

# Reasoning and Deducing

MARKOS VALARIS  
UNSW Sydney  
m.valaris@unsw.edu.au

What exactly is reasoning? While debate on this question is ongoing, most philosophers seem to agree on at least the following: reasoning is a *mental process* operating on contents, and which consists in *adopting* or *revising* some of your attitudes in light of others. In this paper, I argue that this characterisation is mistaken: there is no *single* mental phenomenon that satisfies both of these conditions. Instead, I characterise two distinct mental phenomena, which I call ‘deducing’ on the one hand, or ‘reasoning’ and ‘inference’ on the other, to play each of these roles. Recognising this division of labour is essential to developing a better understanding of our rational economy.

## 1. Introduction

What exactly is reasoning? While philosophers continue to debate the answer to this question, most seem confident that we know what the question is *about*. In particular, they seem confident that there is a *single* mental phenomenon that is picked out by the term ‘reasoning’ as intended in the question, and that we know which phenomenon it is. As I will argue in this paper, however, this confidence may be misplaced.

Boghossian gives voice to the consensus view, telling us that the phenomenon he has in mind is:

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

Similarly, Broome (2013, p. 234) writes:

Active reasoning is a particular sort of process by which conscious premise-attitudes cause you to acquire a conclusion-attitude.

While the views Boghossian and Broome go on to develop differ in detail, they agree that reasoning is a mental process that takes you, in a norm-governed way, from some initial set of attitudes to a new one. Analogous claims are made or clearly implied by many others, including Balcerak Jackson and Balcerak Jackson (2013), Bird (2018), Hlobil (2015), McHugh and Way (2018), Quilty-Dunn and Mandelbaum (2017), Staffel (2013) and Wedgwood (2006). In short, philosophers typically take reasoning to satisfy both of the following conditions:

- i.       **Mental Process:** Reasoning is a norm-governed process operating on attitudes or their contents.
- ii.       **Change in View:** Reasoning consists in adopting (or revising) some attitudes—such as beliefs and intentions—in light of others.

My aim in this paper is to argue that this is a mistake. Although each of Mental Process and Change in View describes something real about our mental lives, I do not think that there is a *single* mental phenomenon that satisfies both. We often do engage in norm-governed processes that operate on mental contents; but such processes are not, as such, processes of *attitude adoption or revision*. And while each of our beliefs and intentions clearly depends on what else we believe or intend, this is not usefully thought of as any kind of *process*. As I hope to show, getting clear about this division of labour is essential to understanding our inferential practices.

The structure of this paper is as follows. In §2 I clarify the question at issue, and make an intuitive case for a distinction between two mental phenomena corresponding to the two claims identified above. In the terminology I will adopt, ‘deduction’ refers to a mental process, while (as

we will see) ‘reasoning’ and ‘inferring’ do not.<sup>1</sup> In §3 I develop my account of the activity of deduction. §4 develops some of the core contrasts between deduction and reasoning, while §5 elaborates on that account, and shows how the two interact with each other—as they surely do.

Before going on, I should note that some of the issues I discuss have been noticed by others, though not, so far as I can tell, treated systematically in the way I suggest. Harman (1986) emphasises the distinction between ‘argument’ or ‘proof’ on the one hand, and ‘reasoning’ on the other. But, for Harman, ‘argument’ and ‘proof’ seem to be formal or linguistic structures rather than mental acts. Ryle ([1949] 2002), D. G. Brown (1955), and White (1971) argue that reasoning or inference are not processes, but they argue on the basis of (questionable) linguistic intuitions, and Ryle at least seems to miss the reality of what I call ‘deduction’. Rumfitt (2011, 2015) recognises a distinction between deduction and inference that parallels my own, but he has little to say about inference (I do, however, draw on his account of deduction).

## 2. Reasoning and deducing: a first pass

My aims in this section are, first, to clarify the question this paper seeks to address, and, second, to make an intuitive case for the distinction I aim to draw.

Shortly after the passage cited earlier, Boghossian gives an example of the phenomenon he intends to discuss (similar examples abound in the literature):

Consider the following episode of thought, which I will call (Rain):

On waking up one morning I recall that:

(1) It rained last night.

I combine this with my knowledge that:

(2) If it rained last night, then the streets are wet.

---

<sup>1</sup> My use of these labels is to some extent stipulative, though it partially follows Rumfitt (2015, pp. 34–38): I doubt that the distinction I am drawing is consistently marked in ordinary language.

to conclude:

So,

(3) The streets are wet.

This belief then affects my choice of footwear. (Boghossian 2014, p. 2)

Trying to stay as non-committal about contested details as possible, this ‘episode of thought’—let’s call it an instance of *deliberation*—encompasses the following. Boghossian begins with two pre-existing beliefs, namely that it rained last night, and that if it rained last night then the streets are wet. By considering their contents, he works out that something further follows from them; and, thanks to his sensitivity to this consequence relation, he adopts the further belief that the streets are wet. Boghossian’s deliberation *as a whole*, therefore, seems to illustrate both Mental Process and Change in View. Our question is whether there is a *single* thing that Boghossian does—a single mental occurrence—in virtue of which it does so.<sup>2</sup>

On the face of it, the answer seems to be ‘no’: there is no single thing that Boghossian does that satisfies both Mental Process and Change in View. To see this, start by noting that the first aspect of Boghossian’s deliberation—namely, his working out what follows from claims (1) and (2)—seems to be easily, and frequently, dissociable from any attitudes to propositions. You can take any set of hypotheses, towards which you may have no particular attitude, and work out what (in context) follows from them—this is what I will call ‘deduction’.<sup>3</sup> I might, for example, consider the hypotheses that God is omnipotent, omniscient, and infinitely good and yet evil still

---

<sup>2</sup> Wright (2014, p. 28) also suggests that (what he calls) ‘inference’ is only *part* of Boghossian’s deliberation. But he then turns his attention to other matters, and does not expand on what else is going on.

<sup>3</sup> One might suggest that *thinking about P*, or *considering P*, themselves constitute attitudes towards *P*, so that my claim is (trivially) false. But, clearly, this would be only a terminological victory: to qualify as a ‘reasoned change in view’, a mental act must involve a change in the sorts of attitudes that are constitutive of *having a view* in the first place. This paradigmatically includes attitudes that are *settled*, such as beliefs and intentions, and perhaps also attitudes that *present* their objects as being in a certain way, arguably including (depending on one’s views) desires, perceptions, and emotions.

exists, and see what follows. I may do this while lacking or suspending any attitudes towards the original hypotheses, and without any disposition to adopt any particular attitude towards any consequences I deduce from them. The activity of working out the consequences of a set of claims does not appear to necessarily involve—much less *consist in*—adopting or revising any such attitudes.

One might respond that even in such cases I actually do have an attitude towards my original hypotheses, even if it is just one of *agnosticism* or *suspension of judgment*. Even if this is true, however, it only serves to highlight the fact that what attitudes the agent has towards her hypotheses is *irrelevant* to the process of deduction: a theist, an atheist and an agnostic may deduce exactly the same consequences, and in the same way, from the original hypotheses, even if they take incompatible attitudes towards them. This sits uncomfortably with any view that takes reasoning to be a phenomenon that satisfies both Mental Process and Change in View.

Even more strikingly, we can apply our deductive capacities to only partially interpreted and even schematic statements, towards which it is doubtful that we can have attitudes like belief at all. For example, subjects in psychological studies are typically asked to consider what follows from statements such as the following (concerning a hypothetical two-person game): ‘if Allan is in the game, then Betsy is in’ and ‘If Carla is in the game, then David is out’ (Johnson-Laird 2008, p. 45). Notice, however, that in such examples the reference of all referring expressions is left unspecified. If our doxastic attitudes are, in any sense, aimed at the truth, then it is doubtful that we can have any doxastic attitude at all towards such partially interpreted statements, since they lack determinate truth-conditions (perhaps setting aside formal tautologies and contradictions). Nevertheless, this is evidently no barrier to our deductive capacities, as we can deduce the statement ‘Betsy is in the game’ from the given statements without much difficulty. But, to drive the point home, it is highly implausible that doing so consists in adopting a *belief* with the content that Betsy is in the game, since we have no clue who Betsy is.

And the point can be strengthened further. Even authors such as Wright (2014) who recognise that deduction does not need to proceed from (or result in) belief, suggest that it still requires a belief-like attitude of *acceptance*. But I think even this is mistaken: acceptance is simply the wrong sort of thing to figure in deduction. Consider the following Moorean premise:

It is raining, but I do not accept that it is raining.

I take it you can readily deduce ‘I do not accept that it is raining’ from this premise. It is, however, hard to see how you could *rationally accept* the Moorean premise, even just provisionally or for the sake of the argument. And yet this seems to be no barrier to your ability to use it for the purposes of deduction.<sup>4</sup>

What we have seen so far is evidence that the mental process of deducing the consequences of a set of hypotheses does not involve any belief-like attitude of acceptance towards those hypotheses or their consequences. This, however, is consistent with maintaining that there really is a single mental phenomenon that satisfies both Mental Process and Change in View: perhaps I just did not pick the right process to examine yet. Returning to Boghossian’s example, one might suggest that Boghossian’s *adopting a belief* with the content that the streets are wet satisfies both. But I think we should reject this as well.

The fundamental problem is that it is unclear what the relevant process is supposed to be. Admittedly, since Boghossian starts out lacking the belief that the streets are wet and ends up possessing it, there must be some event or other that constitutes his transition from the former state to the latter. And we can grant that this event may take some time: since belief is a complex dispositional state, extensive changes to your psychology may be required before you can be truly

---

<sup>4</sup> On the basis of Moorean phenomena, Barnett (2008) distinguishes between ‘categorical’ and ‘hypothetical’ acceptance. In these terms, the argument here concerns *categorical* acceptance only. This is no real limitation since, so far as I can tell, hypothetical acceptance in Barnett’s sense is not a belief-like attitude at all: hypothetically accepting something appears to involve nothing more than simply conceiving of a way for things to be that makes it true. This is something we can clearly do with regard to the Moorean premise, but it does not involve *accepting* the premise *as* true, in any ordinary sense.

said to believe that  $P$ , and such changes may be drawn out—especially if  $P$  really matters to you. Even in simple cases, perhaps there is some sort of psychological bottleneck to negotiate, for example in moving contents from working to long-term memory. Either way, it may take a while for the idea that  $P$  to *sink in*, as we say.

Nevertheless, such processes are not what philosophers discuss when discussing ‘reasoning’ or ‘inference’. Consider Williams’s (1973) example of a sailor’s mother who, despite the mounting evidence, appears unable to believe that her son has been lost at sea. This is clearly a case where something needs to happen—some process needs to occur—for this belief to sink in. But, whatever it is that might eventually make it possible for the mother to accept the loss of her son, it seems clear that it is not the sort of thing that philosophers are after when they are looking for the ‘process of reasoning’.

What could reasoning *be*, if not a mental process? Consider uses of these terms in examples like the following: ‘scientists infer facts about climate in the distant past from Antarctic ice cores’, or ‘the detective reasons that Alf committed the murder from the fact that the murder weapon was found in his possession’.<sup>5</sup> In such cases, what is at issue is not any process the agents may have gone through on the way to believing as they do, but rather the *grounds on which* they believe as they do. Thus, if the terms used in this way successfully pick out a mental phenomenon, this phenomenon may satisfy Change in View, but not Mental Process.

White (1971, p. 291) writes that ‘inference is not the passage from  $A$  to  $B$ , but the taking of  $B$  as a result of reflection on  $A$ ’ (also quoted by Rumfitt (2015, p. 35)). White’s expression ‘the taking of  $B$ ’ is obscure, and his own positive view on the nature of inference remains unclear. I suggest, however, that we can capture what seems right about White’s characterisation by saying

---

<sup>5</sup> Brown (1955, pp. 351–52) argues that inferring is a state, on the basis of the (alleged) centrality of such uses. I agree with Brown’s view, but I do not think that the linguistic evidence settles the case: the existence of such usages is only suggestive.

that inferring from  $A$  to  $B$  is *believing*  $B$  as a result of reflection on  $A$ . Clearly, believing something by reflecting on reasons for it is a *state* that an agent may be in, not any kind of process.

But isn't there a tension between saying that inference or reasoning in this sense is a *change* in view, and also that it is a state?<sup>6</sup> Not necessarily: your being in state  $S$  can count as a change for you, if  $S$  is a state you were not in before. More importantly, when we are asking about the rationality or (more broadly) the appropriateness of your change to  $S$  from that previous state, what we are most likely asking concerns the *relative merits* of  $S$  compared to that earlier state, not what happened during any transitional period between the two. Norms of reasoning are norms about what attitudes are appropriate given the evidence, not about what happens in-between.

I hope that what I have said so far suffices to indicate that there is, intuitively, an important distinction to be drawn in the area. But recognising that there is intuitive backing for such a distinction is not the same thing as explaining what the distinction is, or where its significance lies. To do this we need to take a closer look at the two phenomena.

### 3. The process of deduction

Let us start with the activity of deducing. To make things more concrete, consider the example I mentioned earlier. You are given the following premises, regarding a two-person game (Johnson-Laird 2008, p. 45):

If Allan is in the game, Betsy is in.

If Carla is in the game, David is out.

And you are asked to determine whether Betsy is in the game. After some thought, you confirm that this is so. We might reconstruct your thinking as follows:

---

<sup>6</sup> White (1971, p. 290) himself seems to be swayed by this apparent tension, contributing to the difficulty in pinning down his actual view.

You consider the possibility that Betsy is not in the game. From this and the first premise, you conclude that Allan is not in the game either. But since the game requires two people, that would mean that both Carla and David are in; and this possibility is ruled out by the second premise. You conclude that Betsy is in.

I take it that this is a paradigmatic example of deduction. Our question, then, is how to understand the nature of this performance.<sup>7</sup> I think cases like this motivate the following description:

- i. Deduction involves considering *possibilities*.<sup>8</sup> The context of the deduction establishes a set of possibilities—in this case, the space of the six possible pairings of four people in a two-person game. Possibilities in the present sense are ways for things to be that are jointly exhaustive, mutually exclusive, and specific in all relevant respects—though typically not *maximally* specific (unlike standard possible worlds). Deduction, then, involves using information contained in your premises to *eliminate* or *exclude* possibilities. In this particular case, you begin by considering those among the original set of

---

<sup>7</sup> The question itself needs clarification. On one reading, it looks like an empirical question for cognitive psychology, not philosophy. But I think there is a legitimate, and different, way of understanding the question, which is less directly beholden to empirical findings. Consider an analogy: if we ask about the nature of chess, we might be asking either for an account of the psychology of chess-playing, or an account of the goals and rules of chess. The two questions are not unrelated: the goal and rules of chess must be constrained by human psychology—otherwise chess would never have been invented—and any psychological story about how humans actually play chess *is* an account of how humans pursue the goals of chess, in accordance with its rules. Nevertheless, they are distinct questions. Empirical evidence is not irrelevant to the latter question, but neither does it directly settle it.

<sup>8</sup> This claim is congenial to an important research program in the psychology of reasoning, namely the ‘mental models’ approach (Johnson-Laird 1983; Johnson-Laird and Byrne 1991). See also Rumfitt (2015, pp. 38-46).

possibilities that do not involve Betsy. Then, using the information contained in the two premises, you conclude that those possibilities are, in fact, excluded.<sup>9</sup>

- ii. When I speak about possibilities being ‘excluded’ or ‘eliminated’, what I have in mind is simply recognising that certain possibilities are *inconsistent with* your premises (though not necessarily *formally* inconsistent). The ability to ‘just see’ such inconsistencies is, I take it, a basic feature of our capacity for contentful thought. Importantly, the fact that we have this capacity does not mean that *every* case of inconsistency is one we can *just see*: most of us cannot just see that the axioms of naïve set theory are inconsistent, for example.

Working out inconsistency in the general case requires *deduction*, of the sort exemplified in our earlier example: the agent *deduces* that Betsy’s not being in the game is inconsistent with her premises, in part by *just seeing* the inconsistency of Carla and David’s both being in the game with the second premise. Needless to say, giving an account of the distinction between inconsistencies we can ‘just see’ and those we need to work out by deduction is a difficult task, and one that I do not expect has a purely formal solution.

But I think the distinction is real.<sup>10</sup>

- iii. The conclusion of a competent deduction is *guaranteed* to be true whenever the premises are: competent deduction is necessarily truth-preserving.<sup>11</sup> Importantly, the ‘guarantee’ is

---

<sup>9</sup> An anonymous referee asks how this account would handle deductions from inconsistent premises, and specifically whether in such cases *every* conclusion would be warranted for the deducer. I do not intend my account to conflict with a classical conception of validity, so I accept that such deductions may be valid (the phenomenon of ‘explosion’). Since our topic is the activity of deducing rather than validity as such, however, we can draw on other resources to explain what is wrong with them: in particular, deduction warranted merely by explosion arguably involve *missing the point* of the activity of deduction, or perhaps *violating its presuppositions*.

<sup>10</sup> Notice also that ‘just seeing’ that some possibilities are inconsistent with your premises may involve processes of its own, just like regular seeing does. The point is just that it does not involve deduction.

<sup>11</sup> As pointed out earlier, it is a feature of deduction that it can employ partially interpreted and schematic premises; in our present example, for instance, we have been given no clue as to who the four characters are. This creates a difficulty, since it is unclear how to evaluate such sentences. But the solution is simple:

relativised to the contextually determined space of possibilities (which may or may not include the actual world, or even the world as the reasoner believes it to be). For instance, notice that in specifying the space of possibilities for the above example, I quietly omitted the empty possibility that no-one is in the game, as the game is cancelled; in most ordinary contexts this possibility would not be part of the set of possibilities to consider.

- iv. Finally, while the thinking involved in our example might be described as conforming to formal inference rules such as *reductio* and double negation elimination, our reconstruction made no reference to such rules. This may still be compatible with deduction involving rule following, insofar as you can follow rules ‘tacitly’, in some sense. Or, alternatively, perhaps formal rules merely *codify* our deductive practices, rather than guide and justify them. I will not discuss this issue here. Note, however, that even in an example as simple as this one, deduction must draw on our understanding of non-logical concepts, such as the concept of being ‘in’ or ‘out’ of a game; so it is questionable whether logical rules can do all the work.

With this picture on the table, we can revisit the points raised in the last section. As promised, this account of deduction partially vindicates the consensus view of reasoning: deduction really is a process operating on contents. At the same time, however, deduction does not involve taking any particular attitude towards those contents. Thus, deduction satisfies Mental Process but not Change in View.

One important consequence of this is that the distinction between ‘theoretical’ and ‘practical’, which philosophers often draw with regard to what they discuss under the label of ‘reasoning’, has no analogue in the case of deduction. Deduction is simply about figuring out the

---

in such cases, the deduction is meant to be truth-preserving for all contextually appropriate uniform interpretations of the uninterpreted terms.

consequences of some propositions, regardless of whether those propositions are the contents of intentions or beliefs, and regardless of what they are about.<sup>12</sup>

One might wonder whether the process of deduction, as described above, is subject to the same regress worries that have plagued the theory of reasoning. The answer seems to be ‘no’. Consider how regress worries are motivated. Focusing on the theoretical case, reasoning is typically supposed to culminate with the agent’s believing her conclusion. The issue, then, is this: we consider an agent, like the Tortoise in Lewis Carroll’s *What the Tortoise Said to Achilles* (1895), who accepts both a set of premises and that some conclusion follows from her premises, and ask how this is supposed to result in her believing that conclusion. It should be clear, however, that no such problem arises for deduction as understood here. Deduction is *only* about establishing that your conclusion is true in all possibilities in which your premises are; no further step of coming to believe your conclusion is required.<sup>13</sup>

But if the outcome of a deduction is not a belief or other attitude towards its conclusion, then what is the *point* of deduction? This is a good question, a fuller answer to which will have to wait. For now, we can begin by noting that deduction is about figuring out *what follows from what*. Assuming that the picture I sketched above is along the right lines, if  $A$  is deducible from  $\Gamma$ , then—in context—it is impossible for all members of  $\Gamma$  to be true while  $A$  is not. Performing a deduction, therefore, results in *taking it that*  $A$  follows from  $\Gamma$ . While spelling out the details may

---

<sup>12</sup> This claim seems related to Anscombe’s (2006) claim that the ‘form’ of practical inference does not differ from that of theoretical inference. Although Anscombe does not distinguish between deduction and inference, her arguments seem more pertinent to deduction than to inference, as understood here.

<sup>13</sup> The argument in this paragraph addresses specifically Carroll-style regress. One could generate a different kind of regress for deduction by questioning the idea that, for some inconsistencies, we can ‘just see’ that they obtain. If deduction requires recognising inconsistencies, while recognising inconsistencies always requires deduction, regress ensues. However, any account of content at all friendly to the idea that grasping a content involves grasping truth conditions (ways of dividing logical space) would appear to rest on just such an ability. Motivating a regress in this way, therefore, risks a general scepticism about content (this is a lesson of Kripke’s (1982) reading of Wittgenstein).

prove tricky, this relation is, fundamentally, a causal one: your taking it that  $A$  follows from  $I$  is a causal outcome of your considering what possibilities are or are not consistent with  $I$ .<sup>14</sup> (Saying this is, of course, compatible with scepticism about the prospects of reductive causal analyses of deduction.)

How should we think of the attitude of taking  $A$  to follow from  $I$ ? In informal terms, we can think of it as a doxastic attitude with conditional force, expressible as ‘ $A$ , given  $I$ ’. But there are difficulties, of a familiar sort, if we try to be more precise. On the one hand, taking it that  $A$  has to be true given  $I$  at least *seems* to have substantive cognitive content. It rules out certain ways for things to be that, antecedently, appeared to you to be open—namely, ways such that all members of  $I$  are true and  $A$  is not. One can certainly be wrong about this sort of thing, for example. On the other hand, however, if your deduction was correct, then the ways for things to be that you end up excluding were never, in fact, genuinely possible. They only *appeared* to you to be possible. So, if we wanted to represent content by the set of possibilities that it rules out, we would have to conclude that every correct deduction has the same content as any other, namely the empty set.<sup>15</sup>

I cannot pursue this issue at length here. Nevertheless, for the sake of the discussion below, let me briefly outline a way to think about it. What the difficulty shows is that the set of

---

<sup>14</sup> The details may prove tricky because it may be unclear where the process ends and the resultant state begins. But the issues here have to do with the beginnings and ends of processes in general, and have nothing in particular to do with deduction or reasoning. (For discussion of these puzzles, see Medlin (1963).) For the purposes of this paper, I set these problems aside; I will simply take it for granted that we can meaningfully speak of a process culminating with, or resulting in, a state.

<sup>15</sup> As an anonymous referee points out, this problem is closely connected to Mill’s famous complaint that logicians fail to explain how the ‘syllogistic process’ can constitute ‘a progress from the known to the unknown’ (Mill 1879, p. 188). Mill’s solution depends on his view of deduction as reducible to syllogistic logic, and so is not acceptable. As I explain below, to find a substantive cognitive role for deduction we need to take into account the *limits* of our rational capacities: it is only because we are limited cognitive agents, who do not always ‘just see’ what is genuinely possible and what not, that deduction is indispensable to us.

‘possibilities’ on which deduction operates should be taken to include not just ways for things to be that are genuinely possible, but also those that merely *appear* to you to be possible, pre-reflectively. For example, since (assuming you are like me) simply reading the constraints listed above for the two-person game does not automatically settle it for you that Betsy must be in the game, accounting for your subsequent deduction requires including among the ways for things to be that you consider some in which Betsy is *not* in the game. Your deduction, then, consists in your working out that such ways for things to be are not genuinely possible.<sup>16</sup>

Now, if the conditional attitudes in which deduction terminates have substantive cognitive content, questions of *justification* may arise for them. Under what conditions is an agent justified in taking it that  $A$  follows from  $I$ , as a result of a deduction?

Different views are possible here. On one kind of view, the justification of such attitudes depends on their origins: for example, conditional beliefs may be justified if they result from dispositions to deduce well, and unjustified otherwise (e.g., Leitgeb (2007)). Alternatively, perhaps the relevant conditional doxastic attitudes are justified (when they are) not in virtue of their origins, but in virtue either of some intrinsic features of the contents of those attitudes themselves, or of the overall coherence of an agent’s entire belief-system.<sup>17</sup> I will not attempt to choose between these views. What I want to do, instead, is focus on some general norms that seem to govern the *process of deduction* itself, rather than the norms that apply to the states it culminates in.

---

<sup>16</sup> For similar remarks see Rumfitt (2008, pp. 69-70). Jago (2012, 2014) argues that there may be no determinate answer as to *which* impossible ways for things to be an agent fails to dismiss pre-reflectively; as a result, he develops an account along the same basic lines, but according to which the content of deduction is vague. This seems compatible with what I say in this paper.

<sup>17</sup> For the former sort of view, see Wedgwood (2012) who argues that some conditional beliefs are rational because they encode arguments constitutive of some of our concepts. Harman (1986) suggests a view of the latter kind.

Of particular interest here are two such norms, namely, *monotonicity* and *transitivity*. Both feature in ‘classical’ views of logical consequence, and both have been questioned in the theory of reasoning. It is instructive to see how the present account motivates those principles as principles of deduction.

Take monotonicity first. The idea is that if you can deduce  $A$  from premises  $\Gamma$ , then you can deduce  $A$  from any expanded list of premises that includes  $\Gamma$ . This is because deduction proceeds by excluding possibilities. By hypothesis,  $A$  is true in all possibilities not ruled out by  $\Gamma$ ; and adding more premises to  $\Gamma$  is not going to get you any *new* possibilities, only exclude some of the remaining ones. So, there is no way to undermine the deduction by adding new premises.

Similar points motivate transitivity. The idea here is that deductions can be chained together: a long deduction is a sequence of shorter ones. For present purposes, we can formulate transitivity as follows: if  $A$  can be deduced from  $\Gamma$ , and  $B$  can be deduced from  $A$ , then  $B$  can be deduced from  $\Gamma$ . As with monotonicity, this is very plausible: there is no way to undermine a deduction by performing more deductions, since deductions *exclude* possibilities without introducing new ones.

Thus, my account of deduction suggests that the structural principles of classical logical consequence hold for deduction. But what about the purported counterexamples?

Consider standard counterexamples to monotonicity. If all you know about Tweety is that he is a bird, you may conclude that Tweety can fly. But given the further information that Tweety is, in fact, a penguin, this conclusion ceases to be rational. In the jargon, the claim that Tweety is a penguin is a ‘rebutting’ defeater of your reasoning. In a different sort of case, suppose you have just proved a difficult theorem in logic. You are reasonably confident in your abilities as a theorem-prover, so that you now believe some complicated formula to be logically true. But then, you are told that someone earlier slipped some drug into your coffee, whose effects include undermining the deductive capacities of those who ingest it while also making them overconfident (Christensen 2007). If you are rational, the story goes, you should now

withdraw your belief in your theorem. The claim that you ingested the drug is an ‘undercutting’ defeater.

However, such cases are not counterexamples to the thesis that *deduction* satisfies monotonicity, *because they are not cases of deduction at all*. Notice that, in order to get the right intuitions about the examples, we need to talk about what is and is not rational to believe. Otherwise, the move from ‘Tweety is a bird’ to ‘Tweety can fly’ would never have been sanctioned in the first place: obviously, that Tweety is a bird does not eliminate all ways for things to be such that Tweety cannot fly. Even more obviously, the undercutting defeater in the second example does not target your conclusion at all, but only your warrant or justification for believing it. This, from the present point of view, is a clear mark that we are not talking about deduction, but about something else: deduction is *indifferent* to any attitude you may or may not have to your conclusion. Similar points apply to purported counterexamples to transitivity as well.<sup>18</sup>

In fact, we can take this as a further reason for drawing the distinction I have been urging. There seems to be *one* reading of terms such as ‘reasoning’ or ‘deducing’ according to which they *are* governed by norms such as monotonicity and transitivity—indeed, these norms can seem inescapable, for an intuitive understanding of what it is to derive one thing from another. But there also seems to be another reading of these same terms, on which they are not similarly constrained: our inferential practices, understood as reasoned changes in view, are notoriously defeasible and holistic. As a result, they cannot be expected to satisfy strong structural principles such as monotonicity and transitivity. As we will see, this intuition is borne out by the account of reasoning sketched in what follows.

---

<sup>18</sup> In addition to worries having to do with defeasibility, there are also objections to transitivity from sorites-type considerations (Rumfitt 2015, pp. 45-46). These are not directly relevant to present concerns.

#### 4. Is reasoning a process?

I have been arguing that deduction ought to be distinguished from reasoning or inference, understood as reasoned change in view. This is the sense in which, in his (Rain) example, Boghossian (2014, 2) reasons that the streets are wet from the premise that it rained last night, or the detective reasons that Alf committed the murder from the fact that a bloody knife with the victim's DNA was found in his possession. So, what should we say about reasoning in this sense?

One difficulty is that, while deduction does not presuppose any particular attitude towards its premises and conclusions, reasoning does: it matters whether you are reasoning to a belief or an intention, for example. Thus, we may not be able to say very much about what reasoning is while abstracting from the nature of the attitude types involved. Nevertheless, we can say enough to highlight the following point: reasoning, in the sense of interest here, is not a mental process.

Reasoning is the sort of thing that is subject to norms of rationality. So, if reasoning were a process extended in time, norms of rationality would exist that govern this process. What would such norms look like? Plausibly, such norms would constrain your attitudes at the end of the process, given your attitudes at its beginning. Thus, they would have the following form (focussing on beliefs again):

If you believe  $A_1, A_2 \dots A_n$  at  $t_1$ , then you are rationally required/permitted to believe  $B$  at a later time  $t_2$ .

But I think it is implausible that there are diachronic norms of this form.<sup>19</sup>

---

<sup>19</sup> The claim is not that there are *no* diachronic norms for belief at all. For example, it may be rational to *maintain* a belief for which you have no present evidence, so long as you have no contrary evidence either. I take no stand on norms of this sort. Thus, my argument falls short of supporting 'time-slice rationality'—a view defended by Hedden (2015) and Moss (2015).

The reason is straightforward: even if my beliefs at  $t_1$  make belief in  $B$  rational, there is no guarantee that my beliefs at  $t_2$  still do so: I may lose some of the earlier beliefs that were essential to  $B$ 's rationality, or acquire new ones that defeat it. Moreover, all this can happen without any lapse in rationality. As a result, however, it may be neither required nor permitted of me to believe  $B$  when  $t_2$  rolls around.<sup>20</sup> For example, even if at  $t_1$  I believe both  $A$  and that if  $A$ , then  $B$ , by  $t_2$  I may have rationally changed my mind, and now believe not- $A$ . Assuming I have no independent grounds for  $B$ , it would seem that rationality does not permit me to believe  $B$  *now*.

Why should we think that there *must* be norms of rationality of this sort? Broome explains his commitment to them by asserting that he is concerned 'with causal processes of reasoning, which take time' (Broome 2013, p. 187). This does not help, since our question is precisely whether reasoning is a process. Similarly, Podgorski (2016) presupposes rather than argues that the evaluation of reasoning is about belief-forming processes. What we need is an argument that there can be pairs of agents who are equivalent in *synchronic* rationality but not in *overall* rationality, and where the difference is best explained by a difference in a process of reasoning. Broome (2015) and Hlobil (2015) purport to describe such cases. Here is Hlobil's example:

Suppose I am doxastically justified in believing at  $t_0$  that *not-A*, that *if B, then A*, and that *if not-B, then not-A*. At time  $t_1$ , I have added the belief that *not-B* to my stock of beliefs, and there is no time between  $t_0$  and  $t_1$  at which I am in a third belief state. Now consider two ways in which this might happen. In the first case, I fallaciously infer *not-B* from *not-A* and *if not-B, then not-A* by affirming the consequent. In the second case, I competently infer *not-B* by modus tollens [...]. In the first case I am less epistemically rational than in the second case. My belief states at  $t_0$  and at  $t_1$  are identical. Hence, I must have violated a

---

<sup>20</sup> Essentially the same argument is given by Boyle (2011, pp. 11–12) and Broome (2007, p. 368). Broome seems to have changed his mind since, as he has more recently defended such norms (Broome 2013, pp. 177–91, 2015).

norm of diachronic rationality in the first case but not in the second. Therefore, there are diachronic norms of epistemic rationality. (Hlobil 2015, pp. 38-39)

Hlobil argues that, since the agent's belief states are identical at the relevant times, the only way the difference in rationality between them could be explained is by a *diachronic* norm: a norm that governs the transition between the two states, and which the agent violates in the first case but not in the second.

The argument, however, is inconclusive, because Hlobil does not do enough to show that the agent's belief states really are the same in the two cases, in all relevant respects. If, for example, we include information about *why* the agent believes what she does in each case, then the agent's states at  $t_1$  differ after all: the agent in the first case believes *not-B* because she believes *not-A* and *if not-B, then not-A*, while in the second because she believes *not-A* and *if B, then A*.<sup>21</sup> In addition, if reasoning is subject to what Boghossian (2014) calls the 'Taking Condition' (the condition that in reasoning an agent must take her premises to support her conclusion), the agent in the first case takes it that *not-B* follows from *not-A* and *if not-B, then not-A*, while in the second case she takes it that it follows from *not-A* and *if B, then A* (see Valaris 2016b). If either of these richer descriptions is correct, there are synchronic differences between the two cases.

Still, one might insist that there simply is an intuition that, in addition to any difference between the agent's doxastic *states*, there is also something *processual* that the agent does correctly in the second case but not the first. But we can accommodate this without accepting that

---

<sup>21</sup> Hlobil (2015, p. 40) considers the possibility of accounting for the difference in rationality by means of synchronic 'basing' constraints, but suggests that fallacious inferences involve additional violations of rationality, over and above violations of such constraints. So far as I can see, Hlobil gives no argument for this. In any case, as we will see we can accommodate something like this idea, by allowing that the agent in Hlobil's first scenario *deduces* fallaciously, as well as violating synchronic norms of rationality. This still falls short of admitting the existence of diachronic norms of rationality, however, because the irrationality of the agent's ensuing doxastic state may still be fully explained by her violation of the synchronic basing constraint, rather than her fallacious deduction.

reasoning is a process. In order to believe that *not-B* on the basis of *not-A* and *if B then A*, the agent in the second case had to *recognise* that *not-B* follows from these two other claims. To do so, she presumably performed an elementary *deduction*.<sup>22</sup> The same goes, of course, for the first case. And it is clear that the two deductions are not equally correct: one of them conforms to the valid pattern of modus tollens, while the other is an instance of affirming the consequent. So, there is, indeed, a processual difference between the two cases; but the difference is in *deduction*, not reasoning.

This discussion highlights a crucial point which, I think, goes a long way towards explaining why the distinction I am developing has proven so easy to miss. Although reasoning *as such* is not a process, instances of deliberation in which reasoning features also commonly feature processes of deduction. In the next section, we will get a better view of how this works, at least for the specific case of deliberating with beliefs.

## 5. Belief and deliberation

My focus in this section will be on reasoning with *full* or *outright* beliefs.<sup>23</sup> As I understand this state, believing that *A* essentially involves (or perhaps even is) being disposed to disregard possibilities inconsistent with *A* for the purposes of deliberation and action. If I believe that my flight leaves at 4 pm then I will plan accordingly, disregarding alternative possibilities such as that my flight leaves at 4 am, or that it has been cancelled.

---

<sup>22</sup> It may seem odd to speak of a *process* of deduction here, as for many of us this would be the sort of entailment we can ‘just see’. This seems fair, but what I am trying to do here is simply accommodate Hlobil’s intuition that his example involves some kind of process of working out what follows from your existing beliefs. If that is denied, then Hlobil’s case collapses anyway.

<sup>23</sup> There are philosophers who argue that the true principles of rationality do not concern full or outright beliefs, but only more fine-grained degrees of belief or credences, understood as subjective probabilities (Christensen 2004). I cannot engage this debate here. For reasons to take belief to be a cognitive state in good standing, see Buchak (2014).

What is it to disregard a possibility, in the relevant sense? I am not sure I can give a non-question-begging characterisation, but the idea is intuitive. It is closely connected to the notions of *expectation* and *surprise*. If at the end of the day I believe that my bicycle is where I left it in the morning, I *expect* to find it there, and I will be *surprised* if it turns out that it is not there.

Importantly, disregarding a possibility does not entail that you no longer regard it as a possibility at all: even if I disregard the possibility that my flight has been cancelled, I do not take it to be *impossible* for my flight to have been cancelled. As we know, the world may contradict even our most cherished beliefs. Accordingly, the notion of belief here described is naturally modelled by imposing a ‘plausibility’ ordering on epistemic possibilities, with belief requiring truth only in the *most plausible* epistemic possibilities, rather than all the accessible ones. This allows that the things you disbelieve are not seen as impossible, only more or less ‘implausible’ (Spohn (2012) examines such models, and their sophisticated ‘ranking’ relatives, in great detail).

Thus, a substantive general theory of rationality for beliefs would consist in large part in an account of when it is rational to disregard possibilities. Perhaps we have some non-statistical sense of what *normally* happens, and this is what guides us as to which possibilities we may rationally disregard and which we may not. On such a view, what we may rationally believe is what is maximally *normal*, given our evidence (Smith 2010, 2016; Valaris 2017). Or, perhaps, which possibilities we may disregard partly depends on pragmatic factors: under some circumstances I can rationally disregard the possibility that my flight has been cancelled, but if the stakes are high enough I may have to take it seriously. For present purposes, we do not need a general account of rational belief. We can focus on the narrower question of when it is rational to believe *B* on the strength of *A*, or (equivalently) when it is rational to *infer B* from *A*.

Suppose you start out with a stock of beliefs, which, as above, is represented by an ordering of possibilities with regard to plausibility. The things that you believe correspond to the possibilities that are ‘live’ for you, or in other words regarded as maximally plausible. Our question, then, is what you can infer once you get to accept a possibly new premise, the

(contingent) proposition  $A$ .<sup>24</sup> From the present point of view, the answer is straightforward, at least in outline. Given that you now believe  $A$ , you are committed to disregarding all possibilities such that not- $A$ . It is then rational for you to believe  $B$ , just in case doing so would not commit you to disregarding any *further* possibilities—just in case, that is,  $B$  is true in all the most plausible  $A$  possibilities.

Notice that it can be rational to infer  $B$  from  $A$ , even if you do not take it to be impossible for  $B$  to be false while  $A$  is true: what matters is just that such possibilities are not live for you, or are rationally disregarded. One consequence of this is that rational inference is neither monotonic nor transitive. The fact that  $B$  is true in all the live possibilities given only that  $A$  does not mean that it will remain true in all the live possibilities given that  $A \& D$ . For example, given that all you believe about Tweety is that he is a bird, perhaps in all the most plausible possibilities consistent with what you believe it is true that Tweety can fly. But the most plausible possibilities in which Tweety is not just a bird, but specifically a penguin, are not such that Tweety can fly. So, while you could reason from the premise that Tweety is a bird to the conclusion that Tweety can fly, you cannot reason from the premises that Tweety is a bird and that Tweety is a penguin to the conclusion that Tweety can fly. Similar examples show that transitivity fails as well (see Spohn (2012, pp. 112–13) for discussion).<sup>25</sup>

---

<sup>24</sup> For full generality, we would need to allow that premises may be accepted more or less ‘firmly’. This creates complexities I cannot address here. For some discussion, see Spohn (2012, pp. 82-83).

<sup>25</sup> The account just outlined follows views generally discussed under the heading of ‘belief dynamics’ (see, e.g., Gärdenfors 1990; Spohn 2012; van Benthem 2011). One might initially be surprised by this, as the label ‘belief dynamics’ may suggest that what is at issue in such accounts is the rationality of *processes* of reasoning. This appearance, however, is misleading: accounts of belief dynamics are not really about any processes. The question they address is what *doxastic state* would be rational of you to adopt, given your evidence and your existing beliefs. They discuss neither the sources of your evidence, nor any process that might be involved in working out what your evidence supports. My discussion here goes further than these, in seeking to address the latter of these two topics: as we will see, the role of deduction in deliberation, on my view, is precisely to determine what your evidence requires of you, given what else you believe.

With this picture of reasoning in place, we can now turn to the question of how reasoning and deduction relate to each other. On the account sketched so far, an agent can reason to  $B$  from  $A$  just in case believing  $B$  does not require her to disregard any *new* possibilities, over and above those she disregards in believing  $A$ . This raises the question, however, of how the agent can decide whether this is the case. This is a question that accounts of belief dynamics in a formal style generally avoid, because they incorporate the idealisation that a rational agent's doxastic states are automatically closed under deductive consequence. But, as noted already in §3, to see the point of deduction we need to take into account the fact that we cannot always 'just see' relations of inconsistency or entailment among ways for things to be. Linking deduction and reasoning in an account of deliberation, therefore, requires relaxing the idealisation of deductive closure as well.

Let us begin by considering, as before, a case where you reason to  $B$  from a premise  $A$ . In our framework, believing  $A$  implies that all ways for things to be that are currently live for you are such that  $A$  is true. Let us further assume that  $A$  entails  $B$ , so that there are no real possibilities such that  $A$  is true but  $B$  fails to be true. The crucial point, now, is that it does *not* follow from this that all ways for things to be that are live for you are, automatically, such that  $B$ : given your limited rationality, ways for things to be in which  $A$  is true but  $B$  is not may still *appear* as possibilities (and, indeed, as *live* possibilities) for you. (Such ways for things to be may be simply silent with regard to  $B$ , or even such that not- $B$  is true.) This means, of course, that at this point you believe  $A$  but do not believe  $B$ . *A fortiori*, you do not yet infer or reason to  $B$  from  $A$ .

Clearly, what needs to happen for you to reason to  $B$  from  $A$  is for you to recognise that ways for things to be such that  $A$  is true but  $B$  is not are not, in fact, possible (as we have assumed,  $A$  entails  $B$ ). And this, in my view, is where deduction comes in. It is as a result of deducing  $B$  from  $A$  that you get to exclude such ways for things to be from consideration. And, of course, assuming that at the end of this process you still believe  $A$ , this means that no ways

for things to be in which  $B$  fails to be true remain live for you. You now believe  $B$ , by reasoning from  $A$ .

A couple of points are important to clarify here. To begin with, one may wonder whether inference or reasoning in this sense must *always* result from deduction. The answer is clearly ‘no’. What matters, on the present account, is simply that you believe your conclusion by *taking it that* it follows from your premises, in the sense that you eliminate ways for things to be such that your premises are true but your conclusion is not. This can be the result of deduction, or it can be something else—for instance, a case of ‘just seeing’ that one thing follows from another.<sup>26</sup>

Relatedly, it should also be clear that the present account has no trouble accommodating cases where you take it that your conclusion follows from your premises when in fact it does not. In this case, among the ways for things to be such that  $A$  is true and  $B$  is not that you eliminate will be some that are real possibilities, and which therefore you ought not have eliminated. This can be the result of a faulty deduction, or perhaps some other mistake (for example, it may seem to you as if you can *just see* inconsistencies that are not there).

Furthermore, while my discussion so far has focused on cases where your premises entail your conclusion (that is, deductively valid inferences), it is important to see that the present framework is not restricted to such cases. For example, seeing that Raji is about to strike a match, you infer that the match will light. The inference here is not deductively valid: it is not inconsistent with the match’s being struck that it will not light. How, then, are you able to infer that the match will light, on the present account?

---

<sup>26</sup> Is there a limit to the sorts of things that may result in your taking it that  $B$  follows from  $A$ , in the relevant sense? Could a bump on the head result in your taking it that  $B$  follows from  $A$ , and hence in your inferring  $B$  from  $A$ ? I doubt that we should impose such a limit. If a bump on the head causes you to suddenly become aware of entailment relations that you did not see before, then it will have expanded the realm of inferences available to you. Depending on how the details are fleshed out, it may be *irrational* for you to trust these new inferences, but this does not show that they are not inferences at all.

To begin with, as we already saw, inferences of this sort can be rational, because although your conclusion is not true in *all* possibilities in which your premises are, it may be true in all possibilities in which your premises are *and which are currently live* for you. Thus, for example, we may suppose that you are not currently exercised by the possibility that the match is wet, or that a very powerful magnetic field is present, and so on. This is why it is still the case that *deduction* can enable you to infer that the match will light. Within the restricted set of possibilities that are live for you, the match's not lighting *is* inconsistent with its being struck. Recognising this inconsistency is enough, therefore, for you to infer that the match will light. This is so despite the existence of genuine possibilities in which the match is struck but does not light.

But doesn't this simply amount to adding further premises—for instance, that the match is not wet, no powerful magnetic fields are present, and so on—so as to turn a non-deductive inference into a deductively valid one? It does not. This objection misses the distinction between the *premises* of an inference and its *background assumptions* or *presuppositions*. This is a familiar distinction, and one I believe is worth preserving: our everyday epistemic practices, for example, seem to presuppose that we are not disembodied brains in vats, but it is far from obvious that this should be included as a premise in virtually all of our inferences. This is not just a claim about psychological plausibility; it is also, and primarily, a point about the different epistemic roles that premises and presuppositions would seem to play in our rational economy (see, for example, Pryor (2004), Wright (2000, 2002)). With this distinction in mind, our earlier example would have you infer that the match will light from the premise that Raji is about to strike it, while assuming in the background that the match is not wet, no powerful magnetic fields are present, and so on.

In closing, let us consider what lessons about the structure of deliberation we can draw from the above. Crucially, the present account of deliberation requires no reference to a process of inferring or reasoning. The only process it appeals to is the process of deduction. As argued in §3 of this paper, the upshot of a deduction from  $\Gamma$  to  $\mathcal{A}$  is a doxastic attitude with content

expressible as ‘ $A$ , given  $I$ ’. The present point is that, in a context where we are assuming you believe  $I$ , this attitude can just *be* your believing  $A$  on the grounds that  $I$  (for similar points see Gibbons (2009), Neta (2013), Valaris (2014)). Once the process of deduction has done its job, in other words, there is no further process (of a relevant sort<sup>27</sup>) that is needed for you to carry out the inference. In the right circumstances, your taking it that your conclusion follows from your premises just *is* your inferring your conclusion from your premises.<sup>28</sup> On the present account, as White (1971, p. 291) puts it, ‘inference is not the passage from  $A$  to  $B$ , but the taking of  $B$  as a result of reflection on  $A$ ’.<sup>29</sup>

## References

- Anscombe, G. E. M. 2006, ‘Practical Inference’, In *Human Life, Action and Ethics: Essays by G.E.M. Anscombe* (Charlottesville, VA: Imprint Academic)
- Barnett, David 2008 ‘Ramsey + Moore  $\neq$  God’, *Analysis* 68: 168–74.
- Benthem, Johan van 2011, *Logical Dynamics of Information and Interaction* (Cambridge: Cambridge University Press)
- Bird, Alexander 2018, ‘Evidence and Inference’, *Philosophy and Phenomenological Research* 96: 299–317
- Boghossian, Paul 2014, ‘What Is Inference?’, *Philosophical Studies* 169: 1–18
- Boyle, Matthew 2011, “‘Making Up Your Mind’ and the Activity of Reason”, *Philosophers’ Imprint* 11: 1–24
- Broome, John 2007, ‘Wide or Narrow Scope?’, *Mind* 116: 359–70

---

<sup>27</sup> This qualification is meant to allow for ‘sinking in’ processes, of the sort discussed in §2.

<sup>28</sup> This, incidentally, shows how the present account can incorporate the ‘taking condition’ on inference (the condition that to infer  $B$  from  $A$  you need to take it that  $A$  supports  $B$ ), without threat of regress. For a similar response, see Valaris (2014).

<sup>29</sup> I would like to thank the Editors of *Mind*, as well as three anonymous referees, for their patience and valuable comments. An earlier draft of this paper was presented at a Workshop on Reasons and Reasoning at the Australian National University, organised by Nic Southwood. I also had the opportunity to present ideas from this paper at a Workshop on Mental Agency in London, organised by Michael Brent and Matthew Soteriou. I greatly benefited from the comments I received on both of these occasions. Finally, I would like to thank Melissa Merritt, who read and commented on several drafts of this paper.

- 2013, *Rationality through Reasoning* (Chichester: Wiley Blackwell)
- 2015, ‘Synchronic Requirements and Diachronic Permissions’, *Canadian Journal of Philosophy* 45: 630–46
- Brown, D. G. 1955, ‘The Nature of Inference’, *The Philosophical Review* 64: 351–69
- Buchak, Lara 2014, ‘Belief, Credence, and Norms’, *Philosophical Studies* 169: 285–311
- Carroll, Lewis 1895, ‘What the Tortoise Said to Achilles’, *Mind* 4: 278
- Christensen, David 2004, *Putting Logic in Its Place: Formal Constraints on Rational Belief*, (New York: Oxford University Press)
- 2007, ‘Does Murphy’s Law Apply in Epistemology? Self-Doubt and Rational Ideals’, *Oxford Studies in Epistemology* 2: 3–31
- Gärdenfors, Peter 1990, *Knowledge in Flux: Modeling the Dynamics of Epistemic States* (Cambridge, MA: MIT Press)
- Gibbons, John 2009, ‘Reason in Action’, In Lucy O’Brien and Matthew Soteriou (eds.), *Mental Actions* (Oxford: Oxford University Press)
- Harman, Gilbert 1986 *Change in View* (Cambridge, MA: MIT Press)
- Hedden, Brian 2015, *Reasons Without Persons: Rationality, Identity, and Time* (Oxford: Oxford University Press)
- Hlobil, Ulf 2015, ‘There Are Diachronic Norms of Rationality’, *Thought: A Journal of Philosophy* 4: 38–45
- Jago, Mark 2012, ‘The Content of Deduction’, *Journal of Philosophical Logic* 42: 317–334
- 2014, *The Impossible: An Essay on Hyperintensionality*, (Oxford: Oxford University Press)
- Johnson-Laird, Philip 1983, *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness* (Cambridge, MA: Harvard University Press)
- . 2008, *How We Reason*, (Oxford: Oxford University Press)
- Johnson-Laird, Philip, and Ruth Byrne 1991, *Deduction*, (Hillsdale: Psychology Press)
- Kripke, Saul 1982, *Wittgenstein on Rules and Private Language: An Elementary Exposition* (Cambridge, MA: Harvard University Press)
- Leitgeb, Hannes 2007, ‘Beliefs in Conditionals vs. Conditional Beliefs’, *Topoi* 26: 115–32
- McHugh, Conor, and Jonathan Way 2018, ‘What Is Reasoning?’, *Mind* 127: 167–196
- Medlin, Brian 1963, ‘The Origin of Motion’, *Mind* 72: 155–75
- Mill, John Stuart 1879, *A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence, and the Methods of Scientific Investigation*, (London: Longmans Green and Co.)
- Moss, Sarah 2015, ‘Time-Slice Epistemology and Action Under Indeterminacy’, *Oxford Studies in Epistemology* 5: 172–94
- Neta, Ram 2013, ‘What Is an Inference?’, *Philosophical Issues* 23: 388–407
- Podgorski, Abelard 2016, ‘A Reply to the Synchronist.’ *Mind* 125: 859–71
- Pryor, James 2004, ‘What’s Wrong with Moore’s Argument?’, *Philosophical Issues* 14: 349–78
- Quilty-Dunn, Jake, and Eric Mandelbaum 2017, ‘Inferential Transitions’, *Australasian Journal of Philosophy* online first, < <https://doi-org/10.1080/00048402.2017.1358754>>
- Rumfitt, Ian 2011, ‘Inference, Deduction, Logic’, in John Bengson and Marc A. Moffett (eds.), *Knowing How: Essays on Knowledge, Mind, and Action* (Oxford: Oxford University Press)
- 2015 *The Boundary Stones of Thought: An Essay in the Philosophy of Logic* (New York: Oxford University Press)
- Ryle, Gilbert (1949) 2002, *The Concept of Mind* (Chicago: University of Chicago Press)
- Smith, Martin 2010, ‘What Else Justification Could Be’, *Noûs* 44: 10–31
- 2016, *Between Probability and Certainty: What Justifies Belief*, (Oxford: Oxford University Press)
- Spohn, Wolfgang 2012, *The Laws of Belief: Ranking Theory and Its Philosophical Applications* (Oxford: Oxford University Press)
- Staffel, Julia 2013, ‘Can There Be Reasoning with Degrees of Belief?’, *Synthese* 190: 3535–3551.

- Valaris, Markos 2014, 'Reasoning and Regress', *Mind* 123: 101–127
- 2017, 'Induction, Normality and Reasoning with Arbitrary Objects', *Ratio* 30: 137–48
- Wedgwood, Ralph 2006, 'The Normative Force of Reasoning', *Noûs* 40: 660–86
- 2012, 'Justified Inference', *Synthese* 189: 273–95
- White, Alan 1971, 'Inference', *The Philosophical Quarterly* 21: 289–302
- Williams, Bernard 1973, 'Deciding to Believe', in *Problems of the Self* (Cambridge: Cambridge University Press)
- Wright, Crispin 2000, 'Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof', *Philosophical Issues* 10: 140–63
- 2002, '(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell', *Philosophy and Phenomenological Research* LXV: 330–48
- 2014, 'Comment on Paul Boghossian, "What Is Inference"', *Philosophical Studies* 169: 27–37