

Inference and Self-Knowledge

Abstract: A growing cohort of philosophers argue that inference, understood as an agent-level psychological process or event, is subject to a “Taking Condition” (Boghossian 2014). The Taking Condition states, roughly, that drawing an inference requires one to take one’s premise(s) to epistemically support one’s conclusion, where “takings” are some sort of higher-order attitude, thought, intuition, or act. My question is not about the nature of takings, but about their contents. I examine the prospects for “minimal” and “robust” views of the contents of takings. On the minimal view, taking one’s premise(s) to support one’s conclusion only requires focusing on propositional contents and putative epistemic support relations between them (cf. Peacocke 1996). On the robust view, taking one’s premise(s) to support one’s conclusion also requires knowledge (or being in a position to have knowledge) of the attitudes one holds toward those contents. I argue that arguments for the Taking Condition do not entail or sufficiently motivate the robust view. Accordingly, contra several philosophers, the Taking Condition does not illuminate a deep relationship between inference and self-knowledge.

Keywords: Inference; Taking Condition; Self-Knowledge; Second-Tier Thought; Cognitive Agency

§1—Introduction

Plausibly, drawing an inference can be an agential phenomenon: it can be a psychological process or event that is predicable of the agent herself rather than of her sub-agential cognitive mechanisms.¹ But what makes inference an agential phenomenon when it is one? A growing cohort of philosophers argues that inference involves agency because drawing an inference requires the agent to (1) “take” on how her premise(s) to confer epistemic support on her conclusion, and (2) to draw her conclusion on the basis of this take.

Paul Boghossian offers the following formulation of this suggestion:

(Taking Condition): Inferring from p to q necessarily involves the thinker *taking* p to support q and drawing q *because* of that fact. (2019, p. 110)²

¹ Philosophers often distinguish agent-level and sub-agential inference via the distinction between system 2 and system 1 processing. See Keith Stanovich and Richard West “Individual Differences in Reasoning: Implications for the Rationality Debate,” *Behavioral and Brain Sciences* 23, 5 (2000): 645-665; See also Daniel Kahneman, *Thinking Fast and Slow* (Macmillan Press, 2011). In this paper all talk of inference/reasoning refers roughly to system 2 (agent-level) inference/reasoning. I say ‘roughly’ because these systems may have blurry lines. See, for this reason, Paul Boghossian’s talk of “system 1.5 and up” inference/reasoning in his “What is Inference?” *Philosophical Studies*, 169 (2014): 1-18.

² See Paul Boghossian, “Inference, Agency, and Responsibility” in *Reasoning: New Essays on Theoretical and Practical Thinking*, eds. Brendan Balcerak Jackson and Magdalena Balcerak Jackson (Oxford University Press,

If the Taking Condition (TC) is true, then agent-level inference (hereafter just *inference* or *reasoning*³) is not a wholly first-order process. Rather, the agent must have an intermediating conception of the quality of epistemic support between her premise(s) and conclusion.

Besides arguing for TC itself, its proponents must also clarify exactly what taking one's premise(s) to support one's conclusion amounts to. On this score, philosophers have variously argued that "takings" are beliefs,⁴ intuitions,⁵ mental actions,⁶ and *sui generis* mental states.⁷ I will not be adding to this particular debate here, nor will I be questioning whether TC is true.⁸ Instead, I will ask a different question about TC, one that concerns the *contents* of takings. Specifically, I will ask: does an agent's taking her premise(s) to support her conclusion require that she have self-knowledge of the attitudes she bears toward her premise(s) and conclusion? Boghossian appears to think that it does. Focusing on theoretical inferences, he espouses:

2019), 101-124. He adapts this proposal from Gottlob Frege's *Posthumous Writings*, edited by Hermes Hans, Friedrich Kambartel, and Friedrich Kaulbach (Oxford: Basil Blackwell 1979). In "Inferring By Attaching Force," *Australasian Journal of Philosophy*, 97, 4 (2019): 701-714, Ulf Hlobil notes that a version of this idea goes as far back as Bertrand Russell, "The Nature of Inference," *The Athenæum* 4694 (1920): 514-15.

³ Some distinguish reasoning from inferring, such as Nicholas Koziolok, "Inferring as a Way of Knowing," *Synthese* (2017), <https://doi.org/10.1007/s11229-017-1632-4>. Like most others, however, I use these terms interchangeably.

⁴ See Christian Kietzmann, "Inference and the Taking Condition," *Ratio*, 31, 3 (2018): 294-302; Nicholas Koziolok, "Inferring as a Way of Knowing"; Andreas Müller, "Reasoning and Normative Beliefs: Not Too Sophisticated," *Philosophical Explorations*, 22, 1 (2019): 2-15.

⁵ See Elijah Chudnoff, *Intuition* (Oxford University Press 2013); John Broome, *Rationality Through Reasoning* (Wiley Blackwell 2013).

⁶ See Ulf Hlobil, "Inferring By Attaching Force."

⁷ Boghossian, "What Is Inference?"

⁸ Skeptics include Kieran Setiya, "Epistemic Agency: Some Doubts," *Philosophical Issues: A Supplement to Noûs*, 23 (2013): 179-198; Crispin Wright, "Comments on Paul Boghossian, 'What Is Inference,'" *Philosophical Studies*, 169 (2014): 27-37; Conor McHugh and Jonathan Way, "Against the Taking Condition," *Philosophical Issues*, 26, *Knowledge and Mind*, (2016): 314-331; Alex Kiefer, "Literal Perceptual Inference," in *Philosophy and Predictive Processing*: 17, eds. Thomas Metzinger and Wanja Wiese (MIND Group, 2017): 1-19; Luis Rosa, "Reasoning Without Regress," *Synthese*, 196 (2017): 2263-2278; Ladislav Koreň "Have Mercier and Sperber Untied the Knot of Human Reasoning?" *Inquiry* (2019), <https://doi.org/10.1080/0020174X.2019.1684988>; Susanna Siegel, "Reasoning Without Reckoning," in *Reasoning: New Essays on Theoretical and Practical Thinking*, eds. Brendan Balcerak Jackson and Magdalena Balcerak Jackson (Oxford University Press, 2019), 15-31; Mark Richard, "Is Reasoning Something the Reasoner Does?" in *Reasoning: New Essays on Theoretical and Practical Thinking*, eds. Brendan Balcerak Jackson and Magdalena Balcerak Jackson (Oxford University Press, 2019), 91-100.

(Self-Awareness Condition): “Person-level reasoning [is] mental action that a person performs, in which he is either aware, or can become aware, of why he is moving from some beliefs to others.”⁹

His view seems to be that this Self-Awareness Condition (SAC) is either a direct upshot of TC or is a different way of articulating it.

On the assumption that talk of self-awareness is interchangeable with talk of self-knowledge (ditto for talk of self-consciousness), and on the assumption that TC indeed leads to or amounts to SAC, TC may underpin an interesting “agentialist” account of self-knowledge.¹⁰ Agentialist accounts of self-knowledge argue that, due to the nature of our agency, self-knowledge cannot be acquired by ordinary empirical methods (though they can sometimes leave open exactly how self-knowledge is acquired instead). Here, an agentialist might claim that, because inference presupposes self-knowledge (as per SAC), at least one ordinary empirical route to self-knowledge—the inferential route—is closed off. At least, it follows that not all self-knowledge can be acquired inferentially, on such a view.¹¹ Some agentialists may even want to argue for something stronger here. After all, if self-knowledge or something close to it is required in order to infer, and if a capacity for inference is basic to our rational agency, then self-knowledge of at least some of our mental states may seem like a necessary rather than contingent

⁹ Boghossian, “What Is Inference?”, 16. The label comes from Siegel, “Inference Without Reckoning”. Others who endorse SAC include Sebastian Rödl, *Self-Consciousness* (Harvard University Press, 2007), Matthew Boyle, “Transparent Self-Knowledge,” *Proceedings of the Aristotelian Society, Supplementary Volumes*, 85 (2011): 223-241, and Kietzmann, “Inference and the Taking Condition”.

¹⁰ I take this label from Brie Gertler, “Self-Knowledge and Rational Agency: A Defense of Empiricism,” *Philosophy and Phenomenological Research*, 96, 1 (2016): 91-109, and Ben Sorgiovanni, “The Agential Point of View,” *Pacific Philosophical Quarterly*, 100, 2 (2018): 549-572. Sorgiovanni further distinguishes between non-substantive and substantive agentialist accounts. Non-substantive agentialist accounts (my focus) account for the functions of self-knowledge while leaving open its source(s). Substantive agentialist accounts, in contrast, argue that rational agency is a source of self-knowledge.

¹¹ Pace Gilbert Ryle, *The Concept of the Mind* (University of Chicago Press 1949; reprinted in 2009 by Routledge). Some epistemologists argue that, while self-knowledge is inferentially acquired, the relevant inferences are sub-agential. See, for example, Quassim Cassam, *Self-Knowledge for Humans* (Oxford University Press, 2015). These accounts are harder to dismiss by way of the present agentialist suggestion.

feature of our psychological lives. In that event, any account of self-knowledge that treats it as a contingent cognitive achievement will be objectionable.

To be sure, this agentialist account may have its limits. First, note that SAC includes a crucial “or can become aware” qualifier. As such, SAC does not entail that we actually have self-knowledge in inferring. Now, as we will see below, this qualifier is dropped by several arguments for SAC. But even in its unqualified version this agentialist argument has its limits. For one thing, it *only* proposes a constitutive connection between rational agency and self-knowledge of those mental states that figure into one’s inferences. So pains, tickles, various emotions, and more will not fall under the purview of this agentialist account. Moreover, this account might only deliver a constitutive connection between rational agency and self-knowledge *during the inferential process*, such that it says nothing about self-knowledge of standing attitudes that are not occurrently deployed in one’s inferences.¹²

Despite the limits of this agentialist account, it is surely interesting if SAC is true. In what follows, however, I will argue that extant arguments for TC do not establish SAC. In other words, I will argue that SAC is *not* equivalent to (or an upshot of) TC. This means that, even if TC is true, no agentialist conclusions like the above follow. To reach this conclusion, I will evaluate many arguments for TC. These will be arguments that appeal to TC in order to illuminate: (1) the inference/association distinction, (2) the good/bad inference distinction (3) a Moore-paradoxical phenomenon associated with inference, (4) inference as a mental act, (5) inference as involving cognitive agency, and (6) rational responsibility for inferences. I will

¹² On this point, one might respond that (with some further machinery in place, no doubt) SAC grounds a constitutive connection between rational agency and those mental states that are “available” for inferential application, whether or not they are occurrently embedded in an inference at any given time. To see how this might go, see Sydney Shoemaker’s discussion of available belief in “Self-Intimation and Second-Order Belief,” *Erkenntnis*, 71, 1 (2009): 35-51.

argue that none of these arguments lead us from TC to SAC. Before I evaluate these arguments, however, I want to articulate a conception of inference that accepts TC but rejects SAC. Once we have this conception of inference in view, we can see whether the arguments for TC encourage us to accept this conception instead of one that also accepts SAC.

§2—Robust and Minimal Inference

Answering the question of whether taking your premise(s) to support your conclusion requires self-knowledge depends on whether there can seem to be any alternative conception of takings.

Consider, then, the following toy inference schemas:

(Minimal): *p. p provides sufficient epistemic support for q. Therefore, q.*

(Robust): I believe that *p. p provides sufficient epistemic support for q. Therefore, I now believe that q.*

At least on the surface, the difference between (Minimal) and (Robust) is that the former does not involve thoughts about mental states. Nevertheless, (Minimal) seems to involve some sort of appreciation (taking) of epistemic support: it seems to involve what we might call a ‘meta-propositional’ as opposed to a ‘meta-attitudinal’ taking-attitude.¹³ Thus, on what I will call *the minimal view*, there can be instances of agent-level inference that involve nothing over and above what Christopher Peacocke calls “second-tier” thought: thought that is about “relations of support, evidence, or consequence between contents”.¹⁴ Such thought is ‘second’ rather than ‘first’ tier because it represents epistemic relations between propositional contents rather than *just* representing said contents. But it is not (we might say) ‘third-tier’ because it does not also include content about the subject’s own attitudes toward these contents. *Ex hypothesi*, then, a

¹³ Cf. Philip Pettit, “Broome on Reasoning and Rule-Following,” *Philosophical Studies*, 173 (2016): 3373-3384.

¹⁴ Christopher Peacocke, “Our Entitlement to Self-Knowledge: Entitlement, Self-Knowledge and Conceptual Redeployment,” *Proceedings of the Aristotelian Society*, 96, 1 (1996): 117-58. See especially pp. 129-130.

putatively second-tier inference like (Minimal) cannot require the agent to conceptualize p and q as *believed*. It also seems dangerous to allow that p and q are conceptualized in such an inference as *reasons for belief*. Instead, the epistemic concepts deployed in appreciating an epistemic support relation between p and q should squarely concern indicators of what is or is (probably) *true*. Thus, concepts like EVIDENCE and CONSEQUENCE may be involved, as long as these capture relations between propositions only.

Of course, one can harbor different attitudes toward p , and it is true that whether one hopes, believes, or doubts p can have drastic effects on what sorts of inferences one can and should draw from p . Because of this, it may seem that inferring something from p obviously requires self-knowledge, since it might seem to require one to know the specific attitude one has toward it. However, John Broome has articulated an alternative possibility, according to which a second-tier reasoner need only view p “in a believing way” in theoretical inference,¹⁵ or to view it in an intending or desiring way in the practical case. The idea is that viewing a content in a believing, desiring, or intending way is not equivalent to having a higher-order belief about that attitude. This may sound strange, but it need not. After all, first-order propositional attitudes are just the same: one views p in a believing way *by believing it*, for instance. Similarly, second-tier inferences might involve taking-attitudes that involve viewing (merely perceived or actual) epistemic support relations between propositions in a believing way.

In fact, the minimal view has been repeatedly presented as a major objection to different and highly influential agentalist account of self-knowledge, one that argues for a constitutive connection between self-knowledge and our capacity for a specifically “critical” form of

¹⁵ “A Linking Belief is not Necessary for Reasoning,” in *Reasoning: New Essays on Theoretical and Practical Thinking*, eds. Brendan Balcerak Jackson and Magdalena Balcerak Jackson (Oxford University Press, 2019), 32-43.

reasoning.¹⁶ Tyler Burge describes critical reasoning as reasoning undertaken with the aim of evaluating and adjusting one's attitudes so that they better conform to norms of rationality, and he argues that such reasoning requires a non-empirical form of access to one's mental states. The objection from the minimal view is that agents can adjust their attitudes just fine, and in a sufficiently critical way, even if they never conceptualize their attitudes as such and only focus, instead, on propositional contents and relations of epistemic support between them.¹⁷ Now, if SAC is true, then the minimal view is in danger. However, the significance of Burge's agentialist account will also be threatened, since we won't need to focus on *critical* reasoning to establish an important connection between inference and self-knowledge. Because of this, making the case for or against the minimal view is also significant for debates between Burge-style agentialists and their opponents.

Opposed to the minimal view is what I will call *the robust view*. This is the view that I take authors like Boghossian to accept, given their acceptance of SAC. On this view, agent-level inference necessarily requires at least being in a position to appreciate epistemic support relations between one's own inferential mental states themselves. Crucially, this qualifier allows proponents of the robust view to grant that agents do occasionally second-tier infer. What they must argue is only that any agent that can second-tier infer is also in a position to know her inferentially embedded mental states. Nevertheless, as aforementioned, we will examine many arguments for the robust view that appear to drop this qualifier and so aim to establish a stronger connection between inference and self-knowledge.

¹⁶ Burge, "Our Entitlement to Self-Knowledge," p. 98.

¹⁷ See David Owens, *Rationality Without Freedom* (London: Routledge 2000) and "Deliberation and the First-Person," in *Self-Knowledge*, ed. Anthony Hatzimoysis (Oxford University Press 2011), 261-278; Cassam, *Self-Knowledge for Humans*; Annalisa Coliva, *The Varieties of Self-Knowledge* (Palgrave Macmillan: London, England 2016); Anna-Sara Malmgren, "On Fundamental Responsibility," *Philosophical Issues: A Supplement to Noûs*, 29, 1 (2019): 198-213; Broome, "A Linking Belief is Not Necessary for Reasoning."

We can now clarify the question of this paper as the question of whether the minimal view or the robust view is true. Before we turn to evaluating arguments for TC in order to see whether SAC follows, such that the robust view prevails over the minimal view, four final preliminary points are in order.

First, as one might have guessed, I will be focusing (with Boghossian) on *theoretical* inference in what follows and, even more narrowly, on *non-suppositional* inferences where one's premises are one's actual beliefs. These decisions are intended to simplify and shorten discussion, though I cannot rule out the possibility that they might cause me to unwittingly bypass important complications.¹⁸

Second, let me pre-empt a possible concern: if second-tier taking-attitudes are so much as *intelligible*, doesn't it follow that second-tier inference is possible and, hence, that we can second-tier infer? And if so, isn't the robust view obviously false? The answer is *no*. For it is open to a proponent of the robust view to say that, while second-tier taking-attitudes might be *components* of agent-level inference, they must be supplemented with self-knowledge, because of further desiderata for any good account of (agent-level) inference. Alternatively, as aforementioned, proponents of the robust view can grant that we can second-tier infer, while arguing that agents who can second-tier infer must still be in a position to know their inferentially implicated mental states.

Third, I do not deny that *some* inferring might require self-knowledge, as when we reason explicitly about whether we ought to have the beliefs that we do. But now one could ask: why

¹⁸ One possible complication comes from John Broome, who sees taking-attitudes in practical inferences as unintelligible. See his "Comments on Boghossian," *Philosophical Studies*, 169 (2014): 19-25, and "A Linking Belief is not Necessary for Reasoning." For a rejoinder, however, see Markos Valaris, "What Reasoning Might Be," *Synthese*, 194 (2017): 2007-2024. Another possible complication is that practical inference may require self-knowledge for reasons having nothing to do with TC—see Owens, "Deliberation and the First-Person."

can't an agentalist simply argue that some inference requires self-knowledge? I suppose one can argue this, but it should disturb us that such a view trivializes the desired connection between self-knowledge and inference, for it amounts to saying no more than that inferences that are explicitly about one's own mental states require us to have knowledge of them. Importantly, philosophers like Boghossian make a stronger claim: they claim even inferences about matters besides one's own mental states require one to (be in a position to) know how one's mental states epistemically support other mental states that one has.

A fourth and final point concerns the question of how "explicit" one's self-knowledge must be according to the robust view. I have been writing as if self-ascriptive, higher-order beliefs involving mental state concepts must be involved, but perhaps different versions of the robust view could take on different commitments here. For example, one version of the robust view might attribute "implicit" or "tacit" self-knowledge to agent-level reasoners that does not require explicit deployments of mental state concepts. Alternatively, one could argue that, while inferring does require actively forming higher-order beliefs that conceptualize one's lower order mental states, these can lie in the "background" of our consciousness when we draw inferences.¹⁹ I will typically characterize the robust view as the view that inferring requires taking-attitudes with a conceptualized, psychological self-ascriptive structure, consciously foregrounded or otherwise. This is partially to avoid repeatedly appending cumbersome qualifiers about the possibilities of tacit or backgrounded taking-attitudes to claims made in the ensuing discussion.

¹⁹ Müller, "Reasoning and Normative Beliefs: Not Too Sophisticated," 8.

But it is also because I am somewhat skeptical of the idea of tacit self-knowledge,²⁰ and so take as my target the clearest (by my lights) version of the robust view.²¹

§3—Arguments for the Taking Condition

It is now time to look at the arguments for TC, and to see whether they entail or sufficiently motivate SAC, and hence motivate the robust view over the minimal view. I will consider six such arguments in total (I run two together in §3.1) and show that they do not motivate the robust view.

§3.1—Inference (Good and Bad) and Association

Inference is but one way for agents to form a mental state. A non-inferential cause of many mental states is *association*. An agent might associate by being subject to some “bizarre psychology experiment” where she is conditioned to think that the sun will one day explode every time she thinks that Donald Trump is the President of the United States.²² Or we might imagine a more ordinary case of a habitual depressive who, whenever she thinks about how much fun she is having, also thinks that there is so much suffering in the world.²³ Both cases

²⁰ See, e.g., Coliva, *The Varieties of Self-Knowledge*, chapter seven, for her critical discussion of Sydney Shoemaker’s appeal to tacit self knowledge in his *The First-Person Perspective and Other Essays* (Cambridge: Cambridge University Press 1996).

²¹ A final preliminary is not a caveat for what follows but, rather, a comment about a different dialectical context in which the robust view might be motivated. I have in mind the possibility that the robust view can solve a problem for TC even if it does not motivate it. Many philosophers worry that TC gives rise to an infinite regress inspired by Lewis Carroll, “What the Tortoise Said to Achilles. *Mind*, 4, 14: 278-280.” The regress worry is that, if taking-attitudes are themselves premises of inferences, they must also be taken by further-taking attitudes to support one’s other premises, *ad infinitum*. There are various TC-friendly strategies for avoiding this regress—see, e.g., Kietzmann, “Inference and the Taking Condition; Müller, “Reasoning and Normative Beliefs: Not Too Sophisticated,”; Marcus, “Inference as Consciousness of Necessity,” *Analytic Philosophy* (2020): 1-19, <https://doi.org/10.1111/phib.12153>. Kietzmann’s particular strategy appeals to the robust view. Here I simply note that his strategy, according to which the regress is halted by conceiving of one’s taking-attitudes as ontologically indistinct from their objects, might be made to work on a second-tier conception of taking-attitudes as well. I also believe—though I cannot argue it here—that Marcus’s strategy does not suffer from the defects that Kietzmann points out for other strategies, and that Marcus’s account of inference does not seem to depend on the robust view.

²² Jake Quilty-Dunn and Eric Mandelbaum. “Inferential Transitions,” *Australasian Journal of Philosophy*, 96, 3 (2018): 532-547.

²³ Boghossian, “Inference, Agency, and Responsibility,” 112.

involve mental states that are in *some* way sensitive to one another. But they are not sensitive in an inferential way. This is why it is not enough to say, as Hilary Kornblith does, that inference is simply a matter of “transitions involving the interaction among representational states on the basis of their content”,²⁴ for the interaction also needs to be of the right sort.

Might one argue that association and inference can be distinguished by pointing out that, in the latter case, one’s mental states stand in epistemic support relations, whereas no such support obtains in the associative case? On the face of it, this proposal can even dispense with TC: it can account for the inference/association distinction in terms of the presence or absence of epistemic support relations, without countenancing any appreciation of these relations on the agent’s part. The problem with this is that epistemic support relations can obtain between attitudes even when they are not occurrently involved in an inference.²⁵ Another problem is that attitudes can be related inferentially despite the *absence* of genuine epistemic support relations between them. After all, as Boghossian says, this seems to be what happens with *bad* inferences.

In light of all this, Boghossian argues that:

...something like a taking-based account seems not only natural, but forced: the depressive’s thinking doesn’t count as reasoning *not* because his first judgment doesn’t *support* his second, but, rather, it would seem, because he doesn’t *take* his first judgment to support his second. The first judgment simply *causes* the second one in him; he doesn’t draw the second one because he takes it to be supported by the first.²⁶

So, TC makes sense of the difference between associative and inferential mental state transitions as well as the difference between good and bad inferences. For, while in both good and bad inferences, the agent takes her premise(s) to support her conclusion, inference is only good when the epistemic support that she takes her premise(s) to provide is actual. Either way, the

²⁴ Hilary Kornblith, *On Reflection* (Oxford University Press 2012), 55.

²⁵ E.g., one might associate thoughts in a way that matches the pattern of a modus ponens inference without actually being a modus ponens inference (Quilty-Dunn & Mandelbaum, “Inferential Transitions 13).

²⁶ Boghossian, “Inference, Agency, and Responsibility”, 112.

psychological relationship between the mental states at issue is not merely associative, because associations do not involve takings.

Supposing that Boghossian is right about TC's indispensability here,²⁷ what have we learned about the essential contents of taking-attitudes? I do not think that we have been led to the robust view. This is because proponents of the minimal view could argue that an agent who thinks that p , and that p supports q , infers q on the basis of appreciating (by her lights) an epistemic support relation between two world-directed propositions. In the good case, an agent takes p to support q and p does support q . In the bad case, she takes p to support q but p does not support q . In this way, it does not seem that she is required to think about her *beliefs* in p and q . Moreover, we can agree with Boghossian that she is not merely associating.

According to Nicholas Koziolk, however, this is too quick. He argues that we should prefer an account of the good/bad inference distinction that appeals to an agent's self-knowledge. To see this, consider first how he understands bad inferences: to infer badly is, on his view, to associate while taking it that you have actually formed a conclusion-belief on the basis of a premise-belief,²⁸ whereas good inferences involve no such mistakes. Koziolk grants that this is a paradoxical-sounding view, since it means describing bad inferences as, in fact, just perverse associations. Still, I will grant his claim that these deserve to be called inferences "if only by a sort of courtesy".²⁹

On this new account, one's second-order perspective on one's attitudes is doing explanatory work, since appealing to this perspective allows us to countenance two mistakes involved in bad inferences: (1) a false belief that one's beliefs have inferentially caused one's

²⁷ Though see Siegel, "Inference Without Reckoning," for a non-TC account of the association/inference distinction.

²⁸ Koziolk, "Inferring as a Way of Knowing," 18.

²⁹ Koziolk, "Inferring as a Way of Knowing," 19.

conclusion in a non-deviant way (they haven't, since one's inference *does* involve deviant causation) and (2) a false belief that one has gained knowledge via her inference (one hasn't, since the cause of one's conclusion-belief is not a rational cause). Crucially, while one surely *lacks* a degree of self-knowledge in making these mistakes, one actually still possesses some self-knowledge. For, on this view, one has self-knowledge insofar as one has knowledge *of the beliefs involved* in the defective inference, even though one lacks self-knowledge of the causal connections between these beliefs.³⁰

This account allows us to make sense of bad inferences other than those where an agent takes *p* to support *q* even though, as it happens, *p* turns out to be false. This is especially interesting if Koziolok is right, as he argues, that inferences where *p* fails to support *q* merely because *p* turns out to be false are *not* actually bad. After all, though we can grant that such inferences fail to yield knowledge, it is still the case that, *were p to be true*, one's inference could yield knowledge. For this reason, Koziolok argues that such inferences are actually good because they are "potentially productive of knowledge".³¹ If so, then Boghossian's understanding of the good/bad inference distinction fails, because the bad inferences he describes are potentially productive of knowledge, and so we may need Koziolok's alternative account. And since Koziolok's account appeals to the agent's self-knowledge, it may seem that we also need to accept the robust view.

Two concerns about Koziolok's account, for my purposes, are these. The first concerns a further feature of his account, which is that he avoids treating taking-attitudes as causal ingredients of inferences and argues instead that "takings are second-order beliefs about

³⁰ Koziolok, "Inferring as a Way of Knowing," 18.

³¹ Koziolok, "Inferring as a Way of Knowing," 11.

inferences...that are formed *in response to* the inferences themselves”.³² So the robust view of TC is not a consequence of his account, because his account actually rejects TC.

Of course, I have simply taken TC for granted in this paper, in order to explore whether TC supports SAC. As a result, one may not be too worried about the above concern; after all, TC may turn out to be false, after all. But even if TC is true, Koziolak may be right that proponents of the minimal view need a new way of accounting for the good/bad inference distinction.

I think that the minimal view can countenance at least two species of bad inference. First, bad inferences can involve agents who fail to accord *the right amount* of evidential weight to p in concluding q , even if they take p to support q to some degree and p really does support q to some degree.³³ In such a case, the agent’s focus may strictly be on these world-directed evidential relations between propositions. Second, bad deductive inferences might have the following structure: (1) an agent takes p to support q , (2) p does not in fact support q , and (3) the inference is *not* “potentially productive of knowledge”. For example, someone who affirms the consequent infers badly, but this does not mean that her conclusion is deviantly caused by an association (as Koziolak would have it).³⁴ Rather, while she genuinely infers, she does so in accordance with a bad first-order inference schema, one that is not potentially productive of knowledge. In other words, affirming the consequent need not amount merely to associating p and q via, say, some sort of conditioning that produces a sort of mental “jogging”³⁵ that *looks like* affirming the consequence. A test for whether one merely associates in a way that looks like affirming the consequent, as opposed to actually affirming the consequent, is this: if an agent’s mental activity

³² Koziolak, “Inferring as a Way of Knowing,” 6.

³³ For a similar suggestion, not tethered to a second-tier conception of inference, see Siegel, “Inference Without Reckoning,” 29.

³⁴ Contra Koziolak, “Inferring as a Way of Knowing,” 19, fn. 29, who argues that there are no fallacious inferences without self-consciousness.

³⁵ Broome, *Rationality Through Reasoning*, 226.

satisfies the pattern of affirming the consequent for different values of p and q , this is evidence that she is not merely associating, since associations are content-based conditionings, whereas inference schemas generalize over different contents.

Perhaps only Koziolak's account can explain the bad inferences he describes. But whether being able to do this is something that *all* reasoners must be able to do simply in virtue of being reasoners is a different matter. The proponent of the minimal view can deny that second-tier reasoners can infer badly in Koziolak's sense while arguing that they can nevertheless draw bad inferences. Moreover, as we have seen, they can do so while accounting for the difference between association and inference. Two essential desiderata of an account of inference are hereby satisfied, but they appear to be satisfied along minimalist lines. The robust view has yet to follow.

§3.2—Inference and Inferential Absurdity

Some philosophers have noted that inferences can figure into a version of Moore's Paradox.

Thus, consider what Ulf Hlobil calls "Inferential Absurdity"³⁶, or INFA:

INFA. It is irrational, and transparently so from the agent's own perspective, to infer B from A1, ..., An and to believe also that these premises don't support B or to suspend judgment on whether they do.

Thus, if an agent infers q from p , but also believes that p does not support q (or is ambivalent about whether this is so), she is manifestly irrational. Hlobil points out that INFA has the air of a Moorean paradox. The reason, he says, is that inferences are *acts*,³⁷ and acts are not content-bearing vehicles that can be in tension with *states* like beliefs (though of course inferences *operate on* content-

³⁶ Hlobil, "Inferring by Attaching Force," 2.

³⁷ Though note that in an earlier paper he does not assume that inferential acts are intentional acts. See Ulf Hlobil "Against Boghossian, Wright, and Broome on Inference," *Philosophical Studies*, 167, 2 (2014): 419-429, specifically 421, fn. 1. Note that he does not assume that inferential acts are intentional acts. I take up the question of whether inferences are intentional acts in §3.3.

bearing states like beliefs). So we face the question: “[h]ow can a doing that seems to have no content be in rational tension with a judgment or a belief?”³⁸

Now, if TC is true, this absurdity is readily explained. For, as Christian Kietzmann puts it, “[i]f inference involves the thinker taking his premises to support his conclusion, this taking-attitude will clash with a belief or judgement that the premises do not support the conclusion”.³⁹

Once again, however, we can ask whether this explanation of INFA also motivates SAC and, hence, the robust view. On my view, the minimalist could accommodate INFA by explaining it in terms that do not require self-awareness, on the reasoner’s part, of the tension between her taking-attitude and her further belief that there is no epistemic support relation between the premise(s) and conclusion of her inference (or her agnosticism about this epistemic support relation). Perhaps this can be done by supposing that, while an agent draws an inference such as:

(Minimal): *p*. *p* epistemically supports *q*. Therefore, *q*.

While simultaneously believing either of:

- (1) *p* may or may not support *q*, or;
- (2) *p* does not support *q*

...then the contents of her mental states will contradict. But, *ex hypothesi*, (Minimal) and (1)-(2) are minimally contentful. They take propositional contents and relations between these as their objects, rather than taking mental states as their objects.

What about the fact, as Hlobil sees it, that inferential absurdities are “transparently so from the agent’s own perspective”? One way to go is to read this as a claim about a second-tier rather than a self-conscious perspective. But even if we are thinking about this perspective as a self-

³⁸ Hlobil, “Against Boghossian, Wright, and Broome on Inference,” 421.

³⁹ Kietzmann, “Inference and the Taking Condition,” 295.

conscious one, the minimalist can also argue that it is unclear whether we must understand Hlobil's insistence on the transparency, to the subject herself, of inferential absurdities as an essential feature of the phenomenon. Hlobil seems to stipulate that it is. But if the minimalist's reply above is roughly correct, then inferential absurdities seem to be one thing and their self-conscious appreciation another. What proponents of the robust view need, then, is an argument for two claims: (1) that this appreciation is both necessarily self-conscious (Hlobil does not argue for this) and (2) that this appreciation is necessary to being an agent-level reasoner.

Unfortunately, (2) seems to implausibly assume that agent level reasoners are vulnerable to inferentially absurd states of mind unless they self-consciously ward them off. It is implausible because warding against irrationality is not always a matter of having to actively prevent one's self from adopting clashing attitudes, for part of what it is to be an agent is to be by and large rational.⁴⁰ But even if some such monitoring is required, why won't 'second-tier' monitoring do?

Perhaps it is worth closing this subsection by noting that other Moore-paradoxical phenomena might also be explained without appealing to an agent's self-knowledge, *even when* the phenomena necessarily invoke an agent's capacity to *self-attribute* mental states. Take, for example, the original Moore's Paradox, which concerns thoughts and utterances like "*p*, but I don't believe that *p*." Such thoughts or utterances seem highly irrational despite the fact that their conjuncts do not formally contradict (just as INFA begins with the thought that inferences and beliefs can stand in rational tension without formally contradicting). Now, as several philosophers have argued,⁴¹ a defining featuring of ordinary, first-personal self-ascriptions of

⁴⁰ Cf. Donald Davidson, "Radical Interpretation," *Dialectica*, 27 (1973): 314-28; McHugh and Way, "Against the Taking Condition," 322.

⁴¹ See Rockney Jacobsen, "Wittgenstein on Self-Knowledge and Self-Expression," *The Philosophical Quarterly*, 46, 182 (1996): 12-30; Kevin Falvey, "The Basis of First-Person Authority," *Philosophical Topics*, 28, 2 (2000): 69-99; Dorit Bar-On *Speaking My Mind: Expression and Self-Knowledge* (New York: Oxford University Press 2004).

one's current mental states—what many call “avowals”⁴²—is their *expressive* dimension. Thus, according to the expressivist view of avowals, I can express my first-order mental states by avowing them, whereas others can only *report* on them. What this means is that my avowal, but not your ascription to me of the same state I avow, can directly manifest the avowed state. Expressivists have appealed to this function of avowals in order to account for our so-called first-person authority.⁴³ But they have also shown that this account of avowals lends itself to a “smooth account” of Moore's paradox (Jacobsen 1996, p. 28).

The expressivist account of Moore's paradox is, roughly, this. An avowal of “*p*, but I don't believe that *p*” *expresses* both a first-order belief that *p* and either (1) a first-order belief that not-*p*, or (2) a first-order agnostic attitude towards *p*. Because of this, the rational tension consists in a first-order “expressive conflict”.⁴⁴ All the while, the utterance contains no contradiction at the level of its semantic meaning. Moore's paradox is dissolved when we recognize that expressive force and semantic meaning can come apart. But notice that the explanation does not invoke any claims about the speaker's self-knowledge. Rather, it invokes claims about a rational tension at the level of her first-order attitudes.

Crucially, these Moorean thoughts and utterances have an irreducibly *de se* conjunct. After all, if the target utterance or thought was merely “*p*, but not *p*” rather than “*p*, but I don't believe *p*”, the irrationality manifest in such an utterance or thought would be visible at the semantic level, and would hardly constitute a puzzle of any sort. My point in raising the expressivist analysis of Moore's Paradox, then, is this. Because it is possible to give an account of it that does not depend on the subject's higher-order perspective on her attitudes (again, the tension exists as

⁴² See Bar-On, *Speaking My Mind*.

⁴³ First-person authority is usually defined as a distinctive presumption of truth owed to or conferred on avowals. Alternatively, it is sometimes defined as the relative indubitability of by avowals.

⁴⁴ Bar-On, *Speaking My Mind*, 217-219.

a first-order expressive conflict), despite the fact that the paradox cannot even be set up without granting that the subject self-ascribes her belief, this should make us doubly confident that INFA can be explained without appeals to self-knowledge. This is because there is at least an apparent possibility of understanding the components of INFA (the inference and the incompatible belief) as second-tier rather than robustly contentful, unlike the original Moore's Paradox.

§3.3—Inference and Practical (Self-)Knowledge in Action

Proponents of TC frequently argue that inference is *active*, and that its activeness explains why inference is attributable to the agent herself. The connection between taking and inferring can now be put as follows: “[a]ppreciating the support relation between premises and conclusion and drawing the conclusion on account of that appreciation seem to be things persons actively engage in. Inference will then count as something persons do because it involves the person-level activity of taking”.⁴⁵

Although I am confident that this thought motivates many philosophers to embrace TC, its implications for characterizing the essential contents of taking-attitudes are not straightforward. One might think that there are straightforward implications if one thinks that there is a constitutive connection between *mental action* and *practical knowledge* of what one is doing. One could take one's cue here from Elisabeth Anscombe, who famously argued that non-observational, non-testimonial, and non-inferential knowledge of what one is doing is constitutive of acting intentionally.⁴⁶ The idea now is that, if the activeness of inference is the activeness of intentional (mental) action, then TC could fall out as a consequence. For, if inference requires practical knowledge of what one is doing, one's taking-attitudes might be the

⁴⁵ Kietzmann, “Inference and the Taking Condition,” 295.

⁴⁶ G. E. M. Anscombe, *Intention*. (Oxford: Blackwell 1963).

very site of such knowledge: taking-attitudes might be the form of practical knowledge-in-infering. Put differently: in infering, one's practical knowledge of what one is doing takes the form of a taking-attitude to the effect that one's conclusion derives epistemic support from one's premise-beliefs.

In reply, one might question whether there really is a constitutive connection between intentional action and practical knowledge (however frequently these may come together as a matter of fact).⁴⁷ One might also wonder whether there can be another description under which one knows what one is doing, in infering, that does not presuppose knowledge of one's mental states.⁴⁸

I will not pursue these possibilities here. Instead, I will begin by raising the possibility that inference is not a species of action at all. Good evidence for this consists, as Kieran Setiya and Casey Doyle argue, in the grammar of inference-talk.⁴⁹ We do not say, for example, that we are in the middle of drawing an inference from p to q , even though we may be in the middle of considering whether p is evidence for q .⁵⁰ This suggests that the term 'infering' and its cognates do not have the grammar of ordinary process verbs, in that they lack intelligible progressive aspects.⁵¹

⁴⁷ "Practical Knowledge and Luminosity," *Mind*, <https://doi.org/10.1093/mind/fzz056>

⁴⁸ Thanks to Dorit Bar-On for this suggestion.

⁴⁹ Setiya, "Epistemic Agency: Some Doubts"; Casey Doyle, *Four Essays on Self-Knowledge*, Dissertation: University of Pittsburgh. <http://d-scholarship.pitt.edu/24970/>

⁵⁰ Doyle, *Four Essays on Self-Knowledge*, 105. For similar observations, see Quilty-Dunn & Mandelbaum, "Inferential Transitions," 14.

⁵¹ Similarly, believing also consists in taking a proposition to be true in some sense or other. But beliefs are states, not actions. (Setiya, "Epistemic Agency: Some Doubts"; Nicholas Koziol, "Belief as an Act of Reason," *Manuscripta*, 41, 4 (2018): 287-318. Perhaps *judgements* are mental actions, but even these do not seem to be *intentional* ones.

But let us suppose, for the sake of argument, that inference may yet be a kind of action.⁵² Another concern is that it does not seem to be a *voluntary* act whether one infers *q* from *p*, since one cannot simply *decide* to infer *q* from *p*; the alternative suggests an implausibly strong form of doxastic voluntarism that I will not bother to argue against here. But it is natural to think that intentional actions are standardly voluntary actions, and so if inferences are not voluntary this may be another source of pressure against the claim that they are intentional.

It might be argued that inferences are intentional despite not being voluntary. On this argument, inferences involve *intentions-in* the act, analogously to how I can intentionally albeit reflexively (and so, in one sense, non-voluntarily) raise my arm as a basketball is being hurled at my head.⁵³ Such actions are not intentional because preceded by a decision, but because they have an appropriate means-end telos. However, it may seem that inferences are *never* such that one can simply decide to perform them, whereas even actions like raising my arm to block a basketball can *sometimes* be. Moreover, even when an action is reflexive and hence not voluntary, it is usually something that can be overridden: I can lower my arm quickly after it reflexively raises. But this is not possible with inference. With inference, one understands that there is an epistemic support relation between two propositions and draws a conclusion without choosing to do so or choosing to stop.⁵⁴ A plausible explanation is that inferences are not intentional actions.

It might be replied that inferences *can* in fact be voluntary, and so the argument that inferences are not intentional because not voluntary is a bad one. For example, David Hunter

⁵² Perhaps Doyle's argument from the non-processual nature of inference to its non-actional nature is too quick if there can be instantaneous mental acts (this, it seems to me, is one way of reading Hlobil in "Inferring By Attaching Force").

⁵³ The example comes from Parent, *Self-Reflection for the Opaque Mind: An Essay in Neo-Sellarsian Philosophy*. (Routledge 2017), 186

⁵⁴ Cf. Marcus, "Inference as Consciousness of Necessity."

focuses on cases where an agent's evidence for or against p is strong enough to license an inference in either direction.⁵⁵ He concludes that it is up to you (i.e., is voluntary) whether you infer p or $\sim p$ from the evidence. But this at most shows that *some* inferences are voluntary. For, in cases where the evidence strongly favours only one conclusion p , Hunter agrees that one cannot voluntarily infer p . So, if voluntariness is a sign of intentional action, it will only be a sign that some corner cases of inference are intentional actions.

More importantly, Hunter himself concedes that inferences are not intentional actions even when they are voluntary. For he follows John Hyman in arguing that the voluntary and the intentional simply have different scopes: voluntary actions depend on a lack of coercion and lack of ignorance, whereas intentional actions involve doing something to satisfy a desire.⁵⁶ Inferences do not aim at desire satisfaction, even if they occur without coercion or ignorance. But since Anscombe's thesis is that there is a connection between intentional action and self-knowledge, showing that inference may sometimes be voluntary is insufficient to generate a tight connection between self-knowledge and inference even in the cases that Hunter focuses on. I conclude that there is no obvious move from Anscombe's practical knowledge thesis to the robust view: inferences may not even be actions, but if they are, and even if they are sometimes voluntary, they do not seem to be intentional.

§3.4—Inference and Cognitive Agency

Even if one agrees that inferences are not intentional actions, one might think that inferences are nevertheless active in *some* sense, and that TC can help us to understand this activeness. In other words, TC might still help us to understand how inference is a site of cognitive agency.

⁵⁵ David Hunter, "Inference as a Mental Act," unpublished.

⁵⁶ John Hyman, *Action, Knowledge and Will* (Oxford University Press 2015).

To see how we might proceed, note first that there exist fairly uncontroversial characterizations of the activeness of *mental states* that do not trade on the notion of an intentional action. Thus, consider Joseph Raz's claim that our "beliefs are a product and an aspect of our active nature because they are responsive to reasons".⁵⁷ The domain of active attitudes has contestable boundaries within and beyond the category of belief, but it is plausible that any attitudes that can figure into inferences are active.

Now, whatever philosophers typically mean when they talk about cognitive agency, Doyle doubts that they are merely parroting Raz's claim that many of our cognitive attitudes are active. This is because the existence of cognitive agency is typically treated as controversial,⁵⁸ whereas the responsiveness of (at least some of) our attitudes to reasons is not. A richer conception of the cognitive agency will require, on Doyle's view, that we do better to explain what makes inference "*attributable* to the subject herself".⁵⁹ And to explain *this*, we must concede (so Doyle argues) that the agent's attitudes are not only rationally sensitive to other states that figure into her inferences, but that they are also "possessed in virtue of the agent's own assessment of their credentials".⁶⁰ With this point in mind, Doyle offers the following conception of inference: "I bring it about that I believe that *p* on the ground that *q* when I believe that *p* because I take it that this is what I should believe". In this way, "my sense of how things should be with my beliefs is explanatory of my believing as I do".⁶¹ Notice that these taking-attitudes are *de se*: it is by taking it that *I ought to believe q* on the basis of *p* that I come to believe *q*. The robust view looms.

⁵⁷ Joseph Raz, "The Active and the Passive," *Aristotelian Society Supplementary Volume*, 71, 1 (1997): 211-228. See also Scanlon's talk of "judgement-sensitive attitudes" in his *What We Owe to Each Other* (Harvard University Press 1998).

⁵⁸ See, e.g., Setiya, "Epistemic Agency: Some Doubts" and Kornblith, "Epistemic Agency" in *Performance Epistemology: Foundations and Applications*, ed. Miguel Ángel Fernández (Oxford University Press 2016), 167-182.

⁵⁹ Doyle, *Four Essays on Self-Knowledge*, 106.

⁶⁰ Doyle, *Four Essays on Self-Knowledge*, 106.

⁶¹ Doyle, *Four Essays on Self-Knowledge*, 107.

One reply is simply that Doyle has not hereby ruled out a minimalist alternative to the picture he presents. Perhaps, then, the minimalist could say that inference involves cognitive agency to the extent that we are the ones who bring our mental states into cognitive contact via second-tier taking attitudes in inferences. On this picture, my belief that q is responsive to my reason p because I take p to provide epistemic support for q , minimally understood.⁶² If I am rational, we can imagine that my inference proceeds accordingly.

To better see why this minimalist alternative actually makes a good deal of sense, notice that it would be strange, if not outright disquieting, if a reasoner's awareness of p as providing rational support for q could never motivate an inference without self-knowledge. For, as David Owens rhetorically puts the point: "if you already have a non-reflective awareness of the reasons which ought to motivate you, how does the judgement that you ought to be moved by them help to ensure that you are so moved?"⁶³ An answer is that they do not. This means that beliefs like q is what I should believe in light of p look like "an idle wheel in our motivational economy".⁶⁴ If this is right, an agent's sense of how things should be with her beliefs is *not* explanatory of her believing as she does. Rather, her awareness *that p supports q* is explanatory. We can grant that the only reason why one would (rationally) take it that q is what one should believe on the basis of p is that one takes p to support q . And it is surely reasonable to believe that one ought to believe q on the basis of p if one takes p to support q . But then the question resurfaces: what additional motivational role, in coming to believe q , is played by this further (self-)judgement?⁶⁵ As I understand Owens' point, neither my sense of what I ought to believe nor the norms that constrain this sense must figure into the contents of my inference.

⁶² I am not arguing that mental states are *only* sensitive to each other in inferential episodes.

⁶³ Owens, *Rationality Without Freedom*, 18.

⁶⁴ Owens, *Rationality Without Freedom*, 18.

⁶⁵ Owens is originally responding to Burge's account of the importance of "critical reasoning" (see §2).

I am not denying that we sometimes place great stock in ensuring, from a self-conscious perspective, that we believe in accordance with our sense of how we ought to believe. I am only denying that such an aim is *constitutive* of agent-level inference, and that this aim must be represented in the contents of our taking-attitudes. Thus, Boghossian goes wrong in suggesting that the constitutive aim of inference is “figuring out what follows or is supported by other things one believes”⁶⁶, where this is one’s conscious aim *de dicto*.⁶⁷ An alternative aim of inference might just be to find out what is true, rather than what follows from what one believes.⁶⁸ There are even further alternatives having to do with persuasion rather than truth-seeking.⁶⁹

§3.5—Inference and Rational Responsibility

Philosophers sometimes claim that an agent’s *responsibility* for her inferences is a key motivator for TC. As Boghossian argues, while inferences are epistemically evaluable; *we* are also responsible for our inferences.⁷⁰ In Markos Valaris’ words: “if you make a bad inference, we can legitimately criticize *you* as having been hasty, irresponsible, biased, and so on”.⁷¹ Put differently, we are *rationaly responsible* for our inferences, such that we can be criticized for violating norms of inference. TC can purportedly explain this, for it enables us to say that we are rationally responsible for our inferences, and can be criticized accordingly, because *our* takings produce our inferences.

⁶⁶ Boghossian, “What is Inference?”, 5.

⁶⁷ For a similar view see Nishi Shah and David Velleman “Doxastic Deliberation,” *Philosophical Review*, 114 (2005): 497-534.

⁶⁸ This comes from McHugh & Way, “Against the Taking Condition,” 325, despite their skepticism that this is the best way to characterize the aim of inference. Consider also Anna-Sara Malmgren: “the inferring agent needn’t think of herself as settling, or trying to settle, what to believe...She just has to try and do it.” (“On Fundamental Responsibility,” 206).

⁶⁹ See, e.g., Hugo Mercier and Dan Sperber “Why Do Humans Reason? Arguments for an Argumentative Theory,” *Behavioral and Brain Sciences*, 34 (2011): 57-111, and *The Enigma of Reason*. (Harvard University Press, 2017)

⁷⁰ Boghossian, “Inference, Agency and Responsibility,” 113.

⁷¹ Valaris, “What Reasoning Might Be,” 2010.

Crucially, if we can be criticized for inferring badly, it follows that we must be able to think about *ourselves* and, accordingly, our inferences *themselves*. So if Boghossian is right that it is constitutive of being a reasoner that we can be criticized for our inferences, then perhaps we really must always understand our inferences robustly: we must understand them as *being drawn by us*, which means understanding them as taking place in a mind. Alternatively, it may be that although the very process of inference can be second-tier, the second-tier reasoner must always be *in a position* to robustly conceptualise her inferences after the fact, insofar as she is already, at the moment of second-tier inference, a possible target for epistemic criticism down the line. In sum, we are rationally responsible for our inferences, but we could not be so if we were not in a position to know our minds. So, as reasoners, we are in a position to know our own minds.⁷²

Besides the concern that this weakens the agentialist's conclusion beyond any point of serious interest, since it never actually requires the agent to *have* self-knowledge, the main problem with this argument is that it doesn't tell us whether all reasoners really are rationally responsible for their inferences. Thus, consider Anna-Sara Malmgren's view:

Small children can engage in effectual deliberation, at least of the primitive sort...But small children and most animals lack the kind of responsibility that grounds rational evaluability...

...One might suggest that what makes the difference is the capacity for second-tier thought. But it's not entirely clear why, by itself, it would. The ability to (say) make judgments to the effect that such-and-such is a reason to ϕ , that this evidence outweighs that, or that a given claim implies another, might make one *better* at deliberation—in that it makes one better at conforming to the governing norms. But how is that, in turn, supposed to explain fundamental responsibility?⁷³

Malmgren speaks of “effectual deliberation” of a “primitive sort” as a kind of inference that is not yet second-tier. If such reasoning is understood as agent-level but not taking-mediated, then

⁷² This brings to light my qualification about the robust view in §1.

⁷³ Malmgren, “On Fundamental Responsibility,” 207.

proponents of TC must deny its possibility. No matter how they manage this, however, the more important bit in our dialectical context is Malmgren's insistence that even having a capacity for second-tier inference does not entail that one has a "fundamental responsibility" for one's inferences or the mental states that figure into them. Malmgren's use of the term 'fundamental responsibility' essentially refers to the appropriateness of being subject to various deontic norms in the formation and maintaining of one's attitudes, such that one is criticisable when one violates them. So Malmgren's view seems to be a relevant alternative to Boghossian's. They adopt a seemingly similar conception of rational responsibility, but Malmgren denies that such responsibility follows directly from our capacity to infer.

According to Malmgren, "[w]hat makes the difference" to being rationally responsible for our inferences or their products "are our *introspective* and *self-reflective* capacities" (2019, p. 207). This may be so, but it can be no help for proponents of the robust view. The reason is straightforward: if rational responsibility requires (a capacity for) self-knowledge, and (a capacity for) self-knowledge is something that must be added to the cognitive repertoire of individuals who are antecedently capable of drawing inferences, then the robust view is false.

Perhaps one could reply by simply *defining* agent-level inferences as those for which we are rationally responsible.⁷⁴ But why should we accept this definition? After all, if Malmgren is right, then it isn't a feature of *inference itself* that explains our rational responsibility for our inferences. Rather, it is an additional set of self-reflective capacities that explains this. In that case, Boghossian is not obviously capturing a desideratum of an account of agent-level inference; rather, he is capturing a desideratum of an account of self-reflective agent-level inference. What Boghossian needs to argue is that the TC should lead us to think that inference,

⁷⁴ Thanks to Ryo Tanaka for pressing this point.

at the agent level, already involves these self-reflective capacities. If he cannot do this, opponents of the robust view are free to deny this ad hoc definitional maneuver.

§4—Conclusion

In this paper I have evaluated six arguments for the robust view, which is the view that agent-level inference requires self-knowledge (or being in a position to have self-knowledge) of one's inferentially embedded mental states. These arguments are implausible. So far, then, agentialists who hope to vindicate deep connections between our agency and our self-knowledge have failed to do so by appealing to TC, even if TC is true.

Notably, this conclusion is not entirely negative, for the arguments of this paper simultaneously confer greater plausibility on Peacocke's suggestion, originally made in response to Burge, that our capacity for higher forms of inference is scaffolded onto a more "primitive" capacity for second-tier inference.⁷⁵ So long as this more primitive capacity captures core desiderata for our account of (agent-level) inference, inference and self-knowledge are not bound together by necessity. No doubt, we are often self-conscious when we infer, or can easily become such, but these do not appear to be necessary conditions for the basic process called inference that so many philosophers have sought to illuminate, this being inference that: is attributable to the agent herself, is distinct from association, has good and bad instances, has intelligible Moore-paradoxical instances, and is active.

⁷⁵ Peacocke, "Our Entitlement to Self-Knowledge: Entitlement, Self-Knowledge and Conceptual Redeployment," 129.