



University  
of Glasgow

Kelp, C. (2021) Theory of inquiry. *Philosophy and Phenomenological Research*, 103(2), pp. 359-384.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

This is the peer reviewed version of the following article:

Kelp, C. (2021) Theory of inquiry. *Philosophy and Phenomenological Research*, 103(2), pp. 359-384, which has been published in final form at <http://dx.doi.org/10.1111/phpr.12719>

This article may be used for non-commercial purposes in accordance with [Wiley Terms and Conditions for Self-Archiving](#).

<http://eprints.gla.ac.uk/219477/>

Deposited on: 1 July 2020

Enlighten – Research publications by members of the University of Glasgow  
<http://eprints.gla.ac.uk>

# Theory of Inquiry

## Introduction

*Epistemology is the theory of inquiry.* This is the central methodological idea I will explore in this paper. I'd like to add this: I will not provide a direct defence of this idea. Rather, I will take it for granted and see what it can do for us. More specifically, I will look at four of the most fundamental questions in epistemology, concerning (i) the *nature* and (ii) *sources* of knowledge as well as (iii) its *value* and (iv) the *extent* to which we possess knowledge. My ambition here is to use the central methodological idea and show that it allows us to give new and promising answers to all four questions in an attractively systematic manner. If my arguments successful, they will confirm the central methodological idea, thus providing at least indirect support for it.

One interesting property of inquiry is that it is a type of activity with an aim, and indeed, a constitutive aim. When we are inquiring into a given question, we are trying to settle this question. The idea that inquiry is an activity with a constitutive aim is of central importance for the present project. In fact, it holds the key to the answers to our questions concerning the nature, sources, value and extent of knowledge. To show how is the task for the remainder of this paper.

More specifically here is how I will proceed. Section 1 argues that knowledge is the aim of inquiry. Section 2 uses the idea that inquiry is an activity with a constitutive aim and that knowledge is its constitutive aim to develop a non-reductive 'network' analysis relating inquiry, knowledge and belief. Section 3 develops a couple of substantive constraints on sources of knowledge and deals with the Gettier problem. Section 4 uses the idea that activities with constitutive aims constitute value domains and the idea that knowledge is the constitutive aim of inquiry to argue that knowledge has final value in the epistemic domain and uses this idea to address a range of value problems in epistemology. Section 5 turns to the problem of scepticism. I argue that the central methodological idea serves to motivate denying the principle that knowledge transmits across competent deduction, which is widely believed to be at the heart of the problem of scepticism. In addition, I argue that we can have basic knowledge of the denials of sceptical hypothesis through an ability to recognise that certain possibilities could not easily obtain.

## 1. The Aim of Inquiry

Inquiry is an activity with an aim. Here is one way to characterise this aim, which I expect to be widely agreed: inquiry into whether  $p$  aims at settling the question whether  $p$ ; alternatively: inquiry into whether  $p$  aims at properly closing for oneself the question whether  $p$  in the affirmative/negative. At the same time, there is a debate in epistemology over the aim of inquiry. More specifically, there are three prominent rival views that take the aim of inquiry to be knowledge (K-AIM), true belief (T-AIM) and justified belief (J-AIM), respectively.<sup>1</sup> These rival views can be seen as providing substantive accounts of what it takes to settle a question. Which view is correct? In what follows I will first present a number of arguments that there is reason to favour K-AIM and then discuss some objections to K-AIM.

### 1.1 *Knowledge is the Aim of Inquiry*

Consider any activity with an aim (AA). We may place rewards on attaining the aims of AAs. One will be entitled to any reward so placed if and only if one attains the relevant aim. Say, for instance, I offer a hundred dollars for the first person to beat Deep Blue. In this case, you will be entitled to the sum if and only if you are the first to beat Deep Blue. This gives us the following principle:

#### Reward Entitlement

If there is a reward on attaining the aim,  $a$ , of a given AA, then you are entitled to this reward if and only if you do attain  $a$ .

Now consider the following case:

#### The Reward

The T conjecture is a great unresolved issue in mathematics, which is why there is a big prize on settling the question whether (THEOREM =) T does indeed hold. One mathematician works on a proof of T and submits the results to the prize committee at  $t_1$ . Unfortunately, despite the fact that this mathematician proceeded with extreme care and that a number of experts who have read the proof confirmed its soundness before submission, one peer reviewer of the prize committee discovers a subtle error in the proof, perhaps due to a fault in the computer that contributed to generating it. Another mathematician also

---

<sup>1</sup> Champions of K-AIM include e.g. (Kelp 2014, Forthcoming, Millar 2010, Williamson 2000). T-AIM has been defended e.g. in (Lynch 2005, Kvanvig 2003) and J-Aim e.g. in (Davidson 2005, Feldman 2002, Rorty 1995).

works on a proof and submits the results at t2. There are no problems with her proof. In fact, it is perfectly sound.

Who should get the reward for settling that THEOREM? It is intuitively clear that the answer is: the second mathematician. Since the first mathematician's work contained an error, he did not settle that THEOREM. Crucially, between the three live candidates in the literature, only K-AIM can secure this result. To see why, note that the case of the first mathematician is a textbook example of a Gettier case: he has a belief that THEOREM, which is both justified and true, but falls short of knowledge. Since he has a justified true belief, both T-AIM and J-AIM predict that he has attained the aim of inquiry into whether THEOREM holds. And since he was the first to do so, by Reward Entitlement, these views are firmly committed to the claim that he is entitled to the big prize. But, of course, that's the wrong result. In contrast, K-AIM predicts the right result. Since his belief is gettiered, he does not know. So, by K-AIM, he has not attained the aim of the inquiry into whether THEOREM and so isn't entitled to the prize. The second mathematician, in contrast, does come to know the answer. In conjunction with Reward Entitlement, K-AIM predicts, correctly, that she should get the prize. By the same token, there is reason to believe that K-AIM is correct way of unpacking the aim of inquiry, rather than T-AIM or J-AIM.<sup>2</sup>

Let's move on to the second argument. One can be *committed* to attaining successes in AAs. One way in which one can be released from any commitment one may have towards attaining the aim of an AA is by actually pulling off the feat, i.e. by attaining the relevant success. This gives us the following first crucial thesis about AAs:

#### Commitment Release

If, at t, one attains the aim of a given AA, then, at t, one is released from all commitments towards attaining this aim.

Now consider the following case:

#### The Hire

You are a private detective. I have contracted you to find out whether (MURDER =) someone in my family is the murderer of my wife. The first suspect you investigate is my uncle who, let us suppose, has a particularly

---

<sup>2</sup> (Kelp Forthcoming)

strong motive for the deed. Fortunately for you, my uncle credibly admits to having committed the crime upon questioning and even signs a confession in writing. On the basis of this evidence, you come to believe that MURDER, pack your bags, including the confession, get on the next flight to the Caribbean where you intend to take a holiday for the rest of the month. Meanwhile, it becomes widely known that my uncle's claim to have murdered my wife is false. In fact, he has a watertight alibi for the time of the deed and was protecting the perpetrator. At the same time, MURDER is true: someone in my family did indeed murder my wife. It is just that it wasn't my uncle but, say, my cousin. You are currently sipping cocktails in the sun and are entirely unaware of the news about my uncle's confession.

Are you released from the contractual commitment you incurred when I hired you, i.e. to find out whether MURDER? Again, the answer is no. As far as your contract is concerned, what you ought to be doing is work on the case rather than sit on the beach in the Caribbean. To see this, put yourself in my shoes (i.e. in the shoes of your employer). Suppose I just found out about your situation. While I might concede that you have an excuse for no longer working on the case, I could rightly insist that you go back to work and fulfil your contract. Crucially, there is no need to negotiate a new contract with you. The old contract is still binding. None of this would make any sense if you had been released from your contractual commitment. Again, between the three rival views, only K-AIM accommodate this answer. After all, The Hire is also a standard Gettier case: your belief that MURDER is both justified and true. Since you have a justified true belief, both T-AIM and J-AIM predict that you have attained the aim of inquiry into whether MURDER. By Commitment Release, these views are firmly committed to the claim that you are released from your contractual commitment. But, of course, that's the wrong result. In contrast, K-AIM predicts the right result. Since your belief is gettiered, you don't know. So, by K-AIM, you have not attained the aim of the inquiry into whether MURDER and aren't released from your commitment.<sup>3</sup>

Here, then, is the last argument. Often AAs admit of varying degrees of *progress* towards attaining their aims. Progress here is a function of times and distances from attaining the aims. One has made progress on a given AA between  $t_1$  and  $t_2$  if and only if, at  $t_2$ , one is closer to attaining the aim of this AA than at  $t_1$ . Suppose you are about to cross the finish line of the marathon you are currently running. You have

---

<sup>3</sup> (Kelp 2014, Forthcoming).

made progress compared to when you crossed the start line because you are now closer to attaining its aim. It is not hard to see that this definition of progress yields the following thesis:

#### Progress

If, at  $t_2$ , one has not attained the aim of a given AA and if one makes progress towards attaining its aim between  $t_1$  and  $t_2$ , then one has not attained its aim at  $t_1$  either.

With Progress in play, let's return to inquiry once more. In particular, consider the following case:

#### The Insight

The local prison at which you are a guard has received a tip that (DRUGS =) an inmate is in possession of illegal drugs. You are charged with finding out whether this is correct. The first inmate you investigate has a history of drug dealing and abuse. You inspect his cell and find a little bag of white powder under his bed. Moreover, upon questioning, the inmate credibly admits to being in possession of drugs. On the basis of this evidence, at  $t_1$ , you come to believe that DRUGS. At  $t_2$ , however, you discover that the bag does actually not contain drugs and that the confession was false. Perhaps the inmate thought that someone else had planted drugs on him and was confessing to minimise the impending punishment. So, you go back and resume your inquiry. At the same time, DRUGS is true: some inmate is indeed in possession of illegal drugs. It is just that it isn't the one you investigated.

Now here is what I take to be an overwhelmingly plausible claim about The Insight: *you make progress on your inquiry into whether DRUGS between  $t_1$  and  $t_2$*  (henceforth also 'the crucial claim'). By way of support for the crucial claim notice, first, that, at  $t_2$ , you have discovered that a certain piece of evidence that, at  $t_1$ , appears to support the target proposition is misleading. Second, at  $t_2$ , you can tick one person off the list of suspects that you were not able to tick off at  $t_1$ .

The final reason why K-AIM is preferable to the competition is that only K-AIM is compatible with the crucial claim. Here is why: According to the crucial claim, you made progress on the inquiry into whether DRUGS between  $t_1$  and  $t_2$ . By Progress, it follows that at  $t_1$  you did not attain the aim of your inquiry into this question. However, at  $t_1$  you have a justified true belief that DRUGS. According to both T-AIM

and J-AIM, then, you did attain the aim of your inquiry into whether some inmate is holding at t1. The competition is thus unable to accommodate the crucial claim. In contrast, K-AIM encounters no difficulties here. After all, at t1, your belief that DRUGS is gettiered and so does not qualify as knowledge. According to K-AIM, you do not attain the aim of your inquiry into whether DRUGS at t1. As a result, K-AIM is entirely compatible with the crucial claim.<sup>4</sup>

### *1.2 Objections and Replies*

But can K-AIM be correct? In what follows, I will look at a couple of prominent objections to this view. The first aims to show that the aim of inquiry cannot be thought to involve truth. Since knowledge is factive, the argument also targets K-AIM. While versions of it have been given by a number of different champions of this view, it is particularly clearly stated in the following passage by Davidson:

[T]ruths do not come with a “mark”, like the date in the corner of some photograph, which distinguishes them from falsehoods. The best we can do is test, experiment, compare and keep an open mind. But no matter how long and well we and coming generations keep at it, we and they will be left with fallible beliefs. We know many things, and will learn more; what we will never know for certain is which of the things we believe are true. Since it is neither visible as a target nor recognizable when achieved, there is no point in calling truth a goal. (Davidson 2005, 6)

There are actually two arguments in this short passage. Both lay down a condition on aims. The first argument holds that aims must be visible as targets. The second, in contrast, claims that aims must be recognisable to have been achieved when they have been achieved. Crucially, the argument continues, truth doesn't have either property and so cannot be an aim. *A fortiori*, it cannot be the aim of inquiry.

While I could discuss the arguments individually, I won't. One reason for this is that if I manage to successfully argue that these arguments fail, its champions might just try to replace visibility or recognisability with a different condition. Rather, what I'll do is focus on what I take to be the underlying structure of the argument. Note that both conditions first place an epistemic accessibility requirement on aims (visibility, recognisability) and then move on to the claim that truth does not satisfy this accessibility requirement.

---

<sup>4</sup> (Kelp 2014, Forthcoming)

### The Unaccessibility Argument

1. For something to be an aim, it must satisfy some epistemic accessibility requirement.
2. Truth does not satisfy this epistemic accessibility requirement.
3. Hence, truth cannot be an aim.

This argument fails. The reason for this is that there is no way of unpacking the epistemic accessibility requirement that makes both (1) and (2) true. To see this, let's first return to Davidson once more: the above passage suggests that, according to Davidson, in order to satisfy the epistemic accessibility requirement, you must *know for certain* that you have attained the relevant success. This promises to make (2) true, at least on the plausible assumption that there is fairly little in general that can be known for certain.

At the same time, (1) comes out clearly false. Completing a marathon can clearly be an aim. At the same time, it's not possible to know for certain that you have done so when you have done so. After all, despite your best efforts to ensure otherwise, it may always be the case that something went wrong with the equipment used to measure the distance with the result that you stopped running prematurely and so failed to cover the distance. Note also that this doesn't even start mentioning the more extravagant types of error possibility (Cartesian demons, etc.) that are standardly taken to show that it is very hard to know anything for certain.

On the other hand, we can opt for a weaker account of the epistemic accessibility requirement in order to give (1) a better chance of coming out true. For instance, we might unpack the epistemic accessibility requirement in terms of *knowledge* that you have attained the aim. To see why this promises to make (1) true, note that, in the marathon case, you can plausibly know that you have run a marathon, for instance in virtue of knowing that if you have crossed the finish line, you have run a marathon and coming to know that you have crossed the finish line. (I hasten to add that you will not come to know that you have run a marathon if your epistemic environment is insufficiently hospitable in that the equipment used to measure the distance was malfunctioning, Cartesian demons were meddling, etc. Crucially, however, the fact that you may fail to know (e.g. when the environment is insufficiently hospitable) does not entail that you don't frequently come to know (e.g. when it is not).) In this way, the weaker account of the epistemic accessibility requirement gives us a better chance of making (1) come out true.



Unfortunately, now there is excellent reason to think that (2) is false. After all, it is hard to deny that we can come to know the truth. Suppose I wonder whether Fermat's last theorem is true. It is hard to deny that I can come to know that it is true, for instance, by asking a colleague in the mathematics department. Again, this is not to say that I might also fail to acquire this knowledge, for instance, if I am in an epistemically inhospitable environment. But now that I am not, there is no in principle obstacle for me to come to know what I want to find out. As a result, the weaker account of the epistemic accessibility relation will make (2) come out false.

In general, the trouble with the unaccessibility argument is that it is just not clear why it should be any harder to satisfy any given accessibility requirement for truth than for any other aim. That's why ways of unpacking what it takes to recognise something that make (1) true, will make (2) false and vice versa.

Let's move on to the second objection, then, which can be found in the following passage from Marian David:

To say that believing  $p$  is justified or unjustified is to evaluate believing  $p$ , in some sense, as a good thing or as bad thing, as having some positive status or some negative status. The suggestion is that this type of evaluation, epistemic evaluation, is most naturally understood along broadly teleological lines, as evaluating beliefs relative to the standard, or goal, of believing truth and avoiding error [...]

Although knowledge is certainly no less desirable than true belief, the knowledge-goal is at a disadvantage here because it does not fit into this picture in any helpful manner [...] In particular, any attempt to understand justification relative to the knowledge-goal would invert the explanatory direction and would make the whole approach circular and entirely unilluminating. After all, knowledge was supposed to be explained in terms of justification and not the other way round. (David 2001, 154)

Here is my reconstruction of David's argument:

1. If epistemic justification is analysed teleologically as derivative from the epistemic aim, then: if knowledge is the epistemic aim, epistemic justification is analysed teleologically as derivative from knowledge.
2. Epistemic justification is analysed teleologically as derivative from the epistemic aim.

3. If knowledge is the epistemic aim, epistemic justification is analysed teleologically as derivative from knowledge. [1,2]
4. If epistemic justification is analysed as derivative from knowledge, then knowledge cannot be analysed in terms of epistemic justification.
5. Knowledge can be analysed in terms of epistemic justification.
6. Knowledge is not the epistemic aim. [4-6]

Of course, this argument aims to show that knowledge is not the epistemic aim. In contrast, K-AIM is a thesis about the aim of inquiry. However, given the central methodological idea, it's plausible that the two aims coincide. So, it's safe to assume that the argument does mean trouble for the present view.

How convincing is this argument? As a first observation, it is clearly valid. So, let's grant this much. Instead, let's take a closer look at the argument's key premises, i.e. (1), (2), (4) and (5). (1) is clearly analytic and (4) is rock-solid: as David also points out, denying (4) would commit us to an unacceptable form of circular analysis. In contrast, the prospects for denying (2) may look somewhat brighter at first glance. However, at the end of the day, there is excellent reason for thinking that it is not an option after all, at least not once we adopt the central methodological idea. Unfortunately, we will have to wait until Section 4 to see exactly why this is.

If (1), (2) and (4) hold, the argument gives us the following result: if knowledge can be analysed in terms of epistemic justification, then knowledge is not the epistemic aim.

Now, David goes on to argue that since knowledge does admit of analysis in terms of epistemic justification, knowledge isn't the epistemic aim. Crucially, however, no one, David included, has ever provided a single argument that knowledge admits of analysis in terms of justified true belief. All we have ever done on this front hope for the best (and, ever since (Gettier 1963), in the face of continued failure). But, of course, in the absence of a viable argument that knowledge admits of analysis in terms of epistemic justification, the argument fails to make a compelling case against the claim that knowledge is not the epistemic aim.

In fact, there is reason to think that the argument can now be turned on its head. Section 1 has argued that knowledge is the aim of inquiry. These arguments now provide excellent reason to believe that knowledge is the epistemic aim, especially once we keep the central methodological idea that epistemology is the theory of inquiry firmly in mind (see also Section 4). By the same token, there is reason to think that the way forward with the conditional that David's argument establishes is not by applying modus ponens, as David would have it, but rather modus tollens. If so, what

David's argument in conjunction with the arguments for K-AIM really amounts to is an argument that knowledge does not admit of analysis in terms of epistemic justification.

If epistemology is the theory of inquiry, then one key task is to figure out what exactly the aim of inquiry amounts to. This section has done just this. More specifically, I have mounted a case for K-AIM, i.e. that knowledge is aim of inquiry. With this important result in place, it is now time to put the central methodological idea to work. This is exactly what I will do in what follows.

## 2. The Nature of Knowledge

This section develops an account of the nature of knowledge. In view of the discussion of David's argument at the end of the last section, it is natural to expect an account that drops the justification condition on knowledge. While this is indeed one way to go, I don't think it is ultimately viable (see Section 3 and Kelp 2018). Accordingly, I'd like to explore a different option here. To see how this option works, note first that David is presupposing a certain model of philosophical analysis for knowledge, which is nicely characterised in the following passage by Peter Strawson:

### The Dismantling Model

[The Dismantling Model] represents it [i.e. philosophical analysis] as a kind of dismantling of a complex structure into simpler elements, a process which terminates only when you reach pieces which cannot be further dismantled. (Strawson 1992, 19)

Dismantling analyses allow us to understand complex phenomena by breaking them down into simpler elements which are independently understood. Crucially, for this to work as envisaged, the simpler elements in the analysans must enjoy *explanatory priority* over the complex structure in the analysandum. Of course, if we still want to hold on to the dismantling model for knowledge, then we will have to drop the justification condition on knowledge.

Fortunately, the dismantling model is not the only game in town. In fact, it is once again Strawson who proposes an alternative. Here goes:

### The Network Model

Let us imagine, instead, the model of an elaborate network, a system, of connected items, concepts, such that the function of each of item, each concept, could, from a philosophical point of view, be properly understood only by

grasping its connections with the others, its place in the system ... (Strawson 1992, 19)

In what follows, I will develop just such a network analysis. More specifically, I will argue that a network analysis is befits a certain type of case. The key proposal I develop below is that knowledge is (involved in) an instance of just this type of case.

### 2.1 *Activities with Constitutive Aims*

We do many things. One thing we do is engage in types of activities with constitutive aims and norms (henceforth also ACANs for short). Chess is an example of an ACAN. Its constitutive aim is to checkmate the opponent. Among the constitutive norms are norms specifying the starting positions of the various types of pieces, what moves are moves in chess, etc.

To say that ACANs have *constitutive* aims and norms is to say that the norms and aims of a given ACAN are *essential* to it. Anything that does not have these aims and norms will not qualify as a token of this ACAN. That's why throwing chess pieces around does not qualify as playing chess, for instance. Since its constitutive aim and norms are essential to a given ACAN, every token of it can be assessed in terms of them. For instance, in chess we can always ask whether the opponent was checkmated, whether the bishop was moved diagonally, etc.<sup>5</sup>

In what follows, I will argue that there is a class of ACANs – exemplified here by a paradigm case of an ACAN, the game of chess – for which the prospects for a dismantling analysis are dim: we may expect that neither the activity, nor its constitutive aims and norms admit of such analysis. Even so, network analyses are still very much promising for members of this class of ACANs. Or so I will argue.

Let's take a closer look at the game of chess, then. Is it really true that chess is an ACAN that does not admit of dismantling analysis? It might be thought that the answer is no. Consider checkmating, for which we can give the following necessary and sufficient conditions:

---

<sup>5</sup> Note that to say that an *activity* has a certain constitutive aim is not to say that an *agent* engaging in that activity has this aim, too. For instance, you can play a game of chess and aim to lose. In this case, the aim of the game of chess, which is to checkmate your opponent, does not coincide with your aim, which is to be checkmated by your opponent. By the same token, to engage in an ACAN cannot require you to adopt the ACAN's aim. What, then, does it take to engage in an ACAN? This is a hard question that I won't be able to address here. For some proposed answers, see e.g. (Kelp and Simion 2018, Williamson 2000).

## Checkmate

One checkmates one's opponent if and only if one places one's opponent's king under inescapable threat of capture.

Granted, it may initially appear as though Checkmate gives a successful dismantling analysis of checkmating. On reflection, however, there is reason to think that this can't be quite right. To see why not, note that the right-hand side of Checkmate itself contains elements that are complex and so afford further analysis. For instance, we might ask what a king is or what capturing amounts to. Of course, there are interesting things to be said here. For instance, to answer the question what a king is we might specify his starting point and the ways in which he moves. Crucially, however, the project of attempting to understand what a king is without making reference to the game of chess is hopeless. We can't understand what a king is without invoking the claim that it is a type of piece in chess. But, of course, chess is again not a simple phenomenon and our analysis cannot stop here. We'll need to offer an analysis of chess. The trouble is that there is no chance of offering a successful analysis of chess without including the claim that its constitutive aim is to checkmate one's opponent. But now we have gone around in a circle. Our analysis of checkmating forced us to make reference to the game of chess and our analysis of chess forced us to make reference to checkmating. What comes to light, then, is that although the prospects for a successful dismantling analysis of checkmating may initially have looked bright, closer scrutiny reveals that they are really quite dim, if not entirely hopeless.

On the upside, there is little reason for thinking that we must now abandon the attempt to understand the activity and its constitutive aim altogether. After all, a network analysis might still be made to work. In fact, a network analysis seems particularly well suited here. To properly understand chess, we must understand its constitutive norms and aims, including that the constitutive aim of chess is to checkmate one's opponent, the constitutive norm that castling is a type of move in chess, etc. To properly understand these constitutive norms and aims, we must at the very least understand that they are constitutive of chess. For instance, there is no way to properly understand what checkmating amounts to, unless we understand that it has something to do with chess. To properly understand the game of chess and its constitutive aims and norms, we are thus bound to move around in a circle. But, again, this is entirely compatible with a network analysis. And, more importantly, the resulting analysis is no less illuminating for all that.

## 2.2 Inquiry, Knowledge and Belief

In what follows, I will develop a network analysis for knowledge. The key idea is that knowledge is defined in terms of its place in an ACAN that falls into the class of ACANs, exemplified by the game of chess above, which lends itself to network but not dismantling analysis. Let's look at the central ACAN then. Unsurprisingly, it is inquiry into whether  $p$ . Recall that the aim of inquiry into whether  $p$  can be characterised in a lightweight fashion as settling the question whether  $p$ , or, alternatively, as properly closing that question in the affirmative/negative for oneself. Crucially, inquiry is not only an activity with an aim but also an ACAN. The aim of settling questions is not only the aim of inquiry but it is also its *constitutive* aim. Any activity that does not have this aim is ipso facto not the activity of inquiry.

What about the constitutive norms of inquiry? One prima facie plausible candidate is that in inquiry one must form or work toward positioning oneself to form beliefs via epistemic abilities.<sup>6</sup> That's why going to the local brainwashing service and having a certain belief installed will not count as inquiring into the corresponding question. More importantly for present purposes, I want to suggest that there are constitutive norms specifying kinds of move in inquiry, just as in chess there are constitutive norms specifying kinds of move in chess. In particular, one important kind of move is to close inquiry in the affirmative or negative.

Now the suggestion is that inquiry is an ACAN that falls in the same class as chess in the sense that it lends itself to network but not dismantling analysis. In the case of chess, we already have a pretty good idea of what the network of connected phenomena looks like. Since the same is not true of the network in terms of which knowledge is analysed, here is the rough idea: the network relates knowledge, belief and inquiry in the sense that knowledge is the constitutive aim of inquiry and belief is the kind of move in inquiry that closes inquiry for one in the affirmative or negative.<sup>7</sup>

With the rough idea in play, let's look at some of the details. Recall that I said that inquiry into whether  $p$  aims at settling the question whether  $p$ ; or, alternatively, it aims at properly closing the question whether  $p$  for oneself in the

---

<sup>6</sup> I will return to epistemic abilities in Section 3. For detailed discussion see (Kelp 2016, 2017, 2018, 2019a). While I think this is roughly right, alternatives are conceivable. Evidentialists, for instance, might say that inquiry essentially involves gathering evidence (and that's why using the brainwashing service in the case below doesn't qualify as inquiring).

<sup>7</sup> While I will describe enough of the network for my purpose of illuminating the nature of knowledge in terms of its place in it, I will not describe it in its entirety. Most importantly, suspension of judgement will have a place in the network (as another type of move). However, since suspension of judgement is of little relevance for understanding the nature of knowledge, I will set it (as well as any other phenomena that may belong in the network) aside here.

affirmative/negative. The way in which knowledge enters the picture, I suggest, is that it specifies what it takes to attain the constitutive aim of inquiry (as per Section 1) in much the same way as the right-hand side of Checkmate specifies what it takes to attain the constitutive aim of chess. In particular, I want to suggest the following necessary and sufficient conditions for the question settling aim of inquiry:

#### Question Settling

Inquiry into the question whether  $p$  is settled/properly closed for one in the affirmative/negative if and only if one knows that  $p$ /not- $p$ .

I already mentioned that belief is a kind of move in inquiry and, more specifically, that it is the type of move that closes inquiry into whether  $p$  for one in the affirmative or negative:

#### Belief

One believes that  $p$ /not- $p$  if and only if inquiry into whether  $p$  is closed for one in the affirmative/negative.<sup>8</sup>

With these points in play, the network is up and running. Note that, just as in the case of chess earlier, we find necessary and sufficient conditions for some of the relevant phenomena. And, again, this raises the question as to whether we might not hope for relevant reductive analyses after all.

Crucially, just as in the case of chess, the answer is no. To see this, note that the right-hand side of Question Settling is itself complex and would therefore require further analysis. Moreover, we can say interesting things about it, too, including that knowledge requires epistemic justification and belief.

However, first, we have already seen in the discussion of David's argument above that if (i) knowledge is the aim of inquiry and (ii) epistemic justification is analysed teleologically, then the epistemic justification will have to be analysed in terms of knowledge. Since there is excellent reason to think that (i) holds and I am granting (ii) (see also Section 4), the prospects for a dismantling analysis of knowledge in terms of epistemic justification are dim. After all, a dismantling analysis will lead us from knowledge to the teleological epistemic justification on knowledge and from there back to knowledge. Finally, since I also grant that it would be a mistake to deny that knowledge requires epistemic justification (see also Section 3), the result that we

---

<sup>8</sup> For further support for this view see (Friedman 2019, Harman 1986, Hieronymy 2009).

get is that our analysis of knowledge is going to be circular and a dismantling analysis won't be forthcoming.

Moreover, second, by Belief, the necessary and sufficient conditions for belief feature inquiry. But, of course, inquiry is again not a simple phenomenon and our analysis cannot stop here. We'll need to offer an analysis of inquiry. The trouble is that there is no chance to offer a successful analysis of inquiry without including the claim that its constitutive aim is to settle the question at hand. As a result, we are once again tied up in a circle. Our analysis is thus circular on yet another count, which makes the prospects for dismantling even worse.

Yet, again, there is an upside. There is little reason for thinking that we must now abandon the attempt to understand the nature of knowledge altogether. After all, a network analysis might still be made to work. In fact, a network analysis seems particularly well suited here. To properly understand inquiry, we must understand its constitutive aim and norms. And to properly understand its constitutive aim and norms, we must at the very least understand that they are constitutive of inquiry. Just as in the case of chess earlier, we are thus bound to move around in a circle. But again this is entirely compatible with a network analysis. And, more importantly, the resulting analysis is no less illuminating for all that.

### *2.3 Benefits of the Network Analysis*

It may be worth noting that there are a number of theoretical benefits that the network analysis brings in its train, both for the theory of knowledge and for the theory of belief. Here goes.

First and foremost, the network analysis offers an alternative way of understanding the nature of knowledge, once we abandon the ambition of giving a dismantling analysis of knowledge, that is. This means that we will be able to reap the benefits of the central methodological idea whilst not being at a complete loss when it comes to the important task of understanding the nature of knowledge.

Second, the network analysis places the ACAN of inquiry and hence the phenomena of inquiry, knowledge and belief which it connects into a broader class of ACANs (alongside chess) that does not admit of dismantling analysis. In this way, the proposal promises to deliver an important diagnostic insight into why we shouldn't expect knowledge to admit of dismantling analysis in the first place.



Third, the following highly attractive thesis about belief drops right out of the network analysis, to wit, that belief is constitutively governed by a knowledge norm.<sup>9</sup> Here is how. Recall first that, by Question Settling and Belief, knowledge that *p* (not-*p*) is proper belief that *p* (not-*p*). As a result, belief is governed by a knowledge norm. And since, on the present network analysis, the relations between inquiry, knowledge and belief hold as a matter of constitutive fact, this knowledge norm is a constitutive norm of belief.<sup>10</sup>

Fourth, the network analysis allows us to avoid some important downsides that this thesis that belief is constitutively governed by a knowledge norm might be thought to have. For instance, it might be thought that it forecloses a dismantling analysis of knowledge in terms of belief. While this is of course true, it is no longer worrisome. After all, the network analysis also explains why we shouldn't expect a dismantling analysis of knowledge to be viable in the first place. It might also be thought that the thesis about belief commits us to the implausible claim that belief can fully be analysed in terms of knowledge or, at the very least, that knowledge enjoys explanatory priority over belief. Again, given the network analysis, this is incorrect. This is because both are analysed by their place in the network. Since the network places them on an explanatory par, knowledge does not enjoy priority over belief on this front. Since the network also features inquiry, belief cannot be fully analysed in terms of knowledge.

### 3. Sources of Knowledge

The label 'sources of knowledge' is often used to refer to the theory of specific sources of knowledge, such as perception, testimony and so on. This is not what I mean by 'sources of knowledge'. Rather, what I will be focusing on here are general constraints on sources of knowledge. Let me explain.

One interesting fact about ACANs that feature normative properties in their aims is that they often impose further substantive constraints on what we may call sources of success; or, to be more precise, on the means of attaining the ACAN's constitutive aim and on the environment. Consider, for instance, *fairly* winning a chess

---

<sup>9</sup> While I do take this to be a particularly neat consequence of the account, I won't stop to defend it here. Fortunately, this won't be necessary either since what I take to be a compelling case for this claim has already been made in the literature (e.g. Adler 2002, Huemer 2007, McHugh 2011, Simion, Kelp and Ghijsen 2016, Williamson 2000).

<sup>10</sup> See (Simion, Kelp and Ghijsen 2016) for a more detailed account of what means to say that belief is governed by a knowledge norm (in terms of an evaluative norm specifying conditions for Geach-style attributive goodness).

game. To win a chess game fairly it won't be enough that you checkmate your opponent. There are constraints on the means for attaining the aim. For instance, you may not do so with the help of a chess computer (while your opponent doesn't have such help). In fact, it is impossible for you to win fairly in this way. Similarly, there are constraints on the environment. For instance, your competition must not be hampered by performance diminishing drugs (while you aren't). Again, it's impossible for you to win fairly in this situation.

Since the aim of inquiry features a normative property, the question arises as to whether there are substantive constraints on sources of success here as well; or, to be more precise, on the means of attaining its aim and on the environment. There is reason to think that the answer is yes. To see this, note that certain means are clearly unfit for settling certain questions. Think, for instance, of going by the results of the toss of a fair coin, which will give you the true answer with a probability of 0.5. Just as the use of a chess computer will make it impossible for you to win a game of chess fairly, so the use of a fair coin toss will make it impossible for you to settle the question at hand. And, likewise, there are environments that make it impossible to settle certain questions. The perhaps most well-known examples here are various forms of radical sceptical scenarios in which virtually nothing is as it seems. Just as playing a game of chess in which your opponent is hampered by performance diminishing drugs (while you aren't) will make it impossible for you to win fairly, so being in a radical sceptical scenario will make it impossible for you to settle certain questions.

Given that one knows if and only if one attains the aim of inquiry, these considerations provide excellent reason to believe that knowledge features substantive constraints on its sources, i.e. the means of attaining knowledge and the environment. The question remains whether more can be said about these substantive constraints. In what follows, I will briefly sketch my own preferred answer. Again, my strategy is to use theoretical machinery from the general theory of ACANs. In particular, my answers will appeal to accounts of (i) abilities to attain the aims of ACANs, (ii) their exercises and (iii) competent ACAN moves. A word of caution: the account I am going to give may seem rather detailed, perhaps gratuitously so. However, there is a reason for this. I will need this level of detail in the discussion of closure and scepticism in Section 5.

### 3.1 Abilities

Practitioners of ACANs may have the *ability* to attain the aim of a given ACAN. Let's call abilities to attain aims of ACANs 'ACAN abilities'. Now consider the following account of ACAN abilities:

### ACAN Ability

One has an ability to attain the aim for a range,  $R$ , of ACANs and relative to conditions,  $C$ , if and only if one has a way of move production,  $W$ , such that, for any  $S \in R$ , using  $W$  in  $C$  disposes one to attain the aim of  $S$ .<sup>11</sup>

I want to suggest that exercises of ACAN abilities are uses of ways of move production involved in ACAN abilities. Or, more precisely,

### ACAN Exercise

One exercises an ability,  $A$ , to attain the aim for a range,  $R$ , of ACANs and relative to conditions,  $C$ , if and only if one has  $A$  and produces a move via the way of move production at issue in  $A$ .

Finally, here is an account of a competent move in an ACAN:

### Competent ACAN Moves

A move in a given ACAN,  $S$ , is competent if and only if it is produced by an exercise of an ACAN ability to attain the aim for a range,  $R$ , of ACANs and relative to conditions,  $C$ , such that  $S \in R$ .

The question that immediately arises is why we should think these accounts are correct. While I provide answers elsewhere (Kelp 2016, 2017, 2018, 2019a), unfortunately, due to limitations of space, I won't be able to do so here. What I want to do instead is return to the question of how to unpack the substantive conditions on the sources of attaining the aim of inquiry/knowledge that I asked above. With ACAN abilities, their exercises and competent ACAN moves in play, I am now in a position to answer this question. I want to unpack both substantive constraints in terms of the notion of an ACAN ability. More specifically, I want to suggest that the constraints on the means for attaining the aim of inquiry/knowledge is unpacked in terms of a competence condition and the constraint on the environment in terms of the satisfaction of the  $C$  of the ability that produced the belief. In other words,

### Conditions on Knowledge

---

<sup>11</sup> See (Kelp 2016, 2017, 2018, 2019a) for a more detailed account on ACAN Abilities (as well as ACAN Exercise and Competent ACAN Move below).

One knows that  $p$  only if (K1  $\Rightarrow$ ) one's belief that  $p$  is competent and (K2  $\Rightarrow$ ) the C of the ability by means of which one's belief that  $p$  is formed are satisfied.

### *3.2 Why These Conditions?*

Now, even if we are willing to go along with the idea that for ACANs with constitutive aims that feature normative properties there may be substantive constraints on both the means of attaining these aims and the environment; and even if we are willing to grant that inquiry is an ACAN with a constitutive aim featuring a normative property, why should we think that this specific way of precisifying these constraints is correct?

While I don't have anything conclusive to offer in response to this question, there are some considerations that provide relevant evidence. The first is that the above precisification accommodates the motivations for the substantive constraints on the means and environment for attaining the aim of inquiry. For instance, we can now explain why you cannot attain the aim of inquiry via a coin toss on whether  $p$ . A coin toss on whether  $p$  simply couldn't be an ability to know. This in turn is because a coin toss on whether  $p$  simply doesn't dispose one to acquire knowledge. But, of course, if a coin toss on whether  $p$  couldn't be an ability to know, beliefs acquired by a coin toss simply couldn't be competent. Given K1 of Conditions on Knowledge (and given K-AIM), it follows that it's impossible to attain the aim of inquiry via a coin toss on whether  $p$ . Similarly, we can explain why it's impossible to attain the aim of inquiry when radically deceived. Even if the beliefs one forms are competent, the C of the ability to know that produced these beliefs could not be satisfied. This, in turn, is because the disposition to know at issue in this ability simply could not extend to the C of radical deception. By K2 of Conditions on Knowledge (and given K-AIM), it follows that it's impossible to attain the aim of inquiry when radically deceived.

The second reason for thinking that the precisification is correct is that it serves to address one of the central problems in the literature on the nature of knowledge, i.e. the Gettier problem. Gettier cases are cases in which agents lack knowledge because they are in unsuitable epistemic environments. More specifically, they are in environments such that their ways of belief formation do not qualify as abilities to know relative to these environments. At the same time, their ways of belief formation qualify as abilities to know relative to other more suitable environments and so they qualify as abilities to know the target propositions. This means that agents in Gettier cases, whilst holding competent beliefs, are in unsuitable C for the abilities to know that produced their beliefs. By K2 of Conditions on Knowledge, their beliefs will fall short of knowledge.

By way of illustration, consider first the following famous Gettier cases:

### Stopped Clock

Having come down the stairs, you look at the grandfather clock in the hallway, see that it reads 8:22 and on that basis comes to believe that it is 8:22. The clock has an outstanding track-record of functioning properly and you have no reason to think that it is currently not accurate. Your belief is true. It is in fact 8:22. Unbeknownst to you, however, the clock has stopped exactly twelve hours ago.

### Fake Barns

You are driving through the countryside and are currently looking out of the window of his car. You see what appears to be a barn in the field and form a perceptual belief that there is a barn in the field. Unbeknownst to you, you are looking at one of the few real barns in an area peppered with barn façades that are so cleverly constructed as to be indistinguishable from real barns from your position on the road. (Goldman 1976)

In Stopped Clock, you don't know that it's 8:22. Here is why. The fact that you are taking a reading from a stopped clock means that you are not in conditions C relative to which you have the ability to acquire knowledge about the time by taking a reading from the clock. At the same time, the conditions you find yourself in do not prevent you from forming your belief via the way that constitutes an ability to know relative to some (but different) conditions. This means that you are in unsuitable C for the ability that produced your belief about the time. By, K2 of Conditions on Knowledge, your belief that it's 8:22 falls short of knowledge.

In Fake Barns, you don't know that there is a barn in the field. Here is why. The fact that you are in a part of the country in which fake barns predominate means that you are not in conditions C relative to which you have the ability to recognise barns by looking. At the same time, the conditions you find yourself in do not prevent you from forming your belief via the way that constitutes an ability to know relative to some (but different) conditions. This means that you are in unsuitable C for your ability to recognise barns. By K2 of Conditions on Knowledge, your belief that there is a barn in the field falls short of knowledge (Kelp 2017, 2018).

The third reason for accepting this way of making the conditions on knowledge more precise is that it contributes to making an attractive solution to the problem of scepticism available. I will get back to this point in Section 5.

Let's move on to the fourth and last reason then. To see this one, recall that we found it only natural to identify the constraint on the means with the independently plausible justification condition on knowledge. What the above precisification of the constraint on means does, then, is to deliver the following precise account of the justification condition on knowledge:

#### Justified Belief

One's belief that  $p$  is justified if and only if it is formed by an exercise of an ability to know propositions in range  $R$  and relative to conditions  $C$  such that  $p \in R$ .

In a nutshell, the fourth reason for thinking that the above precisification of the constraints on the aim of inquiry is correct is that (modulo the natural identification of the constraint on means with the justification condition on knowledge) it delivers Justified Belief. I have defended Justified Belief elsewhere in considerable detail (Kelp 2016, 2017, 2018). Unfortunately, due to limitations of space, I won't be able to rehearse my arguments here. Still, it will not come as a surprise that I take the fact that we get Justified Belief to be an attractive result.

### 4. The Value of Knowledge

This section uses the central methodological idea to develop a novel account of epistemic value and shows how this account can solve a range of difficult 'value problems' in epistemology.

#### 4.1 ACANs and Normative Domains

I'd like to start by introducing the idea of a critical domain of value (henceforth also 'critical domain' for short) (Sosa 2007). Critical domains are associated with goods or values for that domain. Any critical domain has at least one central value. This central value is fundamental to the domain in the sense that it cannot be explained fully in terms of other values in that domain. Another way of putting this point is that the central value is valuable for its own sake (or 'finally valuable'), *relative to the domain*. The qualification is important because values in critical domains may not enjoy any of the more familiar kinds of value, such as moral, prudential or aesthetic value. Crucially, central values organise the evaluations in the relevant domains in the sense that all other domain specific values will be derivative from the central ones: they can be explained fully in terms of the central values.

The reason I am mentioning all this is that I want to suggest that ACANs constitute critical domains. More specifically, I want to put forth the following eminently plausible idea. ACANs' constitutive aims correspond to central values in the domains they constitute. And, of course, this means that these values are for-their-own-sake or final values relative to these domains. Consider, for instance, chess. The constitutive aim of chess is to checkmate one's opponent. As a result, checkmating one's opponent is a central value in the domain constituted by chess and it is a final value relative to this domain.

One of the fundamental divides in the general theory of norms and values is between teleologists and deontologists. Teleologists take values in a certain domain to be fundamental and explain norms in the same domain in terms of values. Deontologists, in contrast, reverse the direction of explanation. They take norms in a domain to be fundamental and explain values in the same domain in terms of norms. Crucially, it is hard to deny that the critical domains that are constituted by ACANs are teleological in structure. That is to say, the central values of critical domains are fundamental and norms are explained in terms of these values. For instance, in chess, the central value of checkmating explains why players should play in certain ways and not in others.

If ACANs constitute critical domains and if inquiry is an ACAN, then inquiry constitutes a critical domain. Since constitutive aims of ACANs correspond to final values in the domains they constitute, it follows that knowledge is a final value in the domain constituted by inquiry. Since the domains constituted by ACANs are teleological in structure, so is the domain constituted by inquiry. For instance, how we ought to proceed in inquiry is explained in terms of knowledge. Finally, given our central methodological idea that epistemology is the theory of inquiry, there is reason to think that knowledge is a central value in the epistemic domain and so enjoys final epistemic value.

Before moving on, I'd like to return to a point that I left hanging in Section 1. In the discussion of David's argument, I noted that denying the teleological account of justification did not carry much promise, at least on the present account. Now, we are in a position to see why. Since the epistemic domain is teleological in structure, epistemic norms are to be explained in terms of epistemic values. Since justification corresponds to an epistemic norm, it is to be explained in terms of epistemic values. So, the teleological account of justification is not up for discussion.

## 4.2 Value Problems

One of the key tasks of epistemology is to explain the distinctive value of knowledge. While it is widely recognised that it won't be hard to explain why knowledge has some value, it's believed to be difficult to explain why knowledge is *distinctively* valuable. The reason for this, in turn, is that in order to successfully do so, one must solve a range of so-called value problems, viz., roughly, problems of explaining why knowledge is more valuable than true belief that falls short of knowledge.

It is widely believed that knowledge is more valuable than mere true belief. Unfortunately, on reflection, it's not entirely clear exactly why this should be the case. This was noted already by Plato in the *Meno* (1956) where Socrates challenges his interlocutor to explain why exactly we should think that knowledge is more valuable than mere true belief. After all, a mere true belief regarding the right way to Larissa would seem to be just as useful as knowledge of the way to Larissa: both will get us where we want to go. But surely, knowledge will be distinctively valuable only if our impression is right and it is more valuable at least than mere true belief. Accordingly, any satisfactory account of the distinctive value of knowledge must explain how this can be. This gives us:

### The Meno Problem

Any satisfactory account of the value of knowledge must explain why knowledge is in some respect more valuable than mere true belief.

Some have claimed that simply meeting The Meno Problem won't be enough to give a satisfactory account of the value of knowledge. While there is a range of value problems on the market, I'd like to look at what arguably is the most demanding version of the value problem, which is due to Duncan Pritchard (Pritchard *et al.* 2010). According to Pritchard, in order to account for the distinctive value of knowledge, one must explain why knowledge is more valuable than the corresponding true belief that falls short of knowledge not just as a matter of degree, but also as a matter of kind. This gives us:

### The Tertiary Value Problem

Any satisfactory account of the value of knowledge must explain why knowledge is in some respect more valuable than true belief that falls short of knowledge not just as a matter of degree but as a matter of kind. (Pritchard *et al.* 2010, 7-8).



Of course, if we can solve The Tertiary Value Problem, then we will automatically also have solved The Meno Problem (as well as the other value problems on the market). But what, then, does it take to solve The Tertiary Value Problem? What does it take to show that knowledge is more valuable than true belief that falls short of knowledge not just as a matter of degree but as a matter of kind? Pritchard's answer is that we will have to show that knowledge is *finally* valuable, i.e. valuable for its own sake (2010, 8).

The trouble, according to Pritchard, is that the prospects of solving the Tertiary Value Problem are rather dim. He takes it that there is only one candidate that might deliver the goods on the market, which is associated with virtue epistemology. In a nutshell, according to virtue epistemology, knowledge is a cognitive achievement. Since achievements in general are finally valuable, so is knowledge (e.g. Greco 2010). According to Pritchard, the trouble is that virtue epistemology is mistaken: knowledge isn't a cognitive achievement. As a result, the only contender on the market that stood any chance of solving The Tertiary Value Problem in the first place turns out to fail. That's not exactly great news.<sup>12</sup>

One of the benefits of the present approach is that it makes a new solution to The Tertiary Value Problem available, one that does not depend on the idea that knowledge is a cognitive achievement in the sense virtue epistemologists take it to be. In fact, it is not hard to guess what this solution looks like: since knowledge is the constitutive aim of inquiry, it is valuable for its own sake, relative to the epistemic domain; or, in other words, it enjoys final epistemic value. It is easy to see that knowledge enjoys a different kind of value than true belief that falls short of knowledge. After all, true belief that falls short of knowledge does not enjoy final epistemic value. It follows that knowledge has a kind of value that true belief that falls short of knowledge does not enjoy: knowledge has final epistemic value because it is the constitutive aim of inquiry. The Tertiary Value Problem is thus solved and the distinctive value of knowledge can be adequately explained.

## 5. The Extent of Knowledge

The last central epistemological problem I will be looking into here is the problem of scepticism, which will take centre stage in this section. To begin with, consider the following famous sceptical argument:

---

<sup>12</sup> In fact, Pritchard himself takes his argument to motivate a form of scepticism about the distinctive value of knowledge. Contrary to what appears to be widely believed, knowledge is not distinctively valuable.

### The Argument from Ignorance

- AI1. You don't know that ( $\sim$ SH  $\Rightarrow$ ) sceptical hypothesis SH (e.g. that you are not a handless BIV) is false.
- AI2. If you don't know that  $\sim$ SH (e.g. that you are not a handless BIV), then you don't know a corresponding ordinary empirical proposition O (e.g. that you have hands) to be true.
- AI3. Hence, you don't know that O (e.g. that you have hands).

Both premises of The Argument from Ignorance are highly plausible. AI1 is intuitively compelling. You don't know that you are not a handless BIV. How could you? After all, everything would seem to you exactly as it would were you to be a normal handed person. AI2 is motivated by the plausible idea that competent deduction is a way of extending our knowledge. This idea, in turn, is thought to be captured by the following principle:

### Transmission of Knowledge

If (C1) one knows that p, (C2) one competently deduces q from p, and (C3) one comes to believe that q based on one's deduction, then (C4) one thereby comes to know that q.<sup>13</sup>

But, of course, in conjunction, AI1 and AI2 commit us to the highly implausible sceptical claim, AI3, that you don't even know that you have hands.

The Argument from Ignorance is perhaps the strongest sceptical argument on the market. If we want to avoid its sceptical conclusion, we will have to deny at least one of its premises. Let's call responses that deny AI1 'robust' and responses that deny AI2 'concessive'.<sup>14</sup> In my view, neither response is robust enough. As I will argue in the remainder of this section, there is reason to deny both of the argument's premises; or, to be more precise, we will do well to deny both AI1 and Transmission of

---

<sup>13</sup> (Williamson 2000, Hawthorne 2005). This principle also sometimes goes by the name of *Closure*. In contrast with the debate on knowledge, the literature on warrant distinguishes sharply between closure principles and transmission principles. Since the above would clearly be categorised as a transmission principle there, I decided to go with the label *Transmission of Knowledge*.

<sup>14</sup> Prominent champions of robust responses to scepticism include G.E. Moore (1939), Duncan Pritchard (2007), Ernest Sosa (1999) and Timothy Williamson (2000). The classical concessive responses are (Dretske 1970, Nozick 1981).

Knowledge which motivates AI2. I will start my discussion with Transmission of Knowledge.

### *5.1 Transmission of Knowledge*

Let's start with an argument against Transmission of Knowledge. In order to achieve this, I will once more focus on inquiry and, more specifically, on inquiry into specific whether questions.

Recall the lightweight characterisation of the aim of inquiry:

#### Question Settling Aim

One's inquiry into whether  $p$  aims at settling the question whether  $p$ .

Another thesis that I take to be no less plausible is that question-begging excludes question settling. This gives us:

#### No Settling

If one's inquiry into whether  $p$  is question-begging, then it doesn't settle the question whether  $p$ .

Now consider the following two cases involving inquiry:

#### Inquiry 1

You are at the zoo. Currently you are standing in front of the zebra enclosure and see a black-and-white striped equine creature inside. Since you can tell a zebra from the way it looks, you come to know that (ZEBRA  $\Rightarrow$ ) the animal is a zebra. You now want to find out whether ( $\sim$ DISGUISED  $\Rightarrow$ ) the animal is not a cleverly disguised mule. To settle this question you take a hair sample from the animal and perform a DNA analysis. The result of the analysis is that the animal is indeed a zebra. From this, you competently deduce and thereupon come to believe that  $\sim$ DISGUISED.

#### Inquiry 2

You are at the zoo. Currently you are standing in front of the zebra enclosure and see a black-and-white striped equine creature inside. Since you can tell a zebra from the way it looks, you come to know that ZEBRA. You now want to find out whether  $\sim$ DISGUISED. To settle this question you exploit the

entailment from ZEBRA to  $\sim$ DISGUISED to competently deduce and thereupon come to believe that  $\sim$ DISGUISED.

There is a clear difference between Inquiry 1 and 2. While there is nothing wrong with your inquiry in Inquiry 1, in Inquiry 2, your inquiry is problematic. What explains this difference? The by far best answer that I can think of is that, in Inquiry 2, as opposed to 1, your inquiry is *question-begging* (more on this below). Reflection on the difference between Inquiry 1 and 2 thus motivates the following thesis:

### Question-Begging

In Inquiry 2, your inquiry into whether  $\sim$ DISGUISED is question-begging.<sup>15</sup>

Now, one might wonder what all of these rather unexceptional points have to do with Transmission of Knowledge. To see the answer, recall that there is a more substantive characterisation of the aim of inquiry/ in inquiry in terms of knowledge:

### K-AIM

Inquiry into whether p aims at knowledge that p/not-p.<sup>16</sup>

It is easy to see that from Question Settling Aim, Question-Begging, No Settling and Knowledge Sufficiency, it follows that, in Inquiry 2, your inquiry into  $\sim$ DISGUISED does not lead you to knowledge that  $\sim$ DISGUISED. But now recall how your inquiry in Inquiry 2 proceeds: from ZEBRA, which you know, you competently deduce and thereupon come to believe that  $\sim$ DISGUISED. If Transmission of Knowledge holds, your inquiry leads you to knowledge that  $\sim$ DISGUISED. Hence, Transmission of Knowledge fails.

Can we specify a certain *type* of case in which Transmission of Knowledge fails due to question-begging? Yes. To see how, note that the reason why, in Inquiry 2, your inquiry into whether  $\sim$ DISGUISED begs the question generalises. When we settle a given question, we do so in a certain way, W. Crucially, settling the question whether p in way Wp will always require that a certain set of facts  $\Delta_{Wp}$  are already in place.

---

<sup>15</sup> It may be worth noting that even champions of Transmission of Knowledge and the related transmission of warrant principle acknowledge this point (Pryor 2000, Markie 2005, Pritchard 2007).

<sup>16</sup> It may be worth noting that, for the purposes of this argument, all that I need is the sufficiency direction of Knowledge Aim. Since the competition in the literature advocate weaker aims of inquiry, they would all agree on the sufficiency direction of Knowledge Aim. By the same token, the relevant premise of my argument should not raise too many eyebrows.

More specifically, I want to suggest that  $\Delta_{Wp}$  is the set of facts that obtains if and only if what we may call the ‘false positive conditions’ of  $Wp$  do not obtain. By the ‘false positive conditions’ I mean the conditions under which  $W$  produces a belief on whether  $p$  that falls short of knowledge. Anyone who then ventures to settle any further question whether  $q$  by deducing, say,  $q$  from  $p$ , where  $q$  corresponds to some  $i \in \Delta_{Wp}$  will fall foul of question-begging. If, in addition, one knows that  $p$  by having settled the question whether  $p$  in  $Wp$  and has come to believe  $q$  by competent deduction from  $p$ , we will have a case of Transmission of Knowledge failure due to question-begging. We can thus specify a type of case in which Transmission of Knowledge fails due to question-begging (henceforth also ‘The Case’).

With the argument against Transmission of Knowledge in play, I will now return to my preferred account of knowledge. My aim here is to show that this account has the resources to accommodate not only that Transmission of Knowledge fails in the relevant type of case, but also that all remaining instances of Transmission of Knowledge continue to hold. In order to achieve this, I will first argue that, given certain assumptions, my account of knowledge validates an unrestricted version of Transmission of Knowledge. This establishes that the account can license instances of Transmission of Knowledge, including the target ones. Then I will show how the assumptions that commit me to Transmission of Knowledge can be replaced in such a way that Transmission of Knowledge fails just in The Case.

First things first, recall the substantive conditions on knowledge that I developed in Section 3:

#### Conditions on Knowledge

One knows that  $p$  only if (K1 =) one’s belief that  $p$  competent, and (K2 =) the  $C$  of the ability to know that produced one’s belief that  $p$  are satisfied.

#### Competent Belief

One’s belief that  $p$  is competent if and only if it is formed by an exercise of an ability to know propositions in range  $R$  and relative to conditions  $C$  such that  $p \in R$ .

Recall also abilities are intrinsically connected to success by the relevant dispositions. The strength of the success-connection may vary in a number of ways. For instance, there may be variation in the probability of success conditional on the exercise of ability in  $C$ . The thought here is that the higher the probability of success given the exercise of ability in  $C$ , the stronger the ability (relative to  $C$ ). Most importantly for

present purposes, the strongest abilities (relative to C) feature *surefire dispositions*, i.e. dispositions such that triggering (the exercise of ability) in suitable conditions (C) *guarantees* manifestation (success). As a result, for abilities to know featuring surefire dispositions, competent belief in C entails knowledge.

Consider, next the following bridge principles between Transmission of Knowledge and Conditions on Knowledge:

Bridge 1

[C2] and [C3] hold if and only if one believes that  $q$  via the exercise of a deductive epistemic ability (DEA $_q$ ).

Bridge 2

[C1] identifies a member of  $C_{DEA_q}$ .

Now consider the following assumptions:

Assumption 1

The disposition at issue in DEA $_q$  is a surefire disposition.

Assumption 2.

There are no other members of  $C_{DEA_q}$  besides [C1].

Assumption 3.

The range of DEA $_q$  includes all propositions  $q$  one may competently deduce from  $p$ .

My account of knowledge, Bridge 1 and 2, and Assumption 1 – 3 jointly entail (and thus validate) Transmission of Knowledge. Here's why. To begin with, by Bridge 1, Assumption 3 and Competent Belief, if [C2] and [C3] hold, then one will competently believe that  $q$ . Next, by Bridge 2 and Assumption 2, [C1] specifies *the only* members of  $C_{DEA_q}$ . Hence, if [C1] holds, the set of conditions to which DEA $_q$  is relative must be satisfied. Hence, if [C1] – [C3] hold, then one must also competently believe that  $q$  in  $C_{DEA_q}$ . But now recall that for any ability featuring a surefire disposition, competent belief in C entails success. Since, by Assumption 1, the disposition at issue in DEA is a surefire disposition, it follows that if [C1] – [C3] hold, then one must not only competently believe that  $q$  in  $C_{DEA_q}$ , but one is also guaranteed knowledge that  $q$ , i.e. [C4] holds. Given my preferred account of knowledge, Bridge 1 and 2, and

Assumption 1 – 3, we get the result that if [C1] – [C3] hold, then [C4] must hold also. Transmission of Knowledge is thus validated.

Next, I will show how to accommodate failures of Transmission of Knowledge in and just in The Case. The first and most important step is to connect  $\Delta_{Wp}$  with Conditions of Knowledge. Since on my view of knowledge, one has settled the question whether  $p$  if and only if one knows that  $p$ /not- $p$ , it is immensely plausible that the way of settling the question whether  $p$  referred to in The Case, just is an epistemic ability (EA) that has  $p$  in its range at issue in Conditions on Knowledge. In other words, it is immensely plausible that  $Wp = EA_p$ . Moreover, it is also immensely plausible that the set of facts that  $Wp$  requires to be in place for the question whether  $p$  to be settled will just be the set of conditions to which the epistemic ability is relative. In other words, it is also immensely plausible that  $\Delta_{Wp} = C_{EA_p}$ . Given that this is so, we can describe The Case in terms of the conceptual resources of Conditions of Knowledge: (i) one knows that  $p$  via the exercise of an epistemic ability,  $EA_p$  that requires  $C_{EA_p}$  to be in place, (ii) one has come to believe that  $q$  via the exercise of a deductive epistemic ability,  $DEA$ , and (iii)  $q$  corresponds to some  $i \in C_{EA_p}$ .

With The Case so described, here's what we need to do to accommodate the relevant failures of Transmission of Knowledge: we need to disallow instances of Transmission of Knowledge for any  $q$  such that  $q$  corresponds to some  $i \in C_{EA_p}$ . Now, I take it that while we want Transmission of Knowledge to fail, we also want transmission of justification to fail in The Case. After all, the question-begging character of your inquiry not only prevents you from coming to know the conclusion of the relevant deductions, but also from coming to believe it justifiably.

With this point in play, here is what I take to be the most promising way forward: deny Assumption 3. Instead, we may hold that the range of  $DEA_q$  includes only propositions  $q$  such that  $q$  does not correspond to any  $i \in C_{EA_p}$ , alongside the requirement that one may competently deduce  $q$  from  $p$  of course. It is easy to see that this will (i) accommodate the failure of Transmission of Knowledge in The Case, whilst (ii) continuing to allow instances of Transmission of Knowledge to hold in all other cases.<sup>17</sup>

What remains to be seen is that this account of transmission failure is relevant to our task of dealing with The Argument from Ignorance. Fortunately, it is easy enough to see that it is. Sceptical hypotheses are paradigm cases of false positive conditions for our nearly all of our ways of belief formation. As a result, the falsity of

---

<sup>17</sup> For a more detailed discussion and defence of this approach to Transmission of Knowledge failure, see (Kelp 2019b).

sceptical hypotheses are among the conditions that need to be in place for perceptual ways of forming beliefs to settle questions (alternatively: for our epistemic abilities to give us knowledge). On the above account of transmission failure, knowledge will fail to transmit across deductions from ordinary empirical propositions to the denials of sceptical hypotheses. For instance, knowledge will fail to transmit across the deduction from *you have hands* to *you are not a handless BIV*. And, of course, if knowledge fails to transmit from ordinary empirical propositions to the denials of sceptical hypothesis, the key motivation that Transmission of Knowledge provided for AI2 will be lost.

### 5.2 Anti-Sceptical Knowledge

Can't we just deny AI2 and leave it at that? Unfortunately, the answer to this question is no. One important reason for this is that we threaten to be committed to abominable conjunctions such as:

Abomination

I know that I have hands but I don't know that I am not a handless BIV.

No one in their right mind would want to be committed to something abominable. Accordingly, views that do generate such commitments do thereby suffer a strike against them. If all we do to rescue ourselves from the sceptical conclusion AI3 is to deny AI2, whilst continuing to grant the sceptic AI1, we will face just this difficulty (DeRose 1995).

Here is a related problem. Suppose that knowledge is the norm of assertion. If so, by the lights of the present proposal, you may permissibly assert that you have hands and that this entails that you are not a handless BIV. However, you cannot permissibly assert that you are not a handless BIV. To see just why this is bad, suppose someone asked a champion of the present proposal whether they have hands. They'd of course say yes. And when asked whether this entails that they are not handless BIVs, they'd again say yes. But, of course, when asked whether they are not handless BIVs, they couldn't say yes. Rather, they'd have to say something like 'I can't say.' But that is just absurd. As Hawthorne aptly puts it, in this case they are just like "Lewis Carroll's Tortoise, that familiar object of ridicule who was perfectly willing to accept the premises of a *modus ponens* argument but was unwilling to accept the conclusion" (Hawthorne 2005, 32).

I take it that these considerations indicate that we cannot respond to the sceptical argument simply by denying Transmission of Knowledge/AI2. Rather, we



must also deny AI1. The question that immediately arises is how this can be done. In particular, we may wonder how you could possibly know that you are not a handless BIV given that everything would seem just the same if you were.

There is a prominent response to this question available on the market. According to this response, championed most famously by G.E. Moore, you know that you are not a handless BIV based on competent deduction from your knowledge that you have hands. But, of course, the Moorean response to our question requires (relevant instances of) Transmission of Knowledge and so is unavailable to those who want to adopt the above way of denying Transmission of Knowledge. So, it looks as though foes of Transmission of Knowledge face a particularly difficult challenge here.

To see how champions of the present approach to transmission failure can deny AI1, I'd first like to focus a slightly different type of case, which does not involve sceptical hypotheses. Suppose, for instance, someone were to put to us the hypothesis that (TEAPOT =) there is a teapot orbiting the Sun somewhere between the Earth and Mars or that (CREATOR =) the universe was created by a supernatural creator closely resembling a serving of spaghetti with meatballs. As a first observation, I want to insist that we know that these hypotheses are false. At the same time, it doesn't look as though we will be able to produce a conclusive argument against these hypotheses. At the very least, many of us won't be able to do this. So how can we know these facts?

To arrive at what I take to be a promising answer to this question, I'd like to ask how a sane person might react to being presented with such an hypothesis. Here is what I take to be the sane response: "Get real! That's just crazy!" What's going on here? My suggestion is that we have *an epistemic ability that enables us to recognise that certain possibilities could not easily obtain*.<sup>18</sup> It is via an exercise of this ability that we come to know that the possibilities described in TEAPOT and CREATOR could not easily obtain (and hence that  $\sim$ TEAPOT and that  $\sim$ CREATOR are both true). This suggestion offers an attractive explanation of the sane response. Upon consideration of the possibility described in TEAPOT and CREATOR, you come to know that these possibilities could not easily obtain by an exercise of your epistemic ability to recognise possibilities that could not easily obtain. "Get real! That's just crazy!" is an expression of this very knowledge. The fact that my suggestion offers an attractive explanation of the sane response provides some evidence that it is on the right track.

There is further support for my suggestion. Consider first the following two important properties of beliefs produced by recognitional abilities (henceforth

---

<sup>18</sup> Easy possibilities can be understood in terms of what happens at nearby possible worlds: a possibility could not easily obtain if and only if it does not obtain at any nearby possible world.

*recognitional beliefs*). First, recognitional beliefs are cognitively spontaneous, they do not appear to be based on further evidence. For instance, when you form a perceptual recognitional belief that there is a chair before you, you do not appear to base this belief on further evidence such as the appearance as of a chair. Second, in case of recognitional beliefs, we are often unable to articulate how the recognitional ability works in any detail. For instance, someone may have the ability to recognise pieces by Beethoven without being able to articulate how he does so (other than by making the rather uninformative observation that it sounds like Beethoven) (Millar 2010). But now note that, for our beliefs that the possibilities described in TEAPOT and CREATOR could not easily obtain, these two key properties of recognitional beliefs are plausibly satisfied. First, when we acquire these beliefs, we do not appear to base them on further evidence. Rather they are cognitively spontaneous. Second, we may not be able to articulate how we recognise the relevant possibilities as such (other than by making the rather uninformative observation that they just look crazy). The fact that our beliefs fit the profile of recognitional beliefs further confirms my suggestion.<sup>19</sup>

If I am right, we have an epistemic ability that enables us to recognise whether certain possibilities could not easily obtain. It will come as no surprise that it is just this ability that I want to suggest enables us to come to know the denials of sceptical hypotheses. Here is how. In many situations, the above account can be adapted to explain our knowledge of the denials of sceptical hypotheses. For instance, suppose that someone asks you whether you have any siblings. Suppose you respond that you have two and I interject by asking how you can be so sure, after all you might be a BIV. In this case, the sane response is once again something along the lines of: "Get real! That's just crazy!"<sup>20</sup> Alternatively, consider a court case in which the solicitor attempts to undermine the prosecution's case against the defendant by invoking sceptical possibilities. Again, we'd expect a response very much along the lines of "Get real! That's just crazy!" (except, of course, that it will likely be phrased slightly differently). Just as in the TEAPOT and CREATOR cases, what's doing the work here is your epistemic ability to recognise that certain possibilities could not easily obtain, which offers an attractive explanation of the sane response along the following lines:

---

<sup>19</sup> Just how this ability might operate is a fascinating question. It may not be the job for philosophers to answer it, at least not fully, and I will not attempt to do so here. Instead, I will rest content with the observation that, even if we do not have an account of how the ability operates, the above still provides some evidence that we do indeed have this ability.

<sup>20</sup> Note also that the sane response is not, for instance: "Well, I have syblings, so the BIV possibility doesn't obtain." As a result, Moorean responses to the sceptical argument do not fit well with the common sense reaction here.

upon consideration of the possibility that you might be a BIV, you come to know that this possibility could not easily obtain by an exercise of your epistemic ability to recognise possibilities that could not easily obtain.<sup>21</sup> Again, “Get real! That’s just crazy!” is an expression of this very knowledge.<sup>22</sup>

In this way, even champions of the present approach to transmission failure can deny AI1. That is to say, they can allow that we know the denials of sceptical hypothesis. Crucially, they are known not via competent deduction from ordinary empirical proposition but rather via the exercise of a separate recognitional ability.<sup>23</sup>

### **Conclusion**

Recall the central methodological idea that this paper adopted, viz. that epistemology is the theory of inquiry. I have ventured to provide support for this idea by showing that it allows us to provide, in a systematic way, attractive new solutions to a number of old epistemological problems, including the nature, value and extent of knowledge as well as sources of knowledge. I have now completed this project, or, at the very least, I have completed as much as of it as I can within the confines of one paper. It is my hope that I have managed to make the central methodological idea attractive enough to be taken seriously as a novel way of doing epistemology.

One thing I haven’t done is look into how the central methodological idea compares with other methodologies. In particular, one might wonder how the view relates to Williamson’s knowledge first epistemology. Note that on my view, knowledge plays a central role in epistemological theorising and I also give a non-

---

<sup>21</sup> But doesn’t this view commit me to the highly implausible claim that we can know that a certain ticket in a fair lottery won’t win, by means of this very ability to recognise that certain possibilities could not easily obtain? No. To see why not, note that sceptical possibilities do not obtain at nearby possible worlds. For instance, there is no nearby possible world at which you are a BIV. As a result, sceptical possibilities are possibilities that could not easily obtain, and we may well come to know this via the exercise of our epistemic ability to recognise that they could not easily obtain. Possibilities concerning the winning of a fair lottery, in contrast, do obtain at nearby possible worlds. For instance, even if your ticket is a loser at the actual world, there is a nearby possible world at which it is a winner. As a result, they are possibilities that can easily obtain. And since knowledge is factive this means that we cannot come to know of possibilities concerning the winning of a lottery that they could not easily obtain. Thanks to an anonymous referee for pressing me on this.

<sup>22</sup> It is also easy to see that our beliefs here fit the profile of recognitional beliefs just as well as in the TEAPOT and CREATOR cases with the result that my proposal is further confirmed.

<sup>23</sup> It may also be worth noting that invoking this ability in my response to the sceptic is not ad hoc. After all, the evidence for the existence of the ability is already provided by the TEAPOT and CREATOR cases. As far as the response to scepticism is concerned, then, the argument invokes an ability that there is independent reason to think exists.

reductive account of knowledge. Does this mean that my view is just an instance of knowledge first epistemology? No. To see this, note that the key methodological idea of knowledge first epistemology is that it takes the distinction between knowledge and ignorance as the starting point for epistemological theorising (Williamson 2010, 208). The present approach does not embrace this key idea. Here is why. That knowledge is of central importance in epistemological theorising is a function of the methodological idea in conjunction with a set of arguments. On my view, knowledge must earn its keep as something of key significance in epistemological theorising. By the same token, epistemological theorising does not start from knowledge. In this way, the present approach differs from knowledge first epistemology. Of course, this immediately raises the question as to which of the two approaches is preferable. While there is much to be said about this question, I am running out of space quickly now, which I will have to leave this one for another occasion.

### References

- Adler, J. 2002. *Belief's Own Ethics*. Cambridge/MA.: MIT Press.
- David, M. 2001. "Truth as the Epistemic Goal." In *Knowledge, Truth, and Duty*, edited by M. Steup. Oxford: Oxford University Press.
- Davidson, D. 2005. *Truth, Language and History*. Oxford: Clarendon Press.
- DeRose, K. 1995. "Solving the Skeptical Problem." *Philosophical Review* 104: 1–52.
- Dretske, F. 1970. "Epistemic Operators." *Journal of Philosophy* 67: 1007–23.
- Feldman, R. 2002. "Epistemological Duties." In *The Oxford Handbook of Epistemology*, edited by P. Moser. New York: Oxford University Press.
- Friedman, J. 2019. "Inquiry and Belief." *Noûs* 53: 296–315.
- Gettier, E. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23: 121–23.
- Goldman, A. 1976. "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73: 771–91.
- Greco, J. 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.
- Harman, G. 1986. *Change in View*. Cambridge/MA: MIT Press.
- Hawthorne, J. 2005. "The Case for Closure." In *Contemporary Debates in Epistemology*, edited by M. Steup and E. Sosa. Oxford: Blackwell.
- Hieronymi, P. 2009. "Two Kinds of Agency." In *Mental Actions*, edited by L. O'Brien and M. Soteriou. Oxford: Oxford University Press.
- Huemer, M. 2007. "Moore's Paradox and the Norm of Belief." In *Themes from G.E. Moore: New Essays in Epistemology and Ethics*, edited by S. Nuccetelli and G. Seay. The Collected Works of John Stuart Mill. Oxford: Clarendon Press.

- Kelp, C. 2014. 'Two for the Knowledge Goal of Inquiry.' *American Philosophical Quarterly* 51, 227-32.
- — —. 'Justified Belief: Knowledge First-Style.' *Philosophy and Phenomenological Research* 93, 79-100.
- — —. 2017. 'Knowledge First Virtue Epistemology.' Carter, A., Gordon, E. and Jarvis, B. eds. *Knowledge First: Approaches in Epistemology and Mind*. Oxford: Oxford University Press.
- — —. 2018. *Good Thinking: A Knowledge First Virtue Epistemology*. New York: Routledge.
- — —. 2019a. 'How to Be a Reliabilist.' *Philosophy and Phenomenological Research* 98, 346-74.
- — —. 2019b. 'Inquiry and the Transmission of Knowledge.' *Philosophy and Phenomenological Research* 99, 298-310.
- — —. Forthcoming. 'Inquiry, Knowledge and Understanding.' *Synthese*.
- Kelp, C. and Simion, M. 2018. "The C Account of Assertion: A Negative Result." *Synthese*. <https://doi.org/10.1007/s11229-018-1760-5>.
- Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Lynch, M. 2005. *True To Life. Why Truth Matters*. Cambridge/MA: MIT Press.
- Markie, P. 2005. "Easy Knowledge." *Philosophy and Phenomenological Research* 70: 406–16.
- McHugh, C. 2011. "What Do We Aim at When We Believe?" *Dialectica* 65: 369–92.
- Millar, A. 2011. "Why Knowledge Matters." *Proceedings of the Aristotelian Society, Suppl. Vol.* 85: 63–81.
- Moore, G.E. 1939. "Proof of an External World." In *Philosophical Papers*. London: Allen & Unwin, 1959.
- Nozick, R. 1981. *Philosophical Explanations*. Oxford: Oxford University Press.
- Plato. 1956. "Meno." In *Protagoras and Meno*, edited by W.K.C. Guthrie. London: Penguin.
- Pritchard, D. 2007. "How to Be a Neo-Moorean." In *Internalism and Externalism in Semantics and Epistemology*, edited by S. Goldberg. Oxford: Oxford University Press.
- Pritchard, D., A. Millar, and A. Haddock. 2010. *The Nature and Value of Knowledge*. Oxford: Oxford University Press.
- Pryor, J. 2000. "The Skeptic and the Dogmatist." *Noûs*, no. 34: 517–49.
- Rorty, R. 1995. "Is Truth a Goal of Enquiry? Davidson vs. Wright." *The Philosophical Quarterly* 45: 281–300.

- Simion, M., C. Kelp, and H. Ghijsen. 2016. "Norms of Belief." *Philosophical Issues* 26: 375–92.
- Sosa, E. 1999. "How to Defeat Opposition to Moore." In *Philosophical Perspectives*, edited by J. Tomberlin. Oxford: Blackwell.
- — —. 2007. *A Virtue Epistemology. Apt Belief and Reflective Knowledge*. Vol. 1. Oxford: Oxford University Press.
- Strawson, P. 1992. *Analysis and Metaphysics: An Introduction to Philosophy*. Oxford: Oxford University Press.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- — —. 2010. "Knowledge First Epistemology." In *The Routledge Companion to Epistemology*, edited by S. Bernecker and D. Pritchard, 208–18. New York: Routledge.