not a step towards it” (2009, p. 129). Consider then a paradigm case of successful reasoning. I do not yet believe either *q* or not-*q*, deliberate about whether *q* and conclude by inferring *q* from *p*. I conceive of my inferring *q* as amounting to my achieving the aim of my deliberation. And that aim is to come to know either *q* or not-*q*. I will thus only have achieved my aim as I conceive of it if I know and thus believe that *q when* I infer *q*. On the current picture, however, when I conclude reasoning by inferring *q* I do not yet believe *q*. I only do something which will bring it about that I believe *q*, all being well.

---

<sup>70</sup> The suggestion is made in my (2018).

<sup>71</sup> As seen, when one does so one will judge by inferring and one's inference and one's judgment will be one and the same event. I focus on the relation between inference and belief in what follows.

When we conclude reasoning by inferring we may fail to achieve the aim of such reasoning in doing so. One might fail to come to know and perhaps even fail to come to believe what one judges (see sec. 1.4 above). The problem for the view in question is not that it allows that in concluding reasoning by inferring one *might* fail to achieve the aim of one's deliberation in doing so. The worry is rather that if the view is correct then inferring *cannot* amount to achieving the typical aim of reasoning as we conceive of it. If the view in question is correct then whenever I conclude some typical reasoning by inferring I have not yet achieved the aim of my reasoning. I have rather merely taken a "step towards" doing so by doing something which will, all being well, bring it about that my aim is achieved.

The above issue, though not decisive, suffices to motivate a search for alternative conception of the relation between inference and belief when the former yields the latter. Inference, as seen, is a kind of conscious event. Furthermore, as it is common to observe, events can be variously described (e.g. Anscombe 1957; Davidson 2001a, 2001c; Hornsby 1980, 2012, pp. 234–5; Ryle 2009b, 2009c). We saw this when considering the relation between inference and judgment when one comes to believe by inferring above. When one infers  $q$  from  $p$  and one's inference yields belief that  $q$  one judges that  $q$  *by* inferring  $q$ . Surely we should not say that in such a case one both infers and judges where these are distinct events. One infers and judges where both events are *identical* to the event which is one's accepting  $q$ .

In a similar vein Anscombe asks:

Are we to say that the man who (intentionally) moves his arm, operates the pump, replenishes the water-supply, poisons the inhabitants, is performing *four* actions? Or only one? (1957, p. 45)

She suggests that we should not see there as being four distinct events in question here. Rather

moving his arm up and down with his finger round the pump handle *is*, in these circumstances, operating the pump; and, in these circumstances, it *is* replenishing the house water-supply; and, in these circumstances, it *is* poisoning the household (1957, p. 46).

Ryle's (2009b, p. 494), meanwhile, gives the example of someone's winking. Ryle suggests that when such an event occurs it can be given "thicker" or "thinner" descriptions,  $A$  being a thicker description of an event than  $B$  iff one  $A$ s by  $B$ -ing. To give a thin description of the event which is one's winking, Ryle suggests, would be to call it a contracting of one's eyelid. To give a thicker description would be to call it a winking. To give a thicker description still might be to call it a signalling conspiratorially to an accomplice. This does not mean, Ryle suggests, that there are several distinct events in question here: a contracting of the eyelid, a winking and a signalling. There is just one event in question which can be given thicker and thinner descriptions—a single event which can be variously described.

If we accept this picture of event individuation then our alternative conception of the relation between inference and belief when the former yield the latter becomes available. As seen, when inference from  $p$  to  $q$  yields belief that  $q$  one's inference is the conscious event which is one's acceptance of  $q$ . This event can be variously described. It can be described, for instance, as a judging that  $q$ . It might likewise be described as a concluding of reasoning about whether  $q$ . Crucially, we can also say, such an event can be described as a coming to believe that  $q$ . When one infers  $q$  from  $p$ , in the case in question, in doing so one judges that  $q$ , comes to believe  $q$  and perhaps also concludes reasoning about whether  $q$ . As before, this is not because there are several distinct events in question here, an inferring, a judging, a coming to believe, ... There is just one event in question which can be variously described. When inference from  $p$  to  $q$  yields belief in  $q$ , on this view, the event which is one's inference *just is* the event which is one's coming to believe  $q$ . One's inference is *token identical* to the event which is one's coming to believe that which one infers. In such a case one comes to believe that  $q$  by inferring  $q$  from  $p$ . This is so not because one's inference causes one's belief that  $q$  but because the event which is one's inferring  $q$  *just is* the event which is one's coming to believe  $q$  and thus because, in Ryle's terms, to call the event in question a coming to believe that  $q$  is to give a thicker description of the event in question than one gives in calling it an inferring of  $q$ .



When an inference yields a corresponding belief the inference is thus related to the resultant belief as getting married is related to being married and parking on double yellows may be to being a lawbreaker. Getting married does not *cause* a couple to be married, nor does parking on double yellows *cause* one to be a lawbreaker. The event which is the couple's getting married *just is* the event which is their coming to be married. And the event which is one's parking on double yellows may *just be* the event which is one's becoming a lawbreaker. Similarly, when I come to believe by inferring the event which is my inferring *just is* the event which is my coming to believe. The inference is related to the resultant belief as the event of *o*'s coming into being relates to *o* for any *o*. Inferring, on this view, can unproblematically be seen as the achievement of the typical aim of reasoning. To infer *q* from *p*, for instance, can be to achieve the aim of reasoning the aim of which is to figure out whether *q* because to make such an inference can *just be* to come to know that *q*.

So this is the token identity view I suggest: if I come to believe *q* by inferring *q*, my coming to believe *q* is simply the same event as my inferring *q*. When inference yields belief in that which one infers the event which is one's inference is *token identical* to the event which is one's coming to believe.<sup>72</sup> With this view on the table it can be seen how the observation that reasoning can be action yields a response to the challenge to the possibility of doxastic agency. Of concern is how doxastic agency is possible despite that we can neither believe nor come to believe at will. When inference yields belief, I have conceded, one does not infer at will either. But inferences are events of a kind which occur as constituents of processes of reasoning. And reasoning, we can maintain, is a kind of action. Furthermore, I have now suggested, to infer can *just be* to come to believe that which one infers. Given this, reasoning can be action where to reason can, in part, *just be* to come to believe that which one concludes. To reason can, in this way, *just be* to exercise doxastic agency, rather than merely being action of a kind which can lead causally to belief.

In offering this suggestion I have focused on the case in which doxastic agency is exercised via the concluding of reasoning with inference. In fact, there

---

<sup>72</sup> An analogous claim is true of judgment.

is no need to see doxastic agency through reasoning as only being exercised in cases of this kind. We often engage in extended processes of reasoning and in doing so acquire and revise various beliefs. The token identity view, along with the suggestion that reasoning can be action, allows us to say that much of such reasoning can amount to the exercise of doxastic agency.<sup>73</sup>

The token identity view can also be extended in order to bring clarity to the matter of the place of taking beliefs in inference. I have argued that to infer one must have an appropriate taking belief (sec. 4.8). To infer  $q$  from  $p$ , for instance, one must believe that an appropriate inferential relation obtains between  $q$  and  $p$ —that  $q$  follows from  $p$ —and must have such a belief *when* one infers (see sec. 4.8 above). But such beliefs need not be causally or epistemically prior to one’s inferences. In fact, it seems that we can acquire such taking beliefs and can acquire knowledge of such inferential relations *by* making inferences. One might, for instance, suppose that  $p$  and infer  $q$  from it and in doing so come to believe and to know that  $q$  follows from  $p$ . Making inferences appears to be our most basic way of acquiring knowledge of inferential relations (see sec. 4.11 above).

The place of taking beliefs in inference has thus been widely misunderstood. To infer  $q$  from  $p$  one must believe that  $q$  follows from  $p$ , but one’s belief that  $q$  follows from  $p$  need not play a causal role in one’s inferring  $q$  from  $p$ . Drawing on the token identity view, we can offer an alternative picture of the place of taking belief in inference. To do so is to build on suggestions from Valaris and Rödl. We saw, for instance, that Valaris (2014, 2016c, 2016b) suggests that beliefs are not “isolated atoms” and can rather be constitutively related (sec. 3.2). Believing  $q$  on the basis of  $p$ , he suggests, can be constitutive of believing that  $q$  follows from  $p$ . Likewise believing that  $p$  and that  $q$  follows from  $p$  can be constitutive of believing  $q$ . We saw that this suggestion does not yield an account of inference (sec. 3.3). But combining the token identity view with the suggestion that we can acquire taking beliefs by inferring yields a related suggestion. Rather than (or in addition to) seeing there as being constitutive

---

<sup>73</sup> For discussion of which such reasoning amounts to the exercise of doxastic agency see sec. 5.8 below.

relations between beliefs, we can have it that there are even more intimate relations between events of *belief formation*. Namely, they can be token identical. An event of inferring  $q$  from  $p$  can potentially yield belief that  $q$  follows from  $p$ —and must do so if one does not believe that  $q$  follows from  $p$ —because an event of inferring  $q$  from  $p$  can be token identical to an event of coming to believe that  $q$  follows from  $p$ . It is in this way that we can acquire taking beliefs *by* making inferences. Just as inference from  $p$  to  $q$  can yield belief in  $q$  it can yield belief that  $q$  follows from  $p$ . And supposing one infers  $q$  from  $p$  and does not yet believe that  $q$  follows from  $p$  one must come to believe that  $q$  follows from  $p$  *by* doing so, since to infer  $q$  from  $p$  one must have such a taking belief.

The above is, I take it, a refinement of one of Rödl's (2013) suggestions. Rödl's view is that to infer can *just be* to recognise one's grounds as sufficient grounds (see sec. 3.4 above). I noted that one reading of this suggestion is that talk of recognising as sufficient grounds here is intended as reference to a kind of kind of *occurrence* rather than to (what is naturally understood as) a kind of *state*. Thus read Rödl's suggestion is that when one (successfully) infers  $q$  from  $p$  the event which is one's inferring  $q$  is *token identical* to the event which is one's coming to know that  $p$  is sufficient grounds to do so. I also noted that as it stands Rödl's account is too restrictive, leaving out the possibility of inference without judgment. On the account I have offered to infer  $q$  from  $p$  can just be to come to know that  $q$  follows from  $p$ . This suggestion applies in cases in which one does not come to believe  $q$ . It also seems apt to explain how, in cases in which one is judging in inferring, one's doing so can put one in a position to recognise as sufficient one's grounds for doing so. One who comes to know  $q$  by inferring it from  $p$  can in doing so also come to know that  $q$  follows from  $p$ . And one who self-consciously infers  $q$  from  $p$ , who comes to know  $q$  in doing so and who knows that  $q$  follows from  $p$ , would seem to be left in a position to recognise that one's grounds for having it that  $q$  are sufficient (all being well).

### 5.3. Concerns for the token identity view

I have suggested the following view of the relation between inference and belief when the former yields the latter. When one comes to believe  $q$  by inferring  $q$  from  $p$  the event which is one's inferring  $q$  *just is* the event which is one's coming

to believe  $q$ . When inference plays this role it amounts to judgment. A commitment of my view is thus that judgments can be token identical to events of coming to believe what one judges. As suggested, the view can likewise be extended to bring clarity to the question of the place of taking beliefs in inference. To infer  $q$  from  $p$  one must believe that  $q$  follows from  $p$ . If one does not already believe that  $q$  follows from  $p$  then one will come to believe that  $q$  follows from  $p$  *by* inferring where this is a matter of one's inference being *token identical* to the event which is one's coming to believe that  $q$  follows from  $p$ .

Someone might object to the token identity view along the following lines (see Lee forthcoming, sec. 4; O'Shaughnessy 2000, p. 106). It is possible to judge that  $p$  without coming to believe  $p$ . It is also possible to come to believe  $p$  without judging that  $p$ . These events have distinct modal profiles and therefore cannot be identical. But making this objection would amount to confusing event types and event tokens. The event types *judging that  $p$*  and *coming to believe  $p$*  are indeed not identical, but that does not mean that a token of the former type cannot also be a token of the latter type. Likewise, the event types *inferring  $q$  from  $p$* , *coming to believe  $q$*  and *coming to believe that  $q$  follows from  $p$*  are not identical. But that does not mean that a token of the former type cannot also be a token of one or both of the latter types.

Another misguided objection to the token identity view advanced is offered by Boyle (2009, sec. 4, 2011a, sec. 3.5). Boyle argues that judgments cannot be identical to events of coming to believe. He could, in the same way argue the for the analogous claim of inferences. Boyle's argument proceeds as follows. Suppose I come to believe  $p$  by judging that  $p$ . According to the token identity view offered the event which is my judging that  $p$  is identical to the event which is my coming to believe  $p$ . Boyle claims that a dilemma arises for this suggestion when we ask: 'At the time at which I judge that  $p$  do I believe that  $p$ ?' He argues as follows:

If the subject does believe  $p$  at the moment when he judges  $p$  (or throughout the duration of this event, if it takes time), then it seems that his judging is not an event of "making up his mind" or "forming" the belief that  $p$ , for even at the first moment of its taking place, he *already*

believes the proposition in question. (2011a, p. 14 my emphasis, see also 2009, p. 130)

If we suppose that when I judge that  $p$  I believe  $p$ , Boyle claims, then I *already* believe  $p$  when I make this judgment. I therefore cannot be forming the belief by judging. Suppose instead then that I do not believe  $p$  when I judge that  $p$ . In that case the event which is my judging is not itself the event of my forming the belief that  $p$ , for no such belief exists yet (Boyle 2011a, p. 14). At best, on this view, my judgment that  $p$  can cause a later belief that  $p$ .

Boyle's dilemma is an instance of a general dilemma<sup>74</sup> Suppose we think that event  $e$  is the event of  $o$ 's coming into being. We can then ask: 'Does  $o$  exist at the time of  $e$ 's occurrence?' If the answer is 'Yes' then  $e$  is not really an event of  $o$ 's coming into being since  $o$  *already* exists. If the answer is 'No' then  $e$  is again not an event of  $o$ 's coming into being since  $o$  does not exist yet.  $e$  can at best cause  $o$ 's coming into being. For any event  $e$ ,  $e$  cannot be the event of  $o$ 's coming into being. To see what has gone wrong it will be instructive to consider another instance of the general dilemma.

Suppose I am granted such authority that I can fire you simply by saying 'You're fired'. When I do this, I do not cause you to be fired. My saying 'You're fired' rather makes it the case that you are fired. The event which is my saying 'You're fired' just is the event which is your becoming jobless. A version of Boyle's dilemma arises for this suggestion. Suppose that we ask of the event of my saying 'You're fired': 'Are you jobless when I say this or not?' If the answer is 'Yes' then you are *already* jobless and I have not really fired you. If the answer is 'No' then, again, I have not really fired you. The event which is my saying 'You're fired' can at best cause some future event which is your becoming fired.

On inspection, it can be seen that the first horn of this alleged dilemma trades on a mistake. It can be maintained that my saying 'You're fired' does make it the case that you are fired and that you are thereby jobless when I say this. But that does not mean that you are *already* fired when I say 'You're fired'. You are fired when I say this, not before. For you to be *already* fired when I say

---

<sup>74</sup> Boyle only acknowledges that his dilemma "is related to a more general problem about events of starting and stopping" (2011a, p. 14). It is not clear whether what follows is precisely what he has in mind

‘You’re fired’ would require you to be fired at some time prior to when I say it, which is not the case.

An analogous response is available to Boyle’s dilemma. Suppose I judge that  $p$  and come to believe  $p$  by doing so. The event which is my judging that  $p$  is token identical to the event which is my coming to believe  $p$ . We can then ask: ‘When I judge that  $p$  do I believe that  $p$ ?’ The answer is ‘Yes’. This does not entail that I *already* believe that  $p$  when I judge that  $p$ , just as the fact that you are fired when I say ‘You’re fired’ does not entail that you are *already* fired when I say this. For me to *already* believe  $p$  when I judge that  $p$  would be for me to have a temporally prior belief that  $p$ , which I do not in the case in question. Boyle’s dilemma does not arise for the token identity view. The view is thus apt to be appealed to, as above, in order to explain how we can exercise doxastic agency by reasoning despite that inferences are not actions. We can exercise doxastic agency by reasoning because reasoning can be action, because reasoning constitutively involves inference, and because to infer can *just be* to come to believe that which one infers (and, we can add, that that which one infers follows from that which one infers from).

#### 5.4. Reasoning and mental ballistics

It would be natural to object to the above suggestion on the grounds that in offering it I have *assumed* that reasoning is action and is action when it yields belief. I have taken it that such reasoning can unproblematically be maintained to be action. I have assumed that this is so given that there is no challenge to the suggestion that such reasoning can be performed at will analogous to the challenge to the suggestions that we can believe at will, can come to believe at will, and can infer at will when inference yields belief. To take this line is to ignore a challenge to the suggestion that reasoning of the kind in question is, as a whole, action.<sup>75</sup>

I have suggested that we can exercise doxastic agency by reasoning because we can reason where to do so is action, because such reasoning can involve

---

<sup>75</sup> A challenge which I ignored in first stating the above proposal. I said that “it seems unproblematic to say that reasoning is a form of voluntary, intentional action” (2018, p. 16).

inferences as constituents, and because to infer can *just be* to come to believe that which one infers. Given this, to reason can just be to exercise doxastic agency. A difficulty for this suggestion may seem to arise as a result of the fact that the relevant reasoning has inferences as constituents. When one concludes such reasoning by inferring, we have conceded, one does not infer at will. This is commonly taken to force us to accept a view on which the relevant reasoning cannot be action as a whole (Kornblith 2012, pp. 85–90; Setiya 2013, pp. 183–4; Strawson 2003; Valaris 2016b, sec. 5). Only some of the reasoning can amount to action, and the part of the reasoning which counts—that constituted by the concluding inference which yields belief—cannot be. The difficulty here stems from the fact that the relevant reasoning is *wholly constituted* by its constituent acts. Suppose, for instance, that one judges that  $p$ , then infers  $q$  from  $p$ , and finally concludes one’s reasoning by inferring  $r$  from  $p$  and  $q$ . The occurrence of one’s reasoning is *wholly constituted* by the occurrence of the relevant acts of acceptance. To judge that  $p$ , then infer  $q$  from  $p$ , and then infer  $r$  from  $p$  and  $q$ , *just is* to perform the above reasoning from  $p$  to the conclusion that  $r$ . Furthermore, the constituent acts which make up one’s reasoning may not be performed at will. And one’s concluding acceptance of  $r$ , assuming it yields belief in  $r$ , cannot be performed at will. The challenge is then as follows. How could a process of reasoning which concludes with inference to  $r$  be action as a whole when it is wholly constituted by acts some and perhaps all of which are not performed at will?

The above concluding inference’s not being performed at will might be taken to cast doubt on the suggestion that the inference itself could be an action (e.g. O’Shaughnessy 2008b, pp. 543–7). How could an inference which yields belief be an action if such an inference is *not even the kind of thing which one could do at will*? When inference yields belief in that which one infers, the thought is, one’s so inferring cannot be up to oneself. That might be taken to cast serious doubt on the suggestion that such an inference could be an action. Again, a challenge to the suggestion that reasoning such as the above could amount to doxastic agency then results. How could reasoning which concludes with inference to  $r$  be action as a whole when the reasoning is wholly constituted by acts some and perhaps all of which, there is pressure to think, are not themselves

actions? Suppose as above that one judges that  $p$ , infers  $q$  from  $p$  and then infers  $r$  from  $p$  and  $q$ . How could one qualify as reasoning when one infers  $r$  where one's reasoning as a whole is action despite that the present inference, there is pressure to think, cannot be an action?<sup>76</sup>

An advocate of the causalist conception of inference will not be well placed to respond to the above challenge. Inference to  $r$  which yields belief's not being such that it could be performed at will casts doubt on the suggestion that such an act of acceptance could be an action. It likewise casts doubt on the suggestion that reasoning which concludes with such an act could amount to action as a whole and to the exercise of doxastic agency accordingly. According to advocates of the causalist conception of inference, for such an inference to occur is for acceptance of  $r$  to be appropriately caused by prior acceptance. But acceptance of  $r$ 's being caused by a prior act of acceptance, even supposing that the prior act is performed at will, will not make one's acceptance of  $r$  such that it is itself performed at will. Nor will appeal to acceptance of  $r$ 's being caused by a prior act of acceptance suffice to explain why we should see one's act of accepting  $r$  as an action, even supposing that the prior act is itself an action.

Advocates of the causalist conception of inference have it that when one infers and in doing so comes to believe what one infers one's act of accepting that which one infers is the kind of act it is in virtue of its being appropriately caused by a prior act of acceptance. When it comes to explaining how it might be that in concluding reasoning with inference one could be exercising doxastic agency the causalist will thus look to how one's concluding act is caused. But in doing so they will find nothing apt to explain how in inferring one could be exercising doxastic agency. Appeal to the fact that an act of acceptance which yields belief is the causal product of prior acts of acceptance will not suffice to

---

<sup>76</sup> It would not do, in the face of this challenge, to simply insist that the relevant reasoning as a whole is action despite seemingly not being made up of constituent actions, as we can insist that typing is action despite its not being the case that every constituent event need itself be an action (see sec. 5.2 above). When one engages in an activity of typing in doing so one performs a series of actions (e.g. one types particular words and phrases). That remains so even if every constituent event of one's typing is not itself an action. To maintain that such typing is action we just need to deny, as Hornsby does, that such actions "can always be resolved into indefinitely many different sub-actions" (2013, p. 12). The reasoning in question, meanwhile, is such that there is pressure to see it as at best only in part amounting to the performing of a series of actions.



explain how, in performing the act, one could be exercising doxastic agency despite that one is not performing the act at will.

Accepting the above line of reasoning leads us to a version of the picture on which agency is at best connected to belief in that beliefs can be the effects of something agential. When one reasons and concludes by inferring where one's inference yields belief the reasoning as a whole cannot be action. At best "preparatory" acts and constituents of one's reasoning can be actions, where such actions may go on to cause beliefs (Strawson 2003, p. 232; see also Kornblith 2012, pp. 89–90; Setiya 2013, pp. 183–4; Valaris 2016b, sec. 5). Strawson elaborates on and endorses the view in question here. On it, he explains,

the role of genuine action in thought is at best indirect. It is entirely prefatory, it is essentially—merely— catalytic. (2003, p. 231)

Actions can be involved in reasoning, he suggests. But when one reasons the only actions are really things like "setting one's mind at [a] problem" and "shepherd[ing] one's wandering mind back to the previous thought-content" (2003, pp. 231–2). Strawson insists that

action, in thinking, really goes no further than this. The rest is waiting, seeing if anything happens, waiting for content to come to mind, for the 'natural causality of reason' to operate in one. (2003, p. 232)

On his view there is

no action at all in reasoning and judging considered independently of the preparatory, catalytic phenomena just mentioned (2003, p. 232).

Inferences themselves are not actions, on this view, and reasoning as a whole cannot be action either. At best we are capable of actions of a kind which lead causally to inference and belief fixation through "mental ballistics" (2003, p. 241). We are thus not capable of exercising doxastic agency by reasoning as suggested above.

### **5.5. "Making up our minds" and doxastic versus epistemic agency**

The observation that when inference yields belief one cannot be inferring at will, just as one cannot believe or come to believe at will, yields a challenge to the suggestion that to reason can be to exercise doxastic agency. Reasoning

which yields belief is wholly constituted by acts where those which are inferences and yield belief are not performed at will. That casts doubt upon whether such acts and such reasoning as a whole could amount to action. The causalist conception of inference then makes tempting the thought that the locus of agency exercised over one's beliefs in such cases must be in catalytic acts. The causalist conception of inference, that is, leads us to a falling-dominos-like picture of such reasoning on which the only actions involved play an initiating role, causing the reasoning to play out as it does. No further part for the agent to play is admitted.

Once we abandon the causalist conception of inference and endorse the alternative suggested, however, the above move to endorsement of the ballistics picture is no longer tempting. Suppose, for instance, that one reasons and concludes by inferring  $q$  from  $p$ , coming to believe  $q$  in doing so. On the account of inference recommended one's inference is the self-conscious event which is one's act of accepting  $q$  and which occurs as a constituent of the process which is one's reasoning. This event *just is* one's coming to believe  $q$ , and may even be one's coming to know that  $q$ . The event's being the kind of act of acceptance it is, furthermore, is determined by the self-knowledge which it constitutively involves. If we accept this, I suggest, then then it should no longer be tempting to think that if one is in any way active in so coming to believe then we must look to the prior causes of one's act of acceptance for an explanation of the extent to which this is so. Once we stop conceiving such acts of acceptance as events the nature of which is determined by how they are caused then it ceases to be natural to insist that if we are to account for the extent to which one could be active in so coming to believe this must be done via reference to how one's act was caused.

A comparison to potential cases of the exercise of doxastic agency through worldly inquiry might help to bring out the above thought. Suppose one searches for and eventually finds one's keys. One's search might be the result of a decision of one's to search for them. One's finding them might also be the result of one's pausing and attending to the matter of where one has not yet searched. But even supposing this is the case, there is no need to see oneself as active in finding one's keys only insofar as one's discovery is the result of such

causes. One's decision and one's act of attending lead to one's finding one's keys. But there is no temptation to say that if one is to count as exercising agency in doing so this is only because of the causal role played by one's decision and act of attending. There seems to be no reason to think that one is active when engaged in worldly inquiry only because of and insofar as it is initiated by such acts. Similarly, supposing one reasons and concludes by inferring  $q$  from  $p$ , coming to believe  $q$  in doing so, there should be no temptation to see oneself as exercising agency only because and insofar as what occurs is the product of the likes of decisions and acts of attending. Once we stop conceiving of inferences as acts of acceptance the nature of which is determined by how they are caused, furthermore, the temptation to do so is removed.

Conceiving of inference as suggested, rather than in accordance with the causalist conception, does more than remove the temptation to endorse the above ballistics picture. It also yields an alternative picture of the way in which we can be active with respect to our beliefs. Inference thus conceived is just as it needs to be to be such that it is apt to occur as a constituent of a process of reasoning where that reasoning amounts to the exercise of doxastic agency. Inference, on the account suggested, is a conscious event and act of acceptance the nature of which is determined by the self-knowledge which it constitutively involves. It is also an event the occurrence of which suffices for the occurrence of a process of reasoning. To understand how inference plays its role in doxastic agency we need to consider its place in such reasoning. We also need to conceive of the relevant reasoning appropriately.

We would not be doing so, I suggest, if we were to have it that when engaged in reasoning of the sort apt to amount to the exercise of doxastic agency one's aim is typically *to judge* or *to come to believe*, as we might think that one's aim when engaged in deliberation about what *to do* is *to decide* or *to form an intention*. That it is tempting to misconceive of the activity in this way is evidenced in the common description of the relevant activity as "making up one's mind". Clear examples of those who use the locution to describe the exercise of doxastic agency include Boyle (2011a), Chrisman (2016), Hampshire (1975), McDowell (2009), McHugh (2013) and O'Shaughnessy (2008b), to name just a few. I even do the same in my (2018). Paradigm cases of making up

one's mind are cases in which one *deliberates* about what *to do*. Accordingly, in such cases one's aim is *to decide* or to *form an intention*. I might, for instance, try to make up my mind on whether to have cake or the healthy alternative. In doing so my aim is to decide what to do and I will have achieved the aim of my deliberation only once I have decided. If we conceive of the exercise of doxastic agency via reasoning as parallel then we will think that the typical aim of such reasoning will be to judge or to come to believe.<sup>77</sup> Furthermore, with the reasoning thus conceived the observation that inference which yields belief is not performed at will casts doubt on the suggestion that we are capable of actively achieving the aim of such reasoning. Suppose that one reasons where one's aim is to judge or to come to believe. In so reasoning one sets out to judge or to come to believe. One has not actively achieved this aim, one might think, if one's reasoning terminates with an inference which yields belief and where one's so judging is not up to oneself—where in so inferring one does not judge at will. One has not actively done what one set out to do: to judge or to come to believe. At best, one might think, only part of the reasoning in question can amount to an exercise of agency.

Paradigm potential cases of the exercise of doxastic agency via reasoning are not most naturally conceived of in the above way. Such reasoning, that is, is not most naturally conceived of as performed with the aim of *judgment* or *belief*. This is evidenced by the more natural description of such reasoning's being, not 'making up one's mind', but 'figuring out ...', or 'working out ...'.<sup>78</sup> Take, for instance, the case in which one judges that the answer in a Sudoku game is either 3 or 5. One realises that the answer cannot be 3 and concludes one's reasoning on the matter by inferring that the answer is 5. Such a case is far more naturally described as one in which one *figures out* or *works out* what the answer is. As noted (sec. 5.2), one has not figured out whether *p* or worked out whether *p* unless one *knows* whether *p* and thus knows either that *p* or that not-*p*. Furthermore,

---

<sup>77</sup> Boghossian might be read as conceiving of reasoning this way in stating that no causal process counts as inference, unless it consists in an attempt to arrive at a belief by figuring out what, in some suitably broad sense, is supported by other things one believes. (2014, p. 5)

<sup>78</sup> Or at least apparently doing so. For ease of expression I focus on successful reasoning in what follows.

once one knows whether  $p$ , having come to know by reasoning, one has figured out whether  $p$ . This suggests that the typical aim of reasoning of the sort apt to potentially amount to the exercise of doxastic agency is not judgment, or belief, but *knowledge*, and that the aim of such reasoning is achieved if one acquires the relevant knowledge by reasoning. Furthermore, and crucially, with inference conceived as suggested inference can be seen to be just as it needs to be such that it can occur as a constituent of reasoning with such an aim where the reasoning can amount to the active achievement of that aim. Inferences which yield beliefs' not being performed at will presents no challenge to that suggestion. Suppose, for instance, that one reasons about whether  $q$  and concludes by inferring  $q$  from  $p$ . One's aim is to figure out whether  $q$ . One's aim, that is, is knowledge of whether  $q$ . Furthermore, one's concluding with inference to  $q$  can *just be* one's coming to know  $q$ . In performing the reasoning one can thus *just be* achieving the aim of one's reasoning. One's reasoning can amount to one's figuring out whether  $q$  and one can qualify as exercising doxastic agency in coming to know that  $q$  accordingly. In the case one sets out to *figure out whether  $q$*  where doing so requires that by reasoning one acquires *knowledge*. Actively achieving this aim by reasoning does not seem to require that one's concluding act amount to one's judging at will, as it might seem like it would need to in order for one to actively achieve one's aim had one set out to *judge* or to *come to believe*.

The suggestion here is that once we acknowledge that paradigm reasoning of the sort apt to amount to the exercise of doxastic agency is aimed at *knowledge* rather than *judgment* or *belief* it should no longer be tempting to see our inability to infer at will when inference yields belief as preventing such reasoning from amounting to the active achievement of one's aim. Supposing one were to engage in an activity the aim of which is to judge or come to believe—that one set out to judge or to come to believe—it seems that one's inability to judge or infer at will when doing so yields belief means that one will not manage to actively achieve one's aim in the direct way required for one to qualify as exercising doxastic agency (see sec. 1.5 above). If one sets out to judge, we might think, one has not actively achieved one's aim in the direct way required for doxastic agency unless one ends up making a judgment or coming to believe at

will—where one’s doing so is up to oneself. If one instead engages in an activity the aim of which is knowledge—one sets out, say, to find out whether  $q$ —one’s inability to infer at will when inference yields belief does not similarly cast doubt on one’s ability to actively achieve this aim in the direct way required for one to qualify as exercising doxastic agency. If one sets out to find out whether  $q$  one has actively done so so long as one’s reasoning amounts to one’s finding out whether  $q$ , even if the concluding inference which is one’s coming to know  $q$  or not- $q$  is not performed at will—if one’s so inferring is not up to oneself.

The point here may be best put in terms of its being better to conceive of reasoning which amounts to doxastic agency as primarily *epistemic* rather than *doxastic* agency (c.f. Soteriou 2013, Chapter 11). If we conceive of such reasoning as mere doxastic agency then we encourage the thought that the aim of such reasoning is to judge, where doing so yields belief, or to believe. The observation that when reasoning concludes with inference which yields belief one’s inference is not performed at will then makes it tempting to think that such reasoning cannot, as a whole, amount to the active achievement of one’s aim. It cannot amount to one’s achieving one’s aim to judge or to believe where the process through which one achieves the aim is action as a whole. If we conceive of cases of such reasoning as cases in which one sets out to judge or to believe, the thought is, then seeing one’s concluding inference as not being performed at will makes it tempting to think that the process which amounts to one’s achieving one’s aim is at best partially agential. One has not actively achieved one’s aim to judge or to believe, it is tempting to think, if the concluding inference which amounts to a judgment and yields belief is not performed at will—if one’s so judging is not up to oneself.

If we conceive of the relevant reasoning as primarily epistemic agency—as activity the aim of which is *knowledge*—then things look different. With inference so conceived it can be seen how reasoning which concludes with inference where one’s inference yields belief can amount to one’s coming to know—to perform such reasoning can *just be* to achieve the aim of such reasoning. And such reasoning’s not concluding with inference where one’s inference is performed at will does not prevent us from seeing one’s reasoning as amounting to one’s actively achieving one’s aim—from seeing one’s reasoning as a whole

from amounting to action. Suppose, for instance, that one reasons about whether  $q$ , concluding by inferring  $q$  from  $p$  and coming to know  $q$  by doing so. One's concluding inference is not performed at will. When one infers  $q$  from  $p$ , that is, one's doing so is not up to oneself. But there is no pressure to see things otherwise given that what one is actively trying to achieve is *knowledge* of whether  $q$ . To infer at will here, supposing one could do such a thing, would be to infer whilst seeing the constraint of treating  $q$  as true as self-imposed (see sec. 5.1). It would thus not be the way to achieve one's aim of knowledge (see below).

Again, the point here is particularly clear in potential cases of the exercise of doxastic agency through worldly inquiry. Suppose for instance, that I engage in worldly inquiry in order to find out whether my keys are in the house. I might feel around for them in the draw and judge that they indeed are in the draw. My judgment that my keys are in the draw here will not be performed at will—in judging that my keys are in the draw my so judging is not up to myself. But that does not prevent the case from being one in which I find out for myself whether my keys are in the house by engaging in inquiry. For me to find out for myself here it is my inquiry as a whole which needs to amount to action. And I am engaging in action by inquiring despite that the concluding act of my inquiry is not itself performed at will. Worldly inquiry can amount to epistemic agency—that is, to agency the aim of which is knowledge and which potentially amounts to the acquisition of knowledge—regardless of whether all constituent events of one's inquiry are performed at will. The same is true, I suggest, when it comes to the exercise of epistemic agency via conscious reasoning. It is one's inquiry as a whole which must be action if one is to exercise doxastic agency by inquiring. Similarly, to engage in a process of conscious reasoning can be to exercise doxastic agency regardless of whether all constituent events of one's reasoning are themselves action. It is one's reasoning as a whole which must be action.

Doxastic agency's being exercised via action the aim of which is *knowledge* in fact suggests that it would be no help to us when it comes to achieving that aim if inferences which yield beliefs could be performed at will. Were one such that one could make such inferences one would be such that one could infer  $q$  and come to believe  $q$  by doing so regardless of whether it were (actually or

apparently) evident to one that  $q$ . Such inferences would not be apt to yield knowledge that  $q$ .<sup>79</sup> It is because of this, I suggest, that judgment and inferences which yield beliefs' not being such that they can be performed at will does not yield a limit to our capacity to exercise doxastic agency. It is similarly because of this that our inability to believe or come to believe at will is no limit on our capacity to exercise doxastic agency. Doxastic agency is exercised, at least paradigmatically, through action the aim of which is *knowledge*. The ability to believe, to come to believe or to judge at will would be of no help when it comes to successfully engaging in such action.

My suggestion, in sum, is that we can exercise doxastic agency by exercising *epistemic agency*, the aim of such activity being *knowledge* rather than *belief* or *judgment*. Thus conceived it is not tempting to see reasoning's failure to conclude with inferences which are performed at will as preventing it from amounting to the active achievement of one's aim, as it might be tempting to see a failure to conclude at will as preventing a process the aim of which is to judge or to believe from amounting to the active achievement of one's aim. Furthermore, with inference conceived as suggested it is apt to occur as a constituent of reasoning which amounts to successful epistemic agency, since to infer can *just be* to come to know. Likewise, with inference conceived as suggested there is no longer any temptation to see the locus of agency as being in initiating acts when one engages in such reasoning, since when one infers one's act of accepting that which one infers is not an act of a kind the nature of which is determined by the way in which it is caused.

Conceiving of inference in the way recommended yields the above explanation of why the above ballistics picture is tempting but mistaken. It also yields an alternative picture of how we can be active with respect to our beliefs. The picture is one on which we are genuinely capable of doxastic agency. We can exercise doxastic agency by engaging in epistemic agency, where we can do that by reasoning and by inquiring. This, I suggest, amounts to a compelling case in favour of the account of inference offered. Once the alternative, non-

---

<sup>79</sup> McHugh (2011b, pp. 33–4) argues that being able to judge voluntarily would actually be a hindrance.



causalist account of inference offered is on the table as an option it might seem unsurprising that scepticism about doxastic agency and the above ballistics picture of the extent to which we are capable of agency over belief can be avoided. The causalist conception of inference is *made for* the advocate of the above ballistics picture. The causalist conception of inference and the above ballistics picture, that is, are *natural companions*. Faced by the challenge to the possibility of doxastic agency a natural response is that we can exercise it by reasoning. Paradigm reasoning which looks apt to amount to the exercise of doxastic agency concludes with inference which yields belief. But such inference is no more performed at will than we can believe or come to believe at will. On the causalist conception of inference the relevant concluding act of acceptance is conceived of as an event the nature of which is determined by the way in which it is caused by prior acceptances. When it comes to identifying the extent to which one can be exercising agency in coming to believe in the way in question the causalist is thus led to the causal history of one's inference. They are led to the ballistics picture on which we are best capable of actions which lead causally to belief revision. Once we see that there is an alternative account of inference available on which the act of acceptance in question is itself an inference, and an event of a kind which has its nature determined by the awareness which it constitutively involves, this line of reasoning is no longer tempting. It is no longer tempting to see the role of the agent in the above reasoning as at most being the bringing about of one's concluding act.

In response to the above an advocate of the causalist conception of inference might want to embrace scepticism about doxastic agency and the above ballistics picture of how we are active with respect to our beliefs. They might then insist that my alternative's not leading to scepticism need not move them to accept it.<sup>80</sup> However, the grounds for scepticism and for endorsement of the ballistics picture do not hold up once the alternative account of inference and its place in reasoning is made available. The grounds to be a sceptic about doxastic agency are removed by the mere availability of my alternative picture. One cannot then use one's scepticism as one's grounds for resisting the picture.

---

<sup>80</sup> This is how Koziolok reacted to my view in correspondence.

With inference and reasoning conceived as suggested we can see how to reason can just be to exercise doxastic agency where this does not require that one be able to infer at will when inference yields belief.

### 5.6. An alternative conception of doxastic agency

The standard approach to explaining how we can exercise doxastic agency despite that we cannot believe or come to believe at will—despite that when one believes or comes to believe one’s doing so is not up to oneself—is to insist that we can do so because *judgments* are or can be *actions* (e.g. Boghossian 2014; Cassam 2010; McDowell 2009; McHugh 2009, 2017; C. Peacocke 1999; Sosa 2015). This is maintained despite that judgments which yield beliefs are likewise not performed at will. To take this route is one way of attempting to answer the challenge to the possibility of doxastic agency in the way which Boyle (2009) sees as mandatory (see sec. 1.5 above). It is to suggest that we can exercise doxastic agency via the exercise of a special kind of agency. In Boyle’s terms, it is to assume that “defenders of the application of agential notions to belief” must appeal to and give an account of some “*other* notion of agency or control” which is at issue when doxastic agency is exercised (2009, p. 120). It must be explained how we can exercise doxastic agency via appeal to the exercise of agency which is different in kind to that exercised when one, say, raises one’s arm (again, see sec. 1.5 above). The supposed actions which are appealed to in order to explain how we exercise doxastic agency are actions of a kind which cannot be performed at will—actions of a kind which must be such that when they are performed one’s doing so is not up to oneself. It must be explained why we should consider them actions nonetheless.<sup>81</sup> Boyle’s account of doxastic agency itself amounts to a response of the sort in question of a further form. According to Boyle (2009, 2011a) *believing itself* is a kind of exercise of agency despite believing’s not being the kind of thing which we can do at will (see also Hieronymi 2009).

---

<sup>81</sup> See McHugh (2009) for a sustained attempt to do so. It is typical, in arguing for the claim that judgments are or can be actions, to emphasise the respect in which they are or can be responses to reasons, just as ordinary actions are or can be responses to reasons (e.g. McHugh 2009; McDowell 2009, p. 6).

In offering my view I have not actually denied that inferences which yield beliefs can amount to actions. I have instead suggested we can explain how doxastic agency is possible without appealing to special kinds of action as above. We can exercise doxastic agency by reasoning because reasoning can be *ordinary action* and because reasoning can have constituent events which *just are* events of coming to believe. In this way to reason can *just be* to exercise doxastic agency. A focus on *judgment* rather than inference might have obscured this option. If we see judgments as paradigm acts by which doxastic agency is exercised then when it comes to explaining how doxastic agency is possible we will find ourselves appealing to potentially isolated events of a kind which lack characteristic marks of action. I might, for instance, look out of the window and judge that the bird I see is a goldfinch. I might do so without engaging in any reasoning or inquiry. If cases in which doxastic agency is exercised can have such a structure then it becomes tempting to insist that such isolated judgments can be actions in order to save the possibility of doxastic agency. But if we instead focus on *inference* things look very different. As seen, inference is a conscious act of acceptance and conscious event of a kind which always occurs as a constituent of a process of conscious reasoning. Furthermore, I have suggested, such reasoning can be seen to amount to action and to the exercise of doxastic agency without our having to insist that all of the constituents of one's reasoning themselves be actions.

Thinking of things in the above way also makes natural the following stance on the agential status of the relevant inferences. When one reasons and concludes by inferring where one's inference yields belief one qualifies as exercising agency when inferring in virtue of the fact that by inferring one is engaging in a process of reasoning. Such inferences qualify as exercises of agency (whether or not it is appropriate to call them *actions*) in virtue of their occurring as constituents of processes where these processes as a whole amount to actions. The availability of such a stance removes the temptation to respond to scepticism to doxastic agency by insisting that isolated judgments must be actions.

### 5.7. The role of self-consciousness in doxastic agency

It was noted that it is natural to think that one must be self-conscious to be capable of exercising doxastic agency and that to exercise such agency is to exploit one's self-consciousness (sec. 1.5). O'Shaughnessy, for instance, suggests that we are "aware of our own minds" where this allows us to "rationally govern our existences" (2000, p. 112). He insists that our being self-conscious is part of what makes us capable of "[t]ranscending the condition of 'animal-immersion'" and of "mental freedom" where doxastic agency is one kind of manifestation of such freedom (2000, pp. 111–2). It was also seen to be puzzling why this should be so if it is in turn the case because we must be *aware of our beliefs* in order to exercise agency over them (sec. 1.5). Knowing what one believes might make it possible to act in ways which lead causally to intended changes in those beliefs. But it is unclear why knowing what one believes should make it possible for one to exercise doxastic agency proper—to settle questions for oneself.

This account of inference suggested, and the constraint that inference is self-conscious which led us to it, is motivated in part by the explanatory work it can be put to. And part of the explanatory work which the account can be put to is that of explaining why self-consciousness is required in order to be capable of doxastic agency and in what way to exercise doxastic agency is to exploit one's self-consciousness. Inference, I have suggested, is a distinctive kind of self-conscious act and a kind of self-conscious event of a kind which occurs, whenever it does, as a constituent of a process of conscious reasoning. Furthermore, to infer can *just be* to come to believe that which one infers (and/or that that which one infers follows from that which one infers from). Given this, to reason can *just be* to exercise doxastic agency. Reasoning can be action of a kind which can amount to the settling of questions for oneself.

Just as inference is a self-conscious constituent of reasoning it seems that reasoning is itself a kind of self-conscious action. The same seems to be the case with inquiry. Inquiry is a kind of self-conscious action. To be reasoning one must know that one is doing so under some guise and likewise to be inquiring one must know that one is doing so under some guise, that is, under some

description of what one is doing. Some might take this to be so on the basis of a generalisation about all action—on the basis of the claim that *all* action is self-conscious (e.g. Anscombe 1957). If one is  $\phi$ -ing where one's  $\phi$ -ing is an action, the thought would be, then one must know that one is  $\phi$ -ing under some description of the occurrence which is one's  $\phi$ -ing. Reasoning and inquiry being kinds of actions they must likewise be self-conscious.

Regardless of whether all action is self-conscious the suggestion that conscious reasoning is self-conscious action is particularly plausible. Reasoning is a kind of *conscious* activity which essentially involves *self-conscious* constituent occurrences. And when one reasons it is one's reasoning as a whole, rather than all of one's reasoning's constituent occurrences, which we must conceive of as action. Reasoning's being a kind of conscious action, whenever one reasons the question of whether one is reasoning seems to be a matter on which one is an authority. One's reasoning's being action of a sort one consciously engages in seems to ensure that one cannot be reasoning unbeknownst to oneself. If one is reasoning then one must know that one is doing so under some guise. It is likewise particularly plausible to think that worldly inquiry is self-conscious. It seems that one is not inquiring unless one knows that one is doing so under some guise. The question of whether one is inquiring, like the question of whether one is reasoning, is a matter on which one is necessarily an authority.

The above suggests the following account of the place of self-consciousness in doxastic agency. Self-consciousness is required for one to be capable of exercising doxastic agency because the exercise of doxastic agency is a kind of self-conscious activity. Doxastic agency is exercised through reasoning and inquiry. And reasoning and inquiry are kinds of self-conscious action. Similarly, to exercise doxastic agency is to exploit one's self-consciousness because doxastic agency is exercised via reasoning and inquiry and because reasoning and inquiry are themselves self-conscious activities.

It is thus a mistake to have it, as it is common to think, that one must be self-conscious in order to be a doxastic agent because knowledge of one's beliefs is required in order for one to exercise agency over them. In order to exercise doxastic agency one need not consider some belief which one has and

engage in action of a kind which will potentially result in a change in whether one holds the belief. All one need do is reason or inquire, where in reasoning and inquiring one's attention is typically directed at the *what is the case* and at *what follows from what*, not at one's beliefs. Such world directed reasoning and inquiry can amount to the exercise of doxastic agency without one's having to consider what one believes.

It has thus likewise been a mistake, as is also common to think, that it is exclusively *higher-order* reasoning by which doxastic agency is exercised—to think that it is only by reasoning about *whether one's beliefs are true or knowledgeable* rather than merely about *what is the case* that we exercise doxastic agency by reasoning. Moran suggests such a picture in claiming that

[o]ne is an agent with respect to one's attitudes insofar as one orients oneself toward the question of one's beliefs (2001, p. 64).<sup>82</sup>

In fact, one can exercise doxastic agency by reasoning or inquiring where this can perfectly well be *first-order* reasoning or inquiry—reasoning about or inquiry into what is the case. Self-consciousness is required for doxastic agency. But this is not so because one must know what one believes in order to exercise agency over one's beliefs. Nor is it so because doxastic agency is exercised via reasoning about whether one's beliefs are true or knowledgeable, where doing so requires knowledge of one's beliefs. Self-consciousness is required for doxastic agency because it is exercised via self-conscious action—via reasoning and inquiry where these are kinds of self-conscious actions.<sup>83</sup>

In the above account of the role of self-consciousness in doxastic agency I claim that *all* reasoning and inquiry is self-conscious action. Reflection on cases of so called “sub-intentional actions” yields a challenge to that claim. The actions of the relevant kind are those which are performed absent-mindedly.<sup>84</sup>

---

<sup>82</sup> To be fair to Moran, he does add that we standardly engage in such reasoning simply by “reflecting on what's true” (2001, p. 64).

<sup>83</sup> McHugh (2013, p. 134) is in agreement that we can exercise doxastic agency via first-order reasoning. What I have added is an explanation of how this is so and how to exercise doxastic agency is nonetheless to exploit one's self-consciousness.

<sup>84</sup> There is a dispute concerning whether all such actions must be intentional or whether they can be actions and yet fail to be performed intentionally. O'Shaughnessy (2008b, Chapter 10) accepts the former view. All actions, he claims, are intentional under some description. Steward (2009), meanwhile, accepts the former view. She has it that not all actions are intentional.

Typical examples of absent-minded actions include idly drumming one's fingers as one talks and twiddling with one's hair as one writes. These are actions of a kind which we can *catch ourselves* doing. Furthermore, it might be thought that this is best captured by saying that they are actions of a kind which we can *come to know* are occurring and thus that they need not be self-conscious (e.g. Steward 2009). Moreover, it seems clear that reasoning and inquiry can occur absent-mindedly—we can catch ourselves reasoning or inquiring having been doing so absent-mindedly. Again, it might be thought that this is best captured by saying that absent-minded reasoning and inquiry are actions of a kind which we can come to know are occurring and thus that reasoning and inquiry need not be self-conscious after all.

Some will want to resist the suggestion that the fact that we can reason and inquire absent-mindedly and catch ourselves doing so is best captured by saying that we can reason and inquire without knowing that we are doing so under *any* guise and then come to know what we are up to. This might instead be captured in terms of shifts in attention (c.f. O'Shaughnessy 2008b, p. 361; Lynch 2014, p. 72). When I reason I do so self-consciously, as is suggested by reasoning's being a kind of conscious action. It may nonetheless be absent-minded—in the periphery of my attention. I can catch myself reasoning, the suggestion would go, by directing my attention at my doing so. In doing this I can come to know what I am doing under “additional conceptual determinations” (O'Shaughnessy 2008b, p. 361). Similarly, when I inquire I do so self-consciously. It may nonetheless be absent-minded and such that I can catch myself inquiring by coming to know what I am doing under additional descriptions.

Even if it is admitted that it is possible to reason and inquire without knowing that one is doing so—without knowledge under any description of the *process* which is one's reasoning or inquiry—it does not follow that it is possible for non-self-conscious subjects to exercise doxastic agency. In fact, it seems clear that one will only be able to reason and inquire non-self-consciously if one is likewise capable of doing so self-consciously. Alleged cases of non-self-conscious reasoning and inquiry will be cases in which one reasons or inquires absent-mindedly and thus *inattentively*. And one will only be such that one can reason and inquire inattentively such that one can do so unknowingly if one can

potentially likewise do so *attentively* and thus self-consciously. Absent-minded and non-self-conscious reasoning and inquiry (supposing there is such a thing) would only be apt to qualify as actions if they were the kind of thing which upon catching oneself doing them one may be able to continue to do them self-consciously. Something which is not even the kind of thing which one can do self-consciously does not seem to be something which we should say is a kind of action. Even if it is admitted that not all reasoning and inquiry is self-conscious it will thus remain possible to maintain that only a subject who is capable of reasoning and inquiring self-consciously will be capable of exercising doxastic agency by reasoning and inquiring.

Furthermore, and crucially, even if it is admitted that it is possible to reason and inquire without knowing that one is doing so it does not follow that it is possible to exercise doxastic agency without exploiting one's self-consciousness whatsoever. In fact, given the above account of inference it seems clear that it is *not* possible to do so. Even if reasoning can occur without doing so self-consciously it does have essentially self-conscious constituents. Inference, as seen, is a distinctive kind of conscious and self-conscious occurrence. Any reasoning which involves inference will thus manifest one's self-consciousness. Similarly, even if inquiry can be engaged in without one's inquiring self-consciously, inquiry does have essentially self-conscious constituents. Even if it is admitted that not all reasoning and inquiry is self-conscious it can thus be maintained that doxastic agency requires self-consciousness, and that one exploits one's self-consciousness whenever one exercises doxastic agency, since reasoning and inquiry are the means to exercise doxastic agency and since reasoning and inquiry essentially involve self-conscious constituents.

Admittedly, to establish the above claim it would need to be established that more than just inference is essentially self-conscious. Focus on inference would need to be lifted and attention directed to the various other kinds of occurrences which can be constitutively involved in reasoning and inquiry. In the reasoning case, for instance, attention would need to be directed to the kinds of occurrences involved in the construction of theories and conducting of thought experiments, these being other potential means to exercise doxastic agency by



reasoning. It is far from clear that all such reasoning must involve inference.<sup>85</sup> Nonetheless, I take to have made plausible the suggestion that just as inference ought to be seen as a distinctive kind of self-conscious occurrence the same will turn out to be true when we turn to the other kinds of acts involved in reasoning and inquiry. By focusing on inference we have focused on a paradigm case and in turn made plausible a more general claim: that reasoning and inquiry necessarily involve self-conscious occurrences.

### 5.8. Incidental agency and the aim of inquiry

Of concern has been how we can exercise doxastic agency despite that we can neither believe nor come to believe at will. Inferences, I have conceded, are not performed at will when they yield beliefs either. But inferences are events of a kind which occur as constituents of processes of reasoning. And reasoning can be seen to amount to a kind of action. Furthermore, I have suggested, to infer can *just be* to come to believe that which one infers. Given this, reasoning can be action where to perform such reasoning *just is* to exercise doxastic agency, rather than merely being action of a kind which leads causally to belief (see sec. 5.2 above).

This suggestion might prompt the following objection, engagement with which will help to reveal how to further develop the account offered. The objection is that the account offered might seem to entail that actions which *incidentally* involve judgments as constituents amount to exercises of doxastic agency. It might be thought, however, that they do not. Performing an action a constituent of which merely happens to be a judgment which yields a belief does not suffice for exercising doxastic agency. Suppose, for instance, that whilst looking around for a lost frisbee I happen to notice a goldfinch at the end of the garden. My noticing may *just be* my coming to believe that there is a goldfinch. But one might think that this event is also a constituent of the process of my actively looking around. We seem to have a case in which an event of coming to believe is a constituent of an action, but not obviously a case in which

---

<sup>85</sup> As Grice (2001, pp. 18–20) notes, some might be hesitant when it comes to labelling such activities reasoning. If they are not then they will turn out to be further ways to exercise doxastic agency in addition to reasoning and inquiry.

doxastic agency is exercised in coming to believe that there is a goldfinch. Or suppose that whilst thinking about my plans for the day I note the time and suddenly realise that I am late for a meeting. One might think that my realising that I am late *just is* my coming to believe that I am late as well as a constituent of the activity which is my thinking about my plans for the day. We again seem to have a case in which an event of coming to believe is a constituent of an action, but perhaps not a case in which doxastic agency is exercised.

In sum the issue is that on the view suggested reasoning and inquiry are means to exercise doxastic agency because they are actions of a sort which can have events of belief fixation as constituents. This suggestion may seem to overgeneralise in telling us that *any* action which happens to have an event of belief fixation as a constituent is an exercise of doxastic agency.

I do not want to respond to the above worry by insisting that any action which has an event of belief fixation as a constituent really is an exercise of doxastic agency. What I take myself to have shown is how it is *possible* to exercise doxastic agency via reasoning and inquiry in the face of challenges. I have not said what *suffices* for the exercise of doxastic agency and rather only suggested that reasoning and inquiry *can* amount to the exercise of doxastic agency, explaining how this is so in the face of challenges. That said, there is a natural suggestion available when it comes to saying why the occurrence of reasoning and inquiry can suffice for the exercise of doxastic agency and why they are our distinctive means to exercise such agency.<sup>86</sup>

Reasoning and inquiry can amount to *intentional actions* under descriptions like ‘figuring out whether *p*’ and ‘finding out whether *p*’. When this is the case the *aim* of one’s reasoning is to figure out or to find out whether *p* and thus to *come to know whether p* (see sec. 5.2 above).<sup>87</sup> Furthermore, one’s reasoning or inquiry can amount to the intentional achievement of this aim. Reasoning and inquiry can be intentional actions the aim of which is to find out whether *p*. One can conclude such reasoning or inquiry by inferring *p* (or not-*p*) where one’s

---

<sup>86</sup> Much of what follows draws on McHugh (2011b).

<sup>87</sup> McHugh (2011b) has it that inquiry always has such an aim. Inquiry is an activity which constitutively aims at truth or knowledge—if an activity does not have such an aim then it is thereby not inquiry (see also McHugh and Way forthcoming; Ryle 2000).

doing so *just is* one's coming to know that  $p$  (or that not- $p$ ). To engage in such reasoning or inquiry can thus *just be* to intentionally achieve the aim of one's reasoning or inquiry. That reasoning and inquiry can amount to the achieving of one's aim of knowledge in this way, I suggest, makes clear that they can amount to the exercise of doxastic agency.<sup>88</sup> The view offered is thus not that any actions with events of belief fixation as constituents amount to doxastic agency. It is rather that reasoning and inquiry's having events of belief fixation as constituents makes them apt to amount to the intentional achievement of one's aim of knowledge. And that reasoning and inquiry can amount to that makes it clear that they can be means to exercise doxastic agency.

It is because reasoning and inquiry can be intentional actions the aim of which is knowledge that it is clear that they can amount to the exercise of doxastic agency. It is likewise because they have such an aim, I would suggest more tentatively, that we can be due credit for coming to know via reasoning and inquiry and due blame for failures to know. Grice suggests that

[f]or activities of a certain sort to be something one can be (more or less) good at, they must be directed towards goals (2001, p. 21).

Reasoning and inquiry's being something which we can do with the aim of knowledge and which can amount to or fail to amount to the intentional achievement of that aim makes it apt to be seen as something which we can be good/bad at, which we can do well/badly, and which we can be due praise/blame for in cases of success/failure.<sup>89</sup>

Reasoning and inquiry's being such that they can be performed in pursuit of the aim of knowledge in a way which makes us apt to be due credit for our

---

<sup>88</sup> I have suggested that it is possible for reasoning and inquiry to amount to the exercise of doxastic agency and now suggested that they do when they amount to the intentional achievement of the aim of one's reasoning or inquiry. I leave open whether more reasoning and inquiry than that which amounts to the intentional achievement of one's aim amounts to the exercise of doxastic agency. I leave open, for instance, whether there is sub-intentional reasoning and inquiry which can amount to doxastic agency (see note 84 above) and whether one can exercise doxastic agency in cases in which knowledge in the relevant domain is not one's aim (e.g. when one reasons about one thing and figures out something else, when one reasons without a specific aim and rather just "follows the evidence", or when one's aim is not knowledge but to make some relevant assumption for practical purposes).

<sup>89</sup> Our being such that we can be due credit when we come to know by reasoning and inquiry and due blame when we fail to do so can in turn be seen as the grounds for our being such that we can be held responsible for our beliefs (see McHugh 2013).

successes and due blame for our failures might be taken to require that reasoning and inquiry be performed in at least apparent awareness of its correctness and of its potentially amounting to means to achieve that aim (e.g. Brewer 1995, p. 243). But just as inferring should not be taken to require prior (real or apparent) awareness of relevant inferential connections (see sec. 4.11 above) reasoning and inquiry should not be seen as having to be performed in the light of prior (real or apparent) awareness of their correctness and of their potentially amounting to means to achieve one's aim (Brewer 1995, sec. 2). Instead, I would suggest, one can manifest one's (real or apparent) appreciation of the correctness of one's reasoning or inquiry and one's appreciation of its being a means to achieve the aim of one's reasoning or inquiry *by reasoning or inquiring* (c.f. Soteriou 2013, Chapter 11). It is because reasoning and inquiry can be performed in the pursuit of knowledge and can be performed in real or apparent awareness of their correctness and of their potentially amounting to means to achieve that aim that reasoning and inquiry can be seen as something which we can be good or bad at, which we can do well or badly, and which we can be legitimately praised for in cases of success and blamed for in cases of failure.

## 5.9. Conclusion

The assumption that inference is self-conscious has yielded an account of inference on which it is a distinctive kind of conscious and self-conscious event, the nature of which, contrary to orthodoxy, is determined not by the way in which it is caused, nor by what it causally involves, but by the self-knowledge which it constitutively involves. What results is:

- A resolution of the ontological difficulty with which we began. To infer  $q$  from  $p$  one must accept  $p$  and then accept  $q$ . The conscious act which is one's inference, when one does so, is the conscious *event* which is one's acceptance of  $q$ . This claim can be reconciled with the fact that the occurrence of an inference requires the occurrence of *two* distinct events of acceptance. One's inference, when one infers from  $p$  to  $q$ , is the event which is one's accepting  $q$  where this event's qualifying as an inference requires that it be preceded by an event which is one's

conscious acceptance of that which one infers from. It has been a mistake to think, as is orthodox, that when one infers  $q$  from  $p$  one's inference is the occurrence which begins with one's acceptance of  $p$  and terminates with one's acceptance of  $q$ —a view commonly expressed with the claim that inference is a *process*.

- An account of how inference plays its cognitive and epistemic role. Inference can yield belief and knowledge, not because inference can potentially lead causally to knowledgeable belief, but because to infer can *just be* to come to believe and to come to know. This reveals that the place of taking beliefs in inference has also been widely misunderstood. To infer one must believe that an appropriate inferential relation obtains between that which one infers and that which one infers from. But this belief need not play a causal role in one's inferring. In fact, one can infer without having a corresponding taking belief prior to one's inferring. In such a case one will acquire such a belief and potentially acquire knowledge of the obtaining of an appropriate inferential relation *by inferring*.

Most crucially, the account of inference also yields:

- An explanation of why it is tempting to move from the challenge to the possibility of doxastic agency to a view on which we are at best capable of “mental ballistics”—of actions of a kind which lead causally to belief revision.
- A way to remove the temptation to move to that view and to instead embrace a view on which to reason can *just be* to achieve the aim of one's reasoning where that aim is *knowledge*. To reason and conclude with inference can just be to achieve one's aim of knowledge because to infer can just be to come to know that which one infers. When one makes such an inference the nature of the occurrence is not determined by the way in which it is caused, but by the self-knowledge it constitutively involves. Accordingly, there should be no temptation to look only to the causal history of the occurrence to account for the extent to which one is active in coming to know in the way in question.

On the view which results we can exercise doxastic agency by reasoning and inquiring where there is no need to see such reasoning or inquiry as needing to terminate with judgments which are performed at will.

- An explanation of why self-consciousness is required for and exploited in the exercise of doxastic agency. This is so, not because one must know what one believes in order to exercise agency over one's beliefs, nor because doxastic agency is exercised via reasoning about one's beliefs. It is rather because the exercise of doxastic agency is, at least paradigmatically, *itself a self-conscious activity* and because it is an activity which necessarily has *self-conscious constituent constituents*.

I take it that the above amounts to a compelling account of the nature of inference and likewise to a compelling account of doxastic agency in the face of challenges to its very possibility. That the account of inference can be put to such explanatory work, I would have it, suffices to motivate the starting presumption that inference is necessarily self-conscious. The account of doxastic agency offered could be extended via consideration of the exercise of doxastic agency via worldly inquiry and by lifting the above focus on inference. Inference, after all, is just one kind of constituent of reasoning and inquiry. Focus on inference, however, has been instructive. I will leave the task of providing a fuller account of doxastic agency in its forms to another day. A clear exemplification of the exercise of doxastic agency, and a paradigm case of its exercise, can be seen by consideration of cases in which reasoning concludes with inference, one's concluding inference yielding belief in and potentially knowledge of that which one infers. Such reasoning can amount to the intentional achievement of one's aim of knowledge.

## Bibliography

- Alston, W. P. (1988). The Deontological Conception of Epistemic Justification. *Philosophical Perspectives*, 2, 257–299.
- Alvarez, M. (2010). *Kinds of Reasons: An Essay in the Philosophy of Action*. Oxford: Oxford University Press.
- Anscombe, G. E. M. (1957). *Intention*. Cambridge MA: Harvard University Press.
- Anscombe, G. E. M. (1971). *An Introduction to Wittgenstein's Tractatus*. St. Augustine's Press.
- Antony, L. M. (1989). Anomalous Monism and the Problem of Explanatory Force. *Philosophical Review*, 98(April), 153–87.
- Barnett, D. J. (2016). Inferential Justification and the Transparency of Belief. *Noûs*, 50(1), 184–212.
- Bennett, O. (1990). Why is belief involuntary? *Analysis*, 50(2), 87–107.
- Boghossian, P. A. (1997). Analyticity. In B. Hale & C. Wright (Eds.), *A Companion to the Philosophy of Language* (pp. 331–368). Oxford: Blackwell.
- Boghossian, P. A. (2000). Knowledge of Logic. In P. A. Boghossian & C. Peacocke (Eds.), *New Essays on the A Priori* (pp. 229–254). Oxford: Oxford University Press.
- Boghossian, P. A. (2001). Inference and Insight. *Philosophy and Phenomenological Research*, 63(3), 633–640.
- Boghossian, P. A. (2003). Blind reasoning. *Aristotelian Society supplementary volume*, 77(1), 225–248.
- Boghossian, P. A. (2014). What is inference? *Philosophical Studies*, 169(1), 1–18.

- Boghossian, P. A. (2015). Reasoning and Reflection: A Reply to Kornblith. *Analysis*. doi:10.1093/analys/anv031
- Boyle, M. (2009). Active Belief. *Canadian Journal of Philosophy*, 39(sup1), 119–147.
- Boyle, M. (2011a). “Making Up Your Mind” and the Activity of Reason. *Philosophers’ Imprint*, 11(17).
- Boyle, M. (2011b). Transparent Self-Knowledge. *Aristotelian Society Supplementary Volume*, 85(1), 223–241.
- Boyle, M. (2012). Essentially rational animals. In G. Abel & J. Conant (Eds.), *Rethinking Epistemology* (Vol. 2, pp. 395–427). Berlin: Walter De Gruyter.
- Brandom, R. B. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge MA: Harvard University Press.
- Bratman, M. E. (1992). Practical Reasoning and Acceptance in a Context. *Mind*, 101(401), 1–16.
- Brewer, B. (1995). Mental Causation: Compulsion by Reason. *Aristotelian Society Supplementary Volume*, 69(69), 237–253.
- Brewer, B. (2011). *Perception and Its Objects*. Oxford: Oxford University Press.
- Broome, J. (2013). *Rationality Through Reasoning*. Oxford: John Wiley & Sons.
- Broome, J. (2014). Comments on Boghossian. *Philosophical Studies*, 169(1), 19–25.
- Call, J. (2001). Object permanence in orangutans (*Pongo pygmaeus*), chimpanzees (*Pan troglodytes*), and children (*Homo sapiens*). *Journal of Comparative Psychology*, 115(2), 159–171.
- Call, J. (2004). Inferences about the location of food in the great apes (*Pan paniscus*, *Pan troglodytes*, *Gorilla gorilla*, and *Pongo pygmaeus*). *Journal of Comparative Psychology*, 118(2), 232–241.



- Carroll, L. (1895). What the Tortoise said to Achilles. *Mind*, 4(14), 278–280.
- Cassam, Q. (2010). Judging, Believing and Thinking. *Philosophical Issues*, 20(1), 80–95.
- Cassam, Q. (2014). *Self-Knowledge for Humans*. Oxford: Oxford University Press.
- Chrisman, M. (2016). Epistemic Normativity and Cognitive Agency. *Nous*. doi:10.1111/nous.12184
- Chudnoff, E. (2014). The Rational Roles of Intuition. In A. Booth & D. Rowbottom (Eds.), *Intuitions* (pp. 9–35). Oxford: Oxford University Press.
- Crane, T. (1995). The Mental Causation Debate. *Proceedings of the Aristotelian Society*, 69, 211–36.
- Crane, T. (2013). Unconscious Belief and Conscious Thought. In U. Kriegel (Ed.), *Phenomenal Intentionality* (pp. 156–173). Oxford: Oxford University Press.
- Crowther, T. (2011). The Matter of Events. *The Review of Metaphysics*, 65(1), 3–39.
- Dancy, J. (2014). From Thought to Action. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Volume 9* (pp. 1–17). Oxford: Oxford University Press. Accessed 31 August 2016
- Davidson, D. (2001a). Actions, Reasons, and Causes. In *Essays on actions and events* (pp. 1–20). Oxford: Oxford University Press.
- Davidson, D. (2001b). Freedom to Act. In *Essays on actions and events* (pp. 63–82). Oxford: Oxford University Press.
- Davidson, D. (2001c). The Individuation of Events. In *Essays on Actions and Events* (pp. 163–180). Oxford: Oxford University Press.

- Dogramaci, S. (2016). Reasoning without Binders: A Reply to Valaris. *Mind*.  
doi:10.1093/mind/fzv155
- Dummett, M. A. E. (1991). *The Logical Basis of Metaphysics*. Cambridge MA:  
Harvard University Press.
- Evans, G. (1982). *Varieties of Reference*. Oxford: Oxford University Press.
- Evans, J. S. B. T. (2003). In two minds: dual-process accounts of reasoning.  
*Trends in Cognitive Sciences*, 7(10), 454–459.
- Fumerton, R. (2004). Epistemic Probability. *Philosophical Issues*, 14(1), 149–164.
- Fumerton, R. (2015). What the Internalist Should Say to the Tortoise. *Episteme*,  
12(Special Issue 02), 209–217.
- Ginet, C. (2001). Deciding to Believe. In M. Steup (Ed.), *Knowledge, Truth and  
Duty* (pp. 63–76). Oxford University Press.
- Golob, S. (2015). Self-Knowledge, Transparency, and Self-Authorship.  
*Proceedings of the Aristotelian Society*, 115(3pt3), 235–253.
- Grice, H. P. (1957). Meaning. *Philosophical Review*, 66(3), 377–388.
- Grice, H. P. (1961). The Causal Theory of Perception, Part I. *Proceedings of the  
Aristotelian Society*, 121, 121–152.
- Grice, H. P. (2001). *Aspects of Reason*. Oxford: Oxford University Press.
- Hampshire, S. (1975). *Freedom of the Individual*. Princeton NJ: Princeton  
University Press.
- Harman, G. (1986). *Change in View*. Cambridge MA: MIT Press.
- Hieronymi, P. (2006). Controlling Attitudes. *Pacific Philosophical Quarterly*, 87(1),  
45–74.
- Hieronymi, P. (2009). Believing at Will. *Canadian Journal of Philosophy*,  
*Supplementary Volume*, 35(sup1), 149–187.

- Hieronymi, P. (2014). Reflection and Responsibility. *Philosophy and Public Affairs*, 42(1), 3–41.
- Hlobil, U. (2016). *What Is Inference? Or the Force of Reasoning*. University of Pittsburgh.
- Hornsby, J. (1980). *Actions*. London: Routledge and Kegan Paul.
- Hornsby, J. (2008). A Disjunctivist Conception of Acting for Reasons. In A. Haddock & F. Macpherson (Eds.), *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press.
- Hornsby, J. (2012). Actions and Activity. *Philosophical Issues*, 22(1), 233–245.
- Hornsby, J. (2013). Basic Activity. *Aristotelian Society Supplementary Volume*, 87(1), 1–18.
- Huemer, M. (2016). Inferential Appearances. In B. Coppenger & M. Bergmann (Eds.), *Intellectual Assurance: Essays on Traditional Epistemic Internalism* (pp. 144–160). Oxford: Oxford University Press.
- James, W. (1976). The Experience of Activity. In *Essays in Radical Empiricism*. Cambridge MA: Harvard University Press.
- Jenkins, D. (forthcoming). Luminosity in the Stream of Consciousness. *Synthese*.
- Jenkins, D. (2018). The Role of Judgment in Doxastic Agency. *Thought: A Journal of Philosophy*, 7(1), 12–19.
- Johnson-Laird, P. (2008). *How We Reason*. Oxford: Oxford University Press.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. London: Penguin.
- Kitcher, P. (2011). *Kant's Thinker*. Oxford University Press.
- Kornblith, H. (2012). *On Reflection*. Oxford: Oxford University Press.
- Korsgaard, C. M. (2008). *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford: Oxford University Press.

- Korsgaard, C. M. (2009). The Activity of Reason. *Proceedings and Addresses of the American Philosophical Association*, 83(2), 23–43.
- Koziolok, N. (2017). Inferring as a way of knowing. *Synthese*, 1–20.  
doi:10.1007/s11229-017-1632-4
- Lee, W. (forthcoming). Reasoning, rational requirements, and occurrent attitudes. *European Journal of Philosophy*. doi:10.1111/ejop.12342
- Leite, A. (2008). Believing One's Reasons Are Good. *Synthese*, 161(3), 419–441.
- Leite, A. (2016). Changing One's Mind: Self-Conscious Belief and Rational Endorsement. *Philosophy and Phenomenological Research*, 93(3).
- Littlejohn, C. (forthcoming). A Plea for Epistemic Excuses. In F. D. J. Dutant (Ed.), *The New Evil Demon Problem*. Oxford University Press.
- Lynch, K. (2014). Self-deception and shifts of attention. *Philosophical Explorations*, 17(1), 63–75.
- Marr, D. (2010). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Cambridge MA: The MIT Press.
- Martin, M. G. F. (2002). The Transparency of Experience. *Mind and Language*, 4(4), 376–425.
- McDowell, J. (1994). *Mind and World*. Cambridge MA: Harvard University Press.
- McDowell, J. (2009). *Having the World in View: Essays on Kant, Hegel, and Sellars*. Cambridge MA: Harvard University Press.
- McGinn, C. (1996). *The Character of Mind: An Introduction to the Philosophy of Mind*. New York: Oxford University Press.
- McHugh, C. (2009). Judging as a non-voluntary action. *Philosophical Studies*, 152(2), 245–269.

- McHugh, C. (2011a). What Do We Aim At When We Believe? *Dialectica*, 65(3), 369–392.
- McHugh, C. (2011b). Exercising Doxastic Freedom. *Philosophy and Phenomenological Research*, 88(1), 1–37.
- McHugh, C. (2013). Epistemic Responsibility and Doxastic Agency. *Philosophical Issues*, 23(1), 132–157.
- McHugh, C. (2017). Attitudinal control. *Synthese*, 194(8), 2745–2762.
- McHugh, C., & Way, J. (forthcoming). What is Reasoning? *Mind*. doi:10.1093/mind/fzw068
- McHugh, C., & Way, J. (2016). Against the Taking Condition. *Philosophical Issues*, 26(1), 314–331.
- Mole, C. (2018). The Process of Inference. In R. Stout (Ed.), *Process, Action, and Experience*. Oxford University Press.
- Moran, R. A. (2001). *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton NJ: Princeton University Press.
- Mourelatos, A. P. D. (1978). Events, processes, and states. *Linguistics and Philosophy*, 2(3), 415–434.
- Mugg, J. (2015). The Dual-Process Turn: How Recent Defenses of Dual-Process Theories of Reasoning Fail. *Philosophical Psychology*, 29(2), 300–309.
- Nes, A. (2016). The Sense of Natural Meaning in Conscious Inference. In T. Breyer & C. Gutland (Eds.), *Phenomenology of Thinking* (pp. 97–115). New York: Routledge.
- Neta, R. (2013). What Is an Inference? *Philosophical Issues*, 23(1), 388–407.

- O'Brien, L. (2005). Self-Knowledge, Agency and Force. *Philosophy and Phenomenological Research*, 71(3), 580–601.
- O'Shaughnessy, B. (2000). *Consciousness and the World*. Oxford: Clarendon Press.
- O'Shaughnessy, B. (2008a). *The Will: Volume 1, Dual Aspect Theory* (2nd ed.). Cambridge: Cambridge University Press.
- O'Shaughnessy, B. (2008b). *The Will: Volume 2, a Dual Aspect Theory* (2nd ed.). Cambridge: Cambridge University Press.
- Owens, D. (2000). *Reason Without Freedom: The Problem of Epistemic Normativity*. Oxford: Routledge.
- Parrott, M. (forthcoming). The Look of Another Mind. *Mind*.
- Peacocke, A. (2017). Embedded mental action in self-attribution of belief. *Philosophical Studies*, 174(2), 353–377.
- Peacocke, C. (1993). How Are A Priori Truths Possible?1. *European Journal of Philosophy*, 1(2), 175–199.
- Peacocke, C. (1999). *Being Known*. Oxford: Oxford University Press.
- Peacocke, C. (2003). Implicit Conceptions, Understanding, and Rationality. In M. Hahn & B. Ramberg (Eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (pp. 117–152). Cambridge MA: MIT Press.
- Petit, O., Dufour, V., Herrenschmidt, M., Marco, A. D., Sterck, E. H. M., & Call, J. (2015). Inferences about food location in three cercopithecine species: an insight into the socioecological cognition of primates. *Animal Cognition*, 18(4), 821–830.
- Pink, T. (2003). Action and Self-Determination. *Intellectica*, 36–7, 247–259.
- Rödl, S. (2007). *Self-Consciousness*. Cambridge MA: Harvard University Press.

- Rödl, S. (2013). The Single Act of Combining. *Philosophy and Phenomenological Research*, 87(1), 213–220.
- Rosa, L. (forthcoming). Reasoning Without Regress. *Synthese*, 1–16.
- Rumfitt, I. (2015). *The Boundary Stones of Thought: An Essay in the Philosophy of Logic*. Oxford: Oxford University Press.
- Ryle, G. (2000). Courses of Action or the Uncatchableness of Mental Acts. *Philosophy*, 75(3), 331–344.
- Ryle, G. (2009a). *The Concept of Mind*. London: Routledge.
- Ryle, G. (2009b). The Thinking of Thoughts: What is “Le Penseur” Doing? In *Collected Essays* (Vol. 2, pp. 494–510). London: Routledge.
- Ryle, G. (2009c). Thinking and Reflecting. In *Collected Essays* (Vol. 2, pp. 479–493). London: Routledge.
- Sartre, J.-P. (2003). *Being and Nothingness: An Essay on Phenomenological Ontology*. London: Routledge.
- Sellars, W. (1967). Some Remarks on Kant’s Theory of Experience. In J. Sicha (Ed.), *Kant’s Transcendental Metaphysics: Sellars’ Cassirer Lecture Notes and Other Essays* (pp. 269–282). Atascero CA: Ridgeview Publishing Company.
- Setiya, K. (2013). Epistemic Agency: Some Doubts. *Philosophical Issues*, 23(1), 179–198.
- Shah, N., & Velleman, J. D. (2005). Doxastic Deliberation. *Philosophical Review*, 114(4), 497–534.
- Shepherd, J. (2015). Deciding as Intentional Action: Control over Decisions. *Australasian Journal of Philosophy*, 93(2), 335–351.

- Siegel, S. (forthcoming). Inference Without Reckoning. In B. B. Jackson & M. B. Jackson (Eds.), *Reasoning: Essays in Theoretical and Practical Thinking*. Oxford University Press.
- Snowdon, P. (1980). Perception, Vision and Causation. *Proceedings of the Aristotelian Society*, 81, 175–192.
- Sosa, E. (2015). *Judgment and Agency*. Oxford: Oxford University Press.
- Soteriou, M. (2013). *The Mind's Construction: The Ontology of Mind and Mental Action*. Oxford: Oxford University Press.
- Steward, H. (1997). *The Ontology of Mind: Events, Processes, and States*. Oxford: Oxford University Press.
- Steward, H. (2009). Sub-Intentional Actions and the Over-Mentalization of Agency. In C. Sandis (Ed.), *New Essays on the Explanation of Action*. London: Palgrave-Macmillan.
- Steward, H. (2013). Processes, Continuants, and Individuals. *Mind*, 122(487), 781–812.
- Stout, R. (1997). Processes. *Philosophy*, 72(279), 19–27.
- Strawson, G. (2003). Mental Ballistics or the Involuntariness of Spontaneity. *Proceedings of the Aristotelian Society*, 103(3), 227–257.
- Stroud, B. (1979). Inference, Belief, and Understanding. *Mind*, 88(350), 179–196.
- Taylor, S. (2015). Is Justification Easy or Impossible? Getting Acquainted with a Middle Road. *Synthese*, 192(9), 2987–3009.
- Thomson, J. J. (1965). Reasons and Reasoning. In M. Black (Ed.), *Philosophy in America*. Ithaca: Cornell University Press.



- Thurrow, J. C. (2008). The a priori defended: a defense of the generality argument. *Philosophical Studies*, 146(2), 273–289.
- Toribio, J. (2011). What We Do When We Judge. *Dialectica*, 65(3), 345–367.
- Valaris, M. (2011). Transparency as Inference: Reply to Alex Byrne. *Proceedings of the Aristotelian Society*, 111(2pt2), 319–324.
- Valaris, M. (2014). Reasoning and Regress. *Mind*, 123(489), 101–127.
- Valaris, M. (2016a). Supposition and Blindness. *Mind*, 125(499), 895–901.
- Valaris, M. (2016b). What reasoning might be. *Synthese*, 1–18.  
doi:10.1007/s11229-016-1034-z
- Valaris, M. (2016c). What The Tortoise Has To Say About Diachronic Rationality. *Pacific Philosophical Quarterly*. doi:10.1111/papq.12172
- Vendler, Z. (1957). Verbs and Times. *The Philosophical Review*, 66(2), 143–160.
- White, A. R. (1971). Inference. *Philosophical Quarterly*, 21(85), 289–302.
- Williams, B. (1970). Deciding to Believe. In B. Williams (Ed.), *Problems of the Self* (pp. 136–51). Cambridge: Cambridge University Press.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wittgenstein, L. (1969). *On Certainty*. Oxford: Basil Blackwell.
- Wittgenstein, L. (2009). *Philosophical Investigations*. Oxford: John Wiley & Sons.
- Wright, C. (2014). Comment on Paul Boghossian, “What is inference.” *Philosophical Studies*, 169(1), 27–37.