

‘Enough is enough’: Austin, knowledge and context

Aaron Baird

A thesis submitted in partial fulfilment  
of requirements for the degree of Master of Philosophy

Faculty of Arts and Social Sciences

University of Sydney

2018

## Abstract

This thesis is concerned with J.L. Austin's work on the topic of empirical knowledge. Austin encourages us to attend to our everyday epistemic and discursive practices, and specifically to the particular circumstances in which we might ordinarily say that a person knows something. I begin by considering what kind of illumination on the topic of empirical knowledge we might expect to get by following Austin's approach, and defend Austin's approach against one influential critique. The focus then shifts to one of Austin's key observations regarding knowledge, namely that knowing is a matter of having done 'enough' for present intents and purposes to establish the truth. I argue that this and other Austinian considerations speak in favour of a contextualist account of knowledge. Finally, I present a novel Austin-inspired response to one particular sceptical puzzle occasioned by what have been referred to as 'arguments from ignorance'.

This is to certify that to the best of my knowledge, the content of this thesis is my own work. This thesis has not been submitted for any degree or other purposes.

I certify that the intellectual content of this thesis is the product of my own work and that all the assistance received in preparing this thesis and sources have been acknowledged.

There are a few people who need thanking. Thanks to David Macarthur, for his supervision over the course of the degree. I moved to Sydney specifically to study under David and am grateful to have had the opportunity to do so. Thanks to my parents, my sisters, and my brother, for providing endless support and encouragement from interstate. Thanks to Ollie, Ned, Dan, Blaise, Alistair, and Elle, for keeping philosophy fun and fresh and spirited. And thanks to the Philosophy Department at the University of Sydney, for providing a stimulating and challenging learning environment.

# **Contents**

<b>Introduction .....</b>	<b>1</b>
<b>Chapter 1. Austin’s approach to thinking about knowledge.....</b>	<b>6</b>
1.1 A sceptical conception of knowledge .....	7
1.2 Austin on knowledge .....	10
1.3 Stroud on Austin and appeals to everyday epistemic practice .....	17
1.4 Stroud miscasts Austin’s approach .....	23
1.5 Summary .....	28
<b>Chapter 2. Austin, relevant alternatives, and a contextualist account of knowledge .....</b>	<b>29</b>
2.1 Knowledge ascriptions and relevant alternatives.....	31
2.2 Contrasting cases .....	36
2.3 A contextualist interpretation of the cases .....	40
2.4 ‘Knowing’ and ‘being green’ .....	43
2.5 Issues with an ‘invariantist’ account of knowledge .....	45
2.6 Assurance-giving and contextualism about knowledge .....	50
2.7 Summary .....	54
<b>Chapter 3. An Austin-inspired response to the ‘argument from ignorance’ .....</b>	<b>55</b>
3.1 The argument from ignorance .....	58
3.2 An Austin-inspired response to the argument from ignorance.....	61
3.3 Why do arguments from ignorance seem compelling?.....	66
3.4 What about closure?.....	70
3.5 Comparison with Mark Kaplan’s Austin-inspired strategy for dealing with arguments from ignorance .....	73
3.6 Comparison with David Lewis’s contextualist response to the sceptical puzzle .....	80
3.6.1 Lewis’s contextualism and the Rule of Attention .....	81
3.6.2 Closure and Lewis’s response to the sceptical puzzle .....	84
3.7 Summary .....	87
<b>Concluding remarks .....</b>	<b>88</b>
<b>Bibliography .....</b>	<b>91</b>

# Introduction

This thesis aims to highlight the continuing relevance and importance of J.L. Austin's work on the topic of empirical knowledge. Before setting out how the thesis aims to do this, I will first say a word of introduction regarding philosophical approaches to thinking about knowledge more generally.

Philosophers as far back as Plato<sup>1</sup> have taken an interest in knowledge – in what we know and what 'knowing' amounts to. When I say that I or someone else knows something, what am I saying to be so? Plato foreshadows one contemporary approach to thinking about knowledge in suggesting that a philosophically satisfying answer to this question should take the form of a non-circular *definition* of what it is to know something, given in terms of jointly necessary and sufficient conditions. A condition will be *necessary* for knowledge if the person cannot know in the absence of that condition being met. And a set of conditions will be jointly *sufficient* for knowledge if a person's satisfying these conditions means that he cannot but know the thing in question.

Plato also provides a suggestion as to what form such an analysis should take. He suggests that to say a person knows something is at least to say the person has a true belief on the matter. For example, to say that Ollie knows that there is a goldfinch in the garden is to say that Ollie at least believes that there is a goldfinch in the garden and that his belief is true. However, Plato recognises that there can be cases in which a person has a true belief on a matter, but doesn't count as knowing the thing in question. Ollie could have a true belief that there is a goldfinch in the garden, and yet not count as knowing that there is a goldfinch in the garden. Ollie might, for example, have come to believe that there was a goldfinch in the garden as a result of seeing a cardboard cut-out of a goldfinch from some distance away. As it turns out, there *was* a goldfinch in the garden, hidden from Ollie's view (behind a fern). In that case Ollie's belief that there was a goldfinch in the garden was true, but no one would want to say that Ollie *knew* there was a goldfinch in the garden. So even if a person's having a true belief is necessary for knowing, it is hardly sufficient. This leads to the thought that knowing is to be understood as "true belief plus something more", to borrow an expression from Donald Davidson (1996, p. 263).<sup>2</sup> But true belief

---

<sup>1</sup> Particularly in the *Theaetetus* (1973) and the *Meno* (2006).

<sup>2</sup> Davidson does not himself endorse the view that the concept of knowledge can be analysed in terms of necessary and sufficient conditions.

plus *what*? This is question exercised Plato and came to dominate the epistemological agenda for a time in the twentieth century. It continues to frame much epistemological work on knowledge today.

The problem is that no analysis along these lines has been successful to date. For every suggested analysis of knowledge in terms of true belief plus some further condition(s), counterexamples have been suggested. In the context of this kind of philosophical program, counterexamples can take one of two main forms. First there are counterexamples aimed at showing that a condition is not a necessary condition for knowledge. To show this, we simply need to come up with a case in which a person is in a position to know something without meeting the condition in question.<sup>3</sup> More common, however, are counterexamples aimed at showing that a certain set of conditions alleged to be jointly sufficient for knowledge are not in fact sufficient. To show this, we simply need to come up with a case in which all of the conditions are met, but in which the person in question still doesn't count as knowing. The most famous example of this strategy in action can be found in Edmund Gettier's (1966) 'Is justified true belief knowledge?' In this paper Gettier provides two examples which have been widely taken to show that a person's having a justified true belief on a matter is not sufficient for that person's knowing the thing in question. Put otherwise, it looks as though a person can have a justified true belief on a matter without knowing the thing in question. Similar counterexamples have been found for other suggested analyses of knowledge.<sup>4</sup>

Not everyone is convinced that analysis should be our aim when it comes to thinking about knowledge, or even that the concept of knowledge *is* analysable in this way.<sup>5</sup> One of the most famous critics of such an approach to thinking about knowledge in recent years is Timothy Williamson (2000), who argues that the long history of failed analyses of knowledge provides inductive evidence against the likelihood of its success.<sup>6</sup> Williamson sits at one end of a line of thinkers at Oxford who have similarly rejected the thought that knowing might be analysed into more basic components. It is a line of thinkers that includes Austin and extends back to John

---

<sup>3</sup> Colin Radford (1966) provides a famous example problematizing the view that belief is a necessary condition for knowledge.

<sup>4</sup> For an overview of some early attempts to overcome the 'Gettier problem', see Shope (1983). For a more recent overview of issues relating to the analysis of knowledge, see Hetherington (1996).

<sup>5</sup> It is a good question whether there are many concepts at all that we might expect to be able to provide a Platonic definition for.

<sup>6</sup> For another influential criticism of 'true belief plus' approaches to analysing knowing, see Zagzebski (1994).

Cook Wilson.<sup>7</sup> For such thinkers it is a mistake to try and analyse a factive<sup>8</sup> state like knowing that  $p$  in terms of a non-factive state like believing that  $p$ . So whatever insights regarding knowledge that Austin furnishes us with, they won't take the form of an analysis in the 'Platonic' sense.

Another common way in which philosophers have approached the topic of empirical knowledge is through the frame of scepticism. There are a variety of sceptical arguments and puzzles purporting to make questionable the idea that we are ever in a position to know anything about the world. One philosophical approach involves assessing the force of such sceptical arguments, seeing what illumination there is to be had on the concept of knowledge through showing where (if at all) such sceptical arguments go wrong. Many sceptical puzzles stem from the thought that a person is in no position to know something if he might be mistaken about the thing in question, combined with the observation that there are few things about which we in principle couldn't turn out to be mistaken about. These two points taken together can appear to make it mysterious how we could ever be in a position to know anything about the world. Austin's work is particularly insightful on this topic, so in Chapter 1 this is where I begin.

But if Austin doesn't provide us with an analysis of what it is to know something, what does he provide? What Austin provides in the first instance is a perspicuous presentation of our epistemic and discursive practices as they relate to the concept of knowledge. In large part this involves calling attention to the kinds of circumstances in which we might ordinarily claim or ascribe knowledge. Austin's view is that getting a clear understanding of how we ordinarily use a word like 'knows' can enable us to see where different sceptical arguments go wrong. At the same time he thinks that such an exercise will give us a greater appreciation for what knowing comes to in different situations. In what follows I aim to tease out some of the insights Austin's approach furnishes us with, and how these insights can be used to defuse at least some sceptical puzzles regarding the concept of knowledge.

In **Chapter 1** I begin with a sceptical conception of knowledge according to which a person only counts as knowing that  $p$  if he is in a position to 'rule out'<sup>9</sup> any conceivable alternative to  $p$ . Such a conception of knowledge looks plausible in the abstract, but appears to have the absurd

---

<sup>7</sup> I owe this observation to Charles Travis and Mark Kalderon (2013). They point out that John McDowell (1995) also belongs in this tradition.

<sup>8</sup> To say that knowing is 'factive' is simply to capture the idea that one cannot know that  $p$  unless  $p$  is so. Ollie cannot know there is a goldfinch in the garden if there isn't a goldfinch in the garden. On this point just about everyone agrees.

<sup>9</sup> I spell out what I have in mind in talking about being in a position to 'rule out' an alternative in 1.1.



implication that we are never in a position to know anything. I then introduce some of Austin's key observations regarding everyday epistemic and discursive practice, observations which provide some insight into how this sceptical conception of knowledge is mistaken. Austin observes that generally when we say that a person knows that  $p$ , what we are saying is that he has done enough for present intents and purposes to establish that  $p$ . What counts as 'enough' in a situation is rough, but it doesn't require that a person be in a position to rule out *any* conceivable alternative to what he is taken to know. Rather, we will generally only take a person's inability to rule out an alternative to count against his claim to know in circumstances in which we have some reason to seriously suppose that the alternative might in fact obtain. Having introduced Austin's views on this matter, the rest of the chapter is concerned with Barry Stroud's influential criticism of Austin's work on knowledge. Stroud argues that Austin's approach merely furnishes us with considerations regarding when it is *appropriate* to say that a person knows something, and that it tells us nothing about what it takes for a person to actually count as knowing something. I argue that Stroud miscasts the import of Austin's observations. Austin is not 'merely' interested in when it is appropriate to say that someone knows something. Rather, he is interested in what a person could reasonably be understood to be saying in saying that a person *knows* in such circumstances. And Stroud provides no good reason for thinking that our understanding of what a person could reasonably be taken to be saying in such circumstances doesn't reflect our understanding of what it takes to count as knowing in such circumstances. So Austin's approach still looks to be a useful way of articulating our understanding of what it takes to count as knowing in various circumstances.

In **Chapter 2** I build on Austin's observations and method, and argue that Austinian considerations speak in favour of a contextualist account of knowledge. I start by suggesting that the notion of a 'relevant alternative' can be used to articulate our understanding of what kind of a position a person is said to be in in being ascribed knowledge in particular circumstances. I then argue that what kind of position a person is said to be in can vary depending on the circumstances in which a speaker ascribes knowledge to the person. The same person might count as knowing that  $p$  on one understanding of what it is to know that  $p$ , and not count as knowing on another understanding of what it is to know that  $p$ . I argue that such an account of knowledge allows us to make better sense of certain aspects of our epistemic and discursive practices.

In **Chapter 3** I consider a sceptical puzzle occasioned by what have come to be known as 'arguments from ignorance'. Reflecting on arguments from ignorance is liable to provoke conflicting intuitions in us regarding what it takes to know something, and this in turn might lead

us to suppose that there is an incoherence latent in our everyday epistemic and discursive practices as they relate to the concept of knowledge. I put forward an Austin-inspired response aimed at showing that we needn't suppose there to be any such incoherence within our practices. Key to this response is the idea that there can be circumstances in which it is neither true nor false to say that a person 'knows' something, and this because in such circumstances we wouldn't be saying anything determinately to be so. I then compare my Austin-inspired response to two other responses to the puzzle. The first is Mark Kaplan's Austin-inspired response to arguments from ignorance, which resembles my own in many respects but also differs on some important particulars. I end by comparing my Austin-inspired account to David Lewis's contextualist response to such puzzles, which I find to be dissatisfying on a number of counts.

By the end of this thesis I aim to have shown that Austin's work furnishes us with resources for thinking about issues relating to the topic of empirical knowledge that continue to occupy philosophical attention today.

# Chapter 1. Austin's approach to thinking about knowledge

In this chapter I introduce Austin's approach to reflecting on the concept of knowledge, an approach that involves reflecting on how the concept functions within everyday epistemic and discursive practice. I also defend Austin's approach against Barry Stroud's famous criticism to the effect that such an approach should not be taken to provide any direct illumination regarding what it takes to know something. Stroud casts Austin as being interested merely in considerations regarding when we would ordinarily think it appropriate to say that someone knows, and he argues that no such considerations can be taken to establish anything regarding what it takes to actually know something. I argue that Stroud miscasts Austin's approach. Austin is first and foremost interested in what we could reasonably be taken to be saying in particular situations where we might say that a person knows. Stroud gives us no good reason to suppose that such considerations do not reflect our understanding of what it takes to count as knowing in such circumstances. He therefore gives us no good reason to suppose that Austin's approach fails to illuminate what it takes to count as knowing in particular circumstances.

I begin by considering a conception of knowledge according to which a person only counts as knowing that  $p$  if he is in a position to rule out any conceivable alternative to  $p$ . Such a conception of knowledge, while compelling in the abstract, has the consequence that no one ever counts as knowing anything. For this reason I will refer to this conception of knowledge as a *sceptical* conception of knowledge.

Next I turn to remarks of Austin's regarding everyday epistemic and discursive practice which can help us see where such a conception of knowledge goes wrong. Austin reminds us that when we claim to know something we are claiming to have done enough for present purposes to establish the truth of the matter, where this doesn't mean being in a position to rule out any conceivable alternative. In some cases a person's inability to rule out an alternative to  $p$  will count against his claim to know that  $p$ , but only in circumstances where we have reason to take seriously the alternative in question.

Having considered Austin's remarks on the subject I then turn to Barry Stroud's influential critique of Austin's approach. Stroud argues that Austin's approach furnishes us only with observations regarding when it is appropriate to say that someone knows something, and that

such observations should not be taken to provide any direct illumination of what it takes to actually know something in a situation. Stroud's argument turns on the idea that there is a distinction to be made between the circumstances under which it is appropriate to say something and the circumstances under which it is true to say something. Stroud notes in particular that its being appropriate to say something does not entail that what is said is true. He takes this to leave room for the possibility that, for all Austin's observations regarding everyday epistemic and discursive practice, the sceptical conception of knowledge might in fact be correct.

I end the chapter by arguing that Stroud miscasts Austin's approach, thus making his criticism seem more plausible than it ought to. Austin is not merely interested in when it is appropriate to say that a person knows something. What he is interested in is what a speaker who says that someone knows something can reasonably be understood to be saying in different circumstances. Stroud gives us no good reason to suppose that such considerations do not speak to our understanding of what kind of position a person needs to be in if he is to count as knowing in a particular situation. Considerations regarding what we can reasonably be taken to be saying in saying that a person knows in particular circumstances are still plausibly one of the best resources available to us for illuminating our understanding of what it takes to know something.

### **1.1 A sceptical conception of knowledge**

In this section I adduce some considerations that appear to push us in the direction of a conception of knowledge which, if correct, would mean that we know little (if anything) of what we ordinarily take ourselves to know. On the conception of knowledge in question, a person only counts as knowing that  $p$  if he is in a position to rule out any conceivable alternative to  $p$ . Being able to 'rule out' an alternative in this context means being in a position to distinguish how things appear from how they would appear if the alternative was so. If this conception of knowledge were correct, it looks as though no one would ever be in a position to know anything. This suggests we need to reconsider the proposed conception of knowledge. In the next section I turn to consider J.L. Austin's observations regarding everyday epistemic practice, which suggest just where the sceptical conception of knowledge goes wrong.

Suppose Ollie looks out his window and sees what he takes to be a goldfinch sitting on a branch. Does Ollie know that what he sees is a goldfinch? We can imagine Ollie's being instead presented with a robotic imitation of a goldfinch, perhaps designed by the military (for clandestine purposes). If this were the case, things would appear much the same to Ollie as they would if he

was presented with an actual goldfinch. He would be in no position to tell that what he was seeing was in fact a robotic imitation without some further investigation into the matter. Nor is there anything in his previous experience to suggest that it is strictly impossible that he might be presented with a robotic imitation instead of a goldfinch. In that case it looks as though, for all Ollie can tell, what he sees might in fact be a robotic imitation. So it looks as though Ollie isn't in a position to know that what he sees is a goldfinch. So much the worse for Ollie in this particular case. What I am interested in articulating further is the reasoning employed in this example.

Let  $S$  be some candidate knower and let  $p$  be something he is to know. The above example might be taken to suggest the following. If at a particular point in time  $S$  is in no position to distinguish between  $p$  being so and some conceivable alternative to  $p$  being so, then  $S$  is in no position at that time to know that  $p$ . Another way of putting this same point is to say that if at some time  $S$  knows that  $p$ ,  $S$  must at that time be in a position to distinguish  $p$  being so from any conceivable alternative to  $p$ . If  $S$  is at some time able to distinguish between how things appear to him and how they would appear if some alternative  $q$  were the case, I will say that  $S$  is in a position to 'rule out' the alternative  $q$ .<sup>10</sup> We thus arrive at the following conception of knowledge: a person  $S$  is only in a position to know something  $p$  if he is in a position to rule out every conceivable alternative to  $p$ .

Suppose we took this to be an accurate characterisation of what it takes to know something. This suggests a strategy for working out what a person is at any point in time in a position to know. The strategy in short would be to consider some candidate knower  $S$  and some  $p$  to be known, and then ask ourselves whether there is a conceivable alternative to  $p$ , one which  $S$  is at that time in no position to rule out. If there is a conceivable alternative that  $S$  is unable to rule out, then we will need to conclude that he does not know that  $p$ . But if it turns out there are no conceivable alternatives to  $p$ , or at least none that  $S$  is unable to rule out, then we can conclude that  $S$  knows that  $p$ .

This is how some commentators<sup>11</sup> have understood the strategy pursued by Descartes (1993) in the *Meditations*. Having realised that many of the things he has in the past taken to be so were in fact false, Descartes sets about the task of working out what (if anything) he is in a position to know. To this end he employs the following 'principle of doubt':

---

<sup>10</sup> I will be making recourse to this notion of 'ruling out' throughout the thesis.

<sup>11</sup> For example see chapter 1 of Stroud (1984).

I shall proceed by setting aside all that in which the least doubt could be supposed to exist, just as if I had discovered that it was absolutely false; and I shall ever follow in this road until I have met with something which is certain, or at least, if I can do nothing else, until I have learned for certain that there is nothing in the world that is certain. (Descartes, 1993, p. 50)

When Descartes talks about finding grounds for doubting that something is so, what he has in mind is locating some conceivable alternative that he is in no position to rule out. If he finds that he can conceive of some such alternative to something which he takes to be so, he resolves to suspend his belief on that matter. Descartes's aim is to find something he takes to be so, about which there is not the slightest room for doubt. However, what emerges in the course of Descartes's reflections is that, for just about anything he takes to be so, he can conceive of alternatives that he is in no position to rule out. Reflecting on his belief that he is currently seated by the fire, Descartes has the realisation that he could conceivably be merely dreaming that he is seated by the fire:

At the same time I must remember that I am a man, and that consequently I am in the habit of sleeping, and in my dreams representing to myself the same things or sometimes even less probable things, than do those who are insane in their waking moments. How often has it happened to me that in the night I dreamt that I found myself seated near the fire, whilst in reality I was lying undressed in bed! At this moment it does indeed seem to me that it is with eyes awake that I am looking at this paper; that this head which I move is not asleep, that it is deliberately and of set purpose that I extend my hand and perceive it; what happens in sleep does not appear so clear nor so distinct as does all this. But in thinking over this I remind myself that on many occasions I have in sleep been deceived by similar illusions, and in dwelling carefully on this reflection I see so manifestly that there are no certain indications by which we may clearly distinguish wakefulness from sleep that I am lost in astonishment. (Descartes, 1993, pp. 46-47)

In this passage Descartes suggests that for any waking experience we might have, it is conceivable that we might have a dream that is indistinguishable to us from this waking experience, at least at that time. A more contemporary suggestion in a similar vein is that for any experience we might have, it is conceivable that our brains might be artificially fed the appropriate signals so as to deceive us into thinking we are having that experience. For example, I might be hooked up to a machine that stimulates my brain in such a way as to make it seem to me as though I am looking at a goldfinch. It is conceivable that the deceptive case might be such that I am unable to distinguish it from a case in which I am actually looking at a goldfinch. Insofar as it is conceivable that this could happen, it seems we have a doubt that is ready-made to be applied in relation to just about any case where I might take something to be so. On the conception of knowledge

under consideration, this would suggest that I am not in a position to know much of anything that I take to be so.

Descartes's suggestion that perhaps we might be dreaming is ready-made to be applied in relation to a whole range of candidate cases for knowledge. However, we needn't rely on anything quite as global as Descartes's dream hypothesis in order to show that the present conception of knowledge leaves us in a position to know next to nothing of what we ordinarily take ourselves to. With a bit of ingenuity it is possible to imagine alternatives for just about anything we might take ourselves to know, alternatives that we are in no position to rule out. All we need to do is, as David Lewis (1996, p. 549) suggests, "let our paranoid fantasies rip". The example with which I began this section, in which Ollie is in no position to rule out his having seen a robotic imitation (instead of a goldfinch), illustrates such an approach in action.

What this suggests is that if the conception of knowledge outlined at the beginning of this section is correct, this will carry the implication that we are not in a position to know much (if anything) of what we ordinarily take ourselves to know. Exactly how little this conception of knowledge leaves us in a position to know is a matter that has been debated since at least the publication of Descartes's *Meditations* through to today. It has, for example, been debated whether this conception of knowledge threatens our claim to know certain mathematical truths. But even if it were shown that there were *some* things we might be in a position to know on this conception of knowledge, it is clear that we will nonetheless still be in the position of denying that we know much of what we ordinarily take ourselves to know. For this reason I will refer to the conception of knowledge outlined in this section as the *sceptical* conception of knowledge.<sup>12</sup> Given that adopting this conception of knowledge puts us in conflict with most of our everyday judgments regarding who is in a position to know what, this speaks in favour of reconsidering whether such a conception of knowledge accurately reflects what it takes to know anything.

## **1.2 Austin on knowledge**

In this section I consider some of J.L. Austin's observations regarding everyday epistemic practice, observations that suggest just where the sceptical conception of knowledge goes wrong. Austin (1961, p. 45) encourages us to reflect on "what sort of thing does actually happen when ordinary people are asked 'How do you know?'" And what emerges when we call to mind our

---

<sup>12</sup> It is worth noting that the conception of knowledge I have outlined is not the only conceivable conception of knowledge with sceptical consequences.

everyday epistemic and discursive practices is that the sceptical conception of knowledge does not accurately reflect the concept of knowledge as it functions within those practices. When a person claims to know something ordinarily, we don't take him to be saying that he is in a position to rule out any conceivable alternative to what he takes to be so. Nor do we ordinarily think that just any conceivable alternative can be adduced as a legitimate challenge to a person's claim to know. Rather, we share a (more or less rough) sense of which doubts can be reasonably brought to bear against someone's claim to know and in which circumstances. In some circumstances a person's inability to rule out an alternative to what he claims to know will be considered grounds for concluding that he is not in a position to know the thing in question. However, if circumstances are not such as to make it apt to take a particular alternative seriously, we won't take a person's inability to rule out the alternative to count against his claim to know. This suggests that the sceptical conception of knowledge discussed in the previous section does not accurately reflect the concept of knowledge as it functions within our epistemic and discursive practices.

Suppose Ollie has been out bird-watching and returns to tell Dan that he has seen a goldfinch. In response to this Dan might ask Ollie how he *knows* that it was a goldfinch he saw. One of Austin's key insights is that Dan's question might be construed in different ways depending on the specific circumstances in which it is asked. Dan might, for example, be interested in finding out where Ollie learned to identify goldfinches. Did he perhaps take a short course in ornithology, or grow up in an environment in which goldfinches were common? Then again, Dan's question might be understood as a request for Ollie to tell him which he features of the bird he went off in identifying it as a goldfinch. Was it the distinctive bird call that gave it away, or the colour of its plumage? In each case we have a (more or less rough) sense for what would ordinarily count as a satisfying response on Ollie's part, for what it would ordinarily take to account for his claim to know that he has seen a goldfinch. And what these responses do, among other things, is differentiate Ollie's position from that of someone who is not in a position to know that he has seen a goldfinch.

Austin goes on to observe that in some cases we are liable to be called out on our response when asked how we know something, liable to have doubts raised about whether we are in fact in a position to know what we have claimed to know. Suppose Ollie responds to Dan that he knows that what he saw was a goldfinch because it had a red head. Dan might in turn make the following objection:

But that's not enough: plenty of other birds have red heads. What you say doesn't prove it. For all you know, it may be a woodpecker. (Austin, 1961, p. 51)



This case, at least on the surface, resembles the example with which I began the previous section, the example that seemed to speak in favour of the sceptical conception of knowledge. Dan suggests an alternative to Ollie's having seen a goldfinch, namely that what he saw was in fact a woodpecker (or some other bird). Furthermore, he suggests that Ollie may not be in a position to rule this alternative out, that he may not be able to distinguish between his having seen a goldfinch and his having seen a woodpecker. How might Ollie go about defending himself against Dan's objection? One way in which Ollie might defend his claim to know that he has seen a goldfinch is to note some further feature of the bird which shows it to have been a goldfinch and not a woodpecker. For example, Ollie might note that the bird also had the distinctive eye-markings of a goldfinch. In many ordinary situations we would think this enough to defend Ollie's claim to know against Dan's objection.

That said, if Dan's charge proved to be correct, and it turned out that Ollie was unable to distinguish between his seeing a goldfinch and his seeing a woodpecker, we would ordinarily think Dan right to conclude that Ollie *wasn't* in a position to know that what he saw was a goldfinch. Austin does not, however, think that this observation regarding everyday epistemic practice speaks in favour of a sceptical conception of knowledge according to which a person must be in a position to rule out *any* conceivable alternative if he is to count as knowing something. When Ollie claims to know that he has seen a goldfinch, he is claiming to have done 'enough' to establish that it was a goldfinch. How much is 'enough' in a situation? Austin writes:

Enough means enough to show that (within reason, and for present intents and purposes) it 'can't' be anything else, there is no room for an alternative, competing, description of it. It does *not* mean, for example, enough to show it isn't a *stuffed* goldfinch. (Austin, 1961, p. 52)

Austin here suggests that in claiming to know that he has seen a goldfinch, a person need not be construed as saying that he has done enough to show that he has not instead seen a stuffed goldfinch. Suppose Ollie is not in a position to show that what he sees isn't a stuffed goldfinch – perhaps it is in a tree at some distance away. To say that Ollie is not able to show that what he sees isn't a stuffed goldfinch is to say that there is nothing available to him at this time which would allow him to distinguish his situation from one in which he is presented with a stuffed goldfinch. Such might be the case if the goldfinch has been still for the entire time that Ollie has been looking at it. In such a case Ollie's experience of seeing a goldfinch would be indistinguishable to him (at that time) from the experience he would be having if what he saw

was instead a stuffed goldfinch. He would not be in a position to rule out<sup>13</sup> his seeing a stuffed goldfinch at that time (at least not without further investigation).

We can imagine circumstances in which we would ordinarily think Ollie right to say he knows that what he sees is a goldfinch, even though he is not in a position to show that what he sees isn't instead a stuffed goldfinch. This is what I take it that Austin's remarks are supposed to remind us. But if Ollie is in no position to show that what he sees isn't a stuffed goldfinch, *mightn't* he in fact be looking at a stuffed goldfinch, for all he can tell? In that case it seems wrong to say that Ollie *knows* that what he sees is a goldfinch. It sounds absurd to say 'Ollie knows that what he sees is a goldfinch, but for all he can tell it might be a stuffed goldfinch'. How does Austin propose getting around this worry? The trick, according to Austin, is recognising that there is a difference between merely conceiving of an alternative to what Ollie takes to be so, and having reason to seriously think that Ollie might be mistaken in this way.

Austin alludes to this distinction in a number of passages throughout 'Other Minds'. For any situation in which we might take another person to be angry, we can imagine things being such that the person was merely feigning anger, and this in such a way that we were unable to distinguish the one experience from the other (at least at that time). Austin observes that the mere conceivability of such alternatives does not in itself mean that we ought seriously to suppose that a person might only be feigning anger. Rather, it is only in certain circumstances where there is some *special suggestion* that a person might be pretending that we will seriously think it apt to suppose that the person might be feigning anger. Austin writes:

These special cases where doubts arise and require resolving, are contrasted with the normal cases which hold the field<sup>1</sup> *unless* there is some special suggestion that deceit &c., is involved, and deceit, moreover, of an intelligible kind in the circumstances, that is, of a kind that can be looked into because motive, &c., is specially suggested. (Austin, 1961, p. 81)

We wouldn't normally take it to count against our claim to know that another person is angry that it is after all *conceivable* that the person could have been feigning anger, and that if he was we wouldn't have been able to tell the difference (at least at that time). Only in certain circumstances would we think it apt seriously to suppose that the person might in fact be feigning anger. And this is tied up with the fact that we don't seriously suppose that normal people will have any reason to feign anger in most situations. Feigned anger is something we only expect to encounter

---

<sup>13</sup> I should stress again that I mean something very specific by 'ruling out' in this context. Being in a position to 'rule out' an alternative to *p* means being in a position to distinguish in some way between one's experience if *p* and one's experience if the alternative instead obtained.

in certain exceptional circumstances. Unless there is something unusual about the circumstances in which we find ourselves, something to suggest that the person might in fact have some reason to feign anger, we won't in general think it apt to take the suggestion that he might be feigning seriously.

Austin makes a similar observation in relation to cases where we might ordinarily take ourselves to know something on the basis of someone else's testimony. For any case in which we might rely on another person's testimony, we can conceive of a situation in which the person is mistaken, or lying to us, indistinguishable from the situation in which what he says is true. But the mere conceivability of such situations does not by itself mean that we ought seriously to suppose that the person might in fact be mistaken or lying to us. Nor will the mere conceivability of such situations normally be taken to count against our claim to know things on the basis of the testimony of others. It is only in certain circumstances where there is some special reason to suppose that the person might be mistaken or lying that such a possibility will be taken to count against our claim to know on the basis of the person's testimony. Austin writes:

Naturally, we are judicious: we don't say we know (at second hand) if there is any special reason to doubt the testimony; but there has to be *some* reason. It is fundamental in talking (as in other matters) that we are entitled to trust others, except in so far as there is some concrete reason to distrust them. (Austin, 1961, p. 50)

Austin doesn't say exactly what he has in mind in talking about a 'concrete reason' here, but his use of the expression suggests this much: the mere conceivability of an alternative to how we take things to be does not by itself suggest we have a concrete reason to think we are mistaken, does not by itself suggest that we ought seriously to suppose we might in fact be mistaken. We can imagine situations illustrating what Austin has in mind. Suppose the station manager told me that the next train leaves at four o'clock, but Lucy says he told her that the next train leaves at five. In that case we mightn't have any immediately available means of working out which (if either) of us received the correct information. There is an alternative I have reason to take seriously, namely that the station manager may have given me the incorrect information. If I were in light of this to tell Lucy that I *know* the train leaves at four o'clock, she would expect me to have done something further to confirm that this was so. If I hadn't done anything further to confirm the details and was still relying on the station manager's word, we would think Lucy apt in pointing out that I was not yet in a position to know that the train leaves at four o'clock.

We can contrast this with a situation in which Lucy hasn't herself been told anything by the station manager. Suppose that returning from my discussion with the station manager I tell Lucy that the

train leaves at four o'clock. If Lucy suggests that I don't know the train leaves at four o'clock, and this because the station manager might have been mistaken, I would ordinarily be within my right to ask Lucy what reason we have to suppose that the station manager might have been mistaken. And generally we would expect Lucy to do something more than to merely say that it is possible that the station manager could be mistaken, or to merely tell some coherent story about how the station manager might have come to give me the incorrect information. That someone in Lucy's position can tell a coherent narrative to this effect is not surprising; nor would we in every situation take her ability to tell such a story to count against my claim to know that the train leaves at four o'clock. Unless we have some reason to seriously suppose that Lucy's story might in fact be true, we wouldn't take it to count against my claim to know that the train leaves at four o'clock. What this example illustrates is the fact that that we take one another to be entitled to claim to know things on authority unless there is a reason to seriously suppose the information might be wrong. We don't in general expect that in claiming to know something on the basis of authority, a person is claiming to be in a position to rule out any conceivable alternative.

That we can in some circumstances be misled by appearances into making false judgements about what is so is one aspect of our fallibility. For any policy of belief-formation we can imagine circumstances in which the policy will lead us astray. We ordinarily think it good practice to judge that one has seen a goldfinch when in good lighting what appears to be a goldfinch lands in a nearby tree. Such a policy won't lead us astray if there is nothing in our immediate environment which might in good lighting be mistaken for a goldfinch. But such a policy may lead us astray if we are in an environment rife with robotic goldfinches. The question, then, is whether we have any reason to suppose that our immediate environment might be different in this respect from how we take it to be. If we have reason to take seriously the idea that there are or might in fact be robotic goldfinches in the vicinity, then we have reason to seriously suppose that someone might be mistaken when he judges that something is a goldfinch. But where we don't take ourselves to have any such reason, we don't ordinarily take the suggestion that a person could be looking at a robotic goldfinch to count against his claim to know that he has seen a goldfinch.

Austin recognises that human judgment is "inherently fallible and delusive". He also wants to maintain that "'When you know you can't be wrong' is perfectly good sense." These are two commitments which are liable to appear in conflict with one another. Austin's strategy for reconciling these two commitments is as follows:

'When you know you can't be wrong' is perfectly good sense. You are prohibited from saying 'I know it is so, but I may be wrong', just as you are prohibited from saying 'I promise I will, but I

may fail'. If you are aware you may be mistaken, you ought not to say you know, just as, if you are aware you may break your word, you have no business to promise. But of course, being aware you may be mistaken doesn't mean merely being aware that you are a fallible human being: it means that you have some concrete reason to suppose that you may be mistaken in this case. Just as 'but I may fail' does not mean merely 'but I am a weak human being' (in which case it would be no more exciting than adding 'D.V.'): it means that there is some concrete reason for me to suppose that I shall break my word. It is naturally *always* possible ('humanly' possible) that I may be mistaken or break my word, but that by itself is no bar against using the expressions 'I know' and 'I promise' as we do in fact use them. (Austin, 1961, p. 66)

This last point is particularly important. Austin is interested in how we use expressions like 'I know', how such expressions function in practice. And he is particularly interested in reminding his reader that the mere fact that it is always possible that we might be mistaken when we judge that something is so is no bar to our using the expression 'I know' when and as we ordinarily do. It is a feature of our everyday epistemic practice that a person ought not to say he knows something if he is aware he might be mistaken. But as Austin points out, there is a substantial difference between being aware that it is *possible* that we might be mistaken and being aware that we might in fact be mistaken about something. To highlight conceivable alternatives we would be unable to distinguish from how we take things to be is merely to remind ourselves that it is possible that we might be mistaken. But being aware that we might in fact be mistaken means having some reason to seriously suppose that a particular alternative might obtain. In the absence of any such reason we will ordinarily take one another to be entitled to say we know.

In the previous section I considered a conception of knowledge according to which a person only counts as knowing something if he is in a position to rule out any conceivable alternative to what he takes to be so. It is a conception of knowledge according to which we would rarely (if ever) be in a position to know anything. In this section I have briefly presented Austin's observations regarding our everyday epistemic and discursive practices as they relate to the concept of knowledge, observations which suggest that the sceptical conception of knowledge does not provide an accurate characterisation of the concept as it functions within these practices. Austin's observations suggest that ordinarily when we say that a person knows something, we are not saying that he is in a position to rule out any conceivable alternative to that which he takes to be so. Rather, we are saying that he has done enough for present intents and purposes to establish the truth of the matter. A person's inability to rule out a certain alternative can in some situations count against his claim to know that something is so, but only where we think it reasonable to take seriously the suggestion that he might be mistaken in just this way. Where we have no reason

to seriously suppose that a particular alternative might obtain, we won't take a person's inability to rule out this alternative to count against his claim to know.

### **1.3 Stroud on Austin and appeals to everyday epistemic practice**

According to the sceptical conception of knowledge considered in 1.2, to say that a person knows that  $p$  is to say that he is in a position to rule out any conceivable alternative to  $p$ . Austin's observations regarding everyday epistemic and discursive practice suggest that the sceptical conception of knowledge is not an accurate characterisation of the concept as it functions within these practices. Austin's observations instead suggest that when we say that a person knows that  $p$  we are to be understood as saying that he has done enough for present intents and purposes to establish that  $p$ . This does not require that the person be in a position to rule out any conceivable alternative to  $p$ . Rather, a person's inability to rule out a certain alternative will only be taken to count against his claim to know in circumstances where we have reason to take seriously the alternative in question.

In this section I consider Barry Stroud's influential critique of Austin's work on knowledge.<sup>14</sup> Stroud argues that Austin's observations about everyday epistemic and discursive practice cannot be taken to provide any direct illumination regarding what it takes to know something. Nor, he argues, should Austin's observations be taken to show that the sceptical conception of knowledge is incorrect. Austin, he says, furnishes us with observations regarding when it would ordinarily be *appropriate* to say that a person does or doesn't know something. Stroud is prepared to grant everything Austin says in this regard.<sup>15</sup> But the fact that it is appropriate in certain circumstances to say that someone knows does not entail that what is said is true. Stroud takes this to mean that for all Austin's observations regarding everyday epistemic practice show, the sceptical conception of knowledge might in fact be correct. Stroud suggests that although this would mean that we are speaking falsely any time we say that someone knows something, it might be that for practical purposes we nonetheless think it appropriate to say such things. In the next section I will argue that Stroud has not done enough to show that Austin's observations shouldn't be taken at face value as illuminating what knowing requires of us.

---

<sup>14</sup> Stroud's full critique of Austin can be found in chapter 2 of *The Significance of Philosophical Scepticism*. (1984)

<sup>15</sup> Stroud (1984, p. 53) says that he is prepared to "grant everything Austin says about what sort of thing does actually happen when ordinary people are asked 'How do you know?', and everything else that could be discovered about how we respond to the questions or would-be challenges of others with respect to our knowledge."

In criticising Austin's approach, Stroud doesn't take issue with any of Austin's observations regarding what it is appropriate to say and do by the lights of everyday epistemic practice. In fact Stroud provides a number of examples intended to further support Austin's observations. Stroud imagines he is at a party and returns from the garden to inform the other guests that he has just seen a goldfinch. Were Stroud to be asked how he knows, we would expect him to set about adducing the kinds of considerations Austin discusses in 'Other Minds'. We would, for example, expect him to say something about how he has come to be in a position to recognise goldfinches when he sees them, or about which features of the bird allowed him to successfully pick it out as a goldfinch, or something else establishing his credentials on the matter. Stroud (1984, p. 49) then asks us to suppose that one of the guests goes on to say 'That's not enough – for all you can tell you might have been dreaming'. Said like this and without any further pretext, Stroud agrees that we wouldn't in such circumstances take what this guest says to count against Stroud's claim to know that he has seen a goldfinch. In the absence of any special reason to suppose that Stroud might have been dreaming, we wouldn't take the guest to have given us any reason to seriously suppose that Stroud has not done enough for present intents and purposes to establish that he has seen a goldfinch, nor would we take the guest to have given us any reason to seriously suppose that Stroud might be mistaken about this.

Stroud also agrees with Austin that we can readily imagine the kinds of circumstances in which we *would* think that the suggestion that Stroud was dreaming provided us with reason to seriously suppose that Stroud might be mistaken in taking something to be so. Stroud (1984, p. 50) considers a situation in which he is lying half-awake in bed in the morning and it seems to him as though someone is calling his name. In a situation like this we would think it apt for Stroud to take seriously the possibility that he might merely have dreamed that someone was calling to him, and so to withhold from concluding that someone was calling to him. Or consider again a situation in which Stroud is at a party and claims to know that there is a goldfinch in the garden. Suppose this time that it comes to light that Stroud has a particularly severe form of narcolepsy, one which causes him intermittently to fall into a deep sleep in which he has rather vivid dreams. Furthermore, suppose that he has been known sometimes to have difficulty distinguishing things he has dreamed from things that have actually happened.<sup>16</sup> In these fairly specific circumstances we might be prepared to take seriously the suggestion that Stroud might have been dreaming, and to take Stroud's inability to rule out this possibility<sup>17</sup> to count against his claim to know.

---

<sup>16</sup> This example is mine, though I take it that Stroud would think it apt.

<sup>17</sup> In these specific circumstances we might think it enough to rule out the particular alternative being suggested if Stroud tells us that he has taken his medication and won't be falling asleep any time soon, or if he corroborates his

Stroud grants Austin's observations regarding when we would ordinarily think it right to say that a person knows something. He also grants Austin's point that ordinarily we don't think that just any conceivable alternative can be adduced to challenge a person's claim to know something. Rather, in many contexts we won't think it apt to take seriously an alternative unless there is some special reason to suppose that the person might be mistaken in just this way. This would seem to show that the sceptical conception of knowledge does not accurately reflect our understanding of what knowing in general requires of us. Stroud, however, argues that none of these and like observations should be taken to provide any illumination regarding what it takes to know something. This, says Stroud, is because all of these observations pertain only to what it is *appropriate* to say in particular circumstances, and there need be no connection between its being appropriate to say that someone knows and its being true to say he knows.

Care is needed in making out Stroud's suggestion that Austin's observations pertain only to the appropriateness-conditions for ascriptions of knowledge and tell us nothing directly about what it actually takes to know anything. At points in his discussion Stroud suggests that his argument turns on the mere fact that we can imagine circumstances in which an appropriately-made ascription of knowledge turns out to be false. For example, Stroud (1984, p. 58) considers a case in which he tells the host of a party that he knows their mutual friend Jones will be there, having just spoken to Jones on the phone. Stroud notes that although we would ordinarily think him warranted in claiming to know that Jones will be at the party, it is conceivable that Jones might nonetheless fail to arrive. In that case Stroud couldn't have known that Jones would be at the party. After all, most everyone agrees that  $p$  must be so if a person is to know that  $p$ . In this case Jones wasn't at the party, so Stroud couldn't have known that he would be at the party. The moral Stroud (1984, p. 59) wants us to draw from the example: there can be situations in which a knowledge-claim that is "justified, reasonable, and appropriate in the circumstances" nonetheless turns out to have been false. In other words, the mere fact that it is appropriate to say something in a situation doesn't entail that what is said is true.

If this was all Stroud's critique came to, it is difficult to see what Austin might have to fear from it. Austin never suggests that any time we reasonably and appropriately claim to know something, this entails that we do in fact know the thing in question. Austin, recall, thinks that it is in some sense *always* possible that we might be mistaken any time we take something to be so. This presumably holds as much for cases in which we take ourselves to know things as when we take

---

story with someone else who was there. In that case we might think it apt to grant Stroud's claim to know, provided there is no further reason to take seriously the suggestion that Stroud might be mistaken.



anything else to be so. Indeed, Austin (1961, p. 66) notes that if his observations regarding everyday epistemic practice are correct, “we are often right to say we *know* even in cases where we turn out subsequently to have been mistaken”. Austin clearly doesn’t think this observation poses any problem for his approach to reflecting on knowledge. The mere fact that we sometimes turn out to have been mistaken when we take ourselves to know doesn’t suggest that we are always – or even very often – mistaken in taking ourselves to know things.<sup>18</sup>

The example does, however, highlight something important about Austin’s approach to thinking about knowledge. Unlike some other philosophers, Austin is not interested in providing a non-circular analysis of the necessary and sufficient conditions for knowing something.<sup>19</sup> In the time since Edmund Gettier (1966) published his paper problematising the ‘justified true belief’ analysis of knowledge, no one has been able to provide an analysis immune to counter-examples.<sup>20</sup> That such an analysis hasn’t been successfully formulated is only worrying if we think that concepts should in general be amenable to this kind of analysis. To eschew such analysis is not to give up on there being anything illuminating to be said about the concept of knowledge. Indeed, Austin can be understood as providing a template for an alternative approach. Rather than attempting to state necessary and sufficient conditions for knowledge, Austin asks us to call to mind the kinds of circumstances in which we would ordinarily think that someone was to be counted as knowing something. Even if we can imagine further considerations that might lead us to retract the judgment, this doesn’t mean that the exercise failed to provide illumination regarding what it takes to count as knowing something in particular circumstances.

As I have said, at least some of Stroud’s examples suggest that his criticism of Austin turns on the fact that it is always conceivable that where we have judged that someone knows something, further considerations may arise which would lead us to retract the judgment. That this is so does not immediately suggest that we ought not to take Austin’s observations regarding everyday epistemic practice to illuminate what it takes to know something. Stroud, however, means to be suggesting something more than the mere fact that we can in particular circumstances prove to have been mistaken in taking ourselves to know something. He goes on to argue that what he sees as a logical gap between appropriateness-conditions and truth-conditions makes room for the possibility that something like the sceptical conception of knowledge is correct, and this in

---

<sup>18</sup> Austin (1961, p. 66): “The human intellect and senses are, indeed, *inherently* fallible and delusive, but not by any means *inveterately* so.”

<sup>19</sup> The *locus classicus* of this approach to thinking about knowledge is Roderick Chisholm’s (1989) *Theory of Knowledge*.

<sup>20</sup> See chapter 1 of Timothy Williamson’s (2000) *Knowledge and its Limits* for one prominent rejection of the idea that analysis should be our goal in studying the concept of knowledge.

spite of Austin's observations regarding how the concept of knowledge functions within everyday epistemic practice.

According to the sceptical conception of knowledge, a person only counts as knowing that  $p$  if he is in a position to rule out any conceivable alternative to  $p$ . If right, such a conception of knowledge would carry the implication that no one is ever in a position to know anything, or at least that the circumstances in which we can truly be said to know anything are radically fewer than we might ordinarily suppose. Stroud thinks this might very well be the case. But if this is so, what might explain the widespread discrepancy between the actual requirements for knowledge and our everyday views regarding who counts as knowing what? Stroud's suggestion: practical exigency. He suggests that practical considerations might lead us to treat it as appropriate to say that we know things even though strictly-speaking we are never in a position to know anything.<sup>21</sup>

Stroud's suggestion is that our everyday judgments regarding who is in a position to know are made under (what he thinks may well be) the distorting influence of practical considerations. We realise that for practical purposes things would become unmanageable if we only recognised people as knowing in situations where they are in a position to rule out any conceivable alternative to that which they take to be so. However, Stroud thinks that on at least one plausible understanding of what philosophical reflection on knowledge calls for, we are to consider the question of who is in a position to know what in a completely detached manner, putting aside any immediate practical interest we might have in the question of whether or not a person knows. It is only in such circumstances, says Stroud, that we can come to appreciate what is necessitated by the concept of knowledge itself. And when reflecting in this way on the question of what it takes to know, many epistemologists (including Stroud) have come to feel that the sceptical conception of knowledge is intuitively compelling. As Stroud notes, when reflecting in this detached manner, it is no longer clear to him that a person needn't be in a position to rule out the (admittedly outlandish) possibility that he is, for example, dreaming in order to count as knowing anything about the world. On the contrary, Stroud comes to find it plausible that a person must really must be in a position to rule out this alternative if he is to count as knowing anything. And he thinks the same goes for any other conceivable alternative we can dream up in a situation.

---

<sup>21</sup> "The practical social purposes served by our assertions and claims to know things in everyday life explain why we are normally satisfied with less than what, with detachment, we can be brought to acknowledge are the full conditions for knowledge." (Stroud, 1984, p. 71)

Stroud suggests that the sceptical conception of knowledge might reflect what knowing actually requires of a person, even though we wouldn't ordinarily expect a person to be in a position to rule out any conceivable alternative before taking him to know something. He also suggests that practical exigency can explain why, having realised what knowledge really requires of us, we don't modify our everyday epistemic practices accordingly. We don't modify our practices to fit with the sceptical conception of knowledge because we recognise that our practices would become unmanageable if we were to adhere to such a conception of what knowledge requires. Having sketched this possibility, Stroud concludes that Austin's observations cannot without further argument be taken to establish anything about what knowing requires of a person. For Austin's observations to carry any implications about what knowing requires of us, it needs to be shown, according to Stroud, that the conditions under which we think it appropriate to ascribe knowledge to a person correspond with the conditions under which a person actually counts as knowing. And he thinks that no amount of observations regarding what it is ordinarily appropriate to say and do can show this. This is because its being appropriate in a situation to say that a person knows carries no implication that something like the sceptical conception of knowledge mightn't be correct.

Stroud's criticism, if right, would suggest that Austin is looking in the wrong place for illumination regarding what kind of position a person needs to be in if he is to count as knowing something. Austin's observations, as Stroud paints them, furnish us only with considerations regarding the circumstances under which it is appropriate to say that a person knows something. They do not tell us anything about what it takes to actually count as knowing something in a situation. It is possible, Stroud thinks, that for all Austin's observations regarding everyday epistemic and discursive practice, the sceptical conception of knowledge may provide an accurate characterisation of the kind of position a person needs to be in if he is to actually count as knowing something. In the next section I argue that Stroud's suggestion isn't in fact plausible, and that he hasn't succeeded in showing that Austin is looking in the wrong place for illumination regarding what knowing requires of us. In painting Austin as being interested *merely* in 'appropriateness-conditions' for knowledge attributions, Stroud mischaracterises Austin's approach. Austin calls our attention to what in particular circumstances a speaker can reasonably be understood to be saying in saying that a person knows something. Stroud has given us no good reason to suppose that what such considerations reveal is anything other than our understanding of what it takes to know in particular situations.

#### **1.4 Stroud miscasts Austin's approach**

Stroud casts Austin as being merely interested in the conditions under which we would ordinarily think it appropriate to say that a person knows something, an approach Stroud argues cannot provide any immediate illumination regarding what knowing involves. This, however, is a misleading characterisation of Austin's philosophical approach. Austin is interested in what we can reasonably be understood to be saying on different occasions for ascribing knowledge to a person, how we can reasonably be held responsible for things being, given what we have said. And he thinks, quite plausibly, that such reflections reveal to us our understanding of what it takes in such circumstances to count as knowing something. Stroud has given us no good reason to suppose that we should think otherwise. Austinian considerations should thus be our first port of call when reflecting on what it takes to count as knowing something in particular circumstances.

Stroud suggests that Austin's observations regarding who we are inclined to count as knowing in various situations do not furnish us with any direct insight into what it takes to know in those circumstances. All these observations establish, according to Stroud, is what we are prepared for practical purposes to let people *pass* for knowing. This, he thinks, is to be contrasted with what we are inclined to think knowing actually requires of a person once we have freed ourselves of any immediate practical concerns and attend to the matter. And Stroud thinks it at least not implausible that we might come to discover that what we are saying is false whenever we ordinarily say that people know things. But the suggestion that we are speaking falsely when we say that a person knows only makes sense on the assumption that what we are saying in those circumstances is to be understood in the way suggested. This is what Austin encourages us to call into question.

What Austin furnishes us with is a method for calling to mind how a speaker in particular circumstances is reasonably to be understood in saying that someone knows. Suppose Dan tells Elle that there is a goldfinch in the garden. Elle is aware that Ollie will be disappointed if he misses it, so she asks Dan if Ollie knows. Dan tells her that Ollie already knows. How is Dan to be understood in this situation? Is Dan to be understood as saying that either he or Ollie has done something specifically to show that the 'goldfinch' is not in fact a robotic imitation? If the circumstances were such that Dan and Elle were aware of some reason to seriously suppose there might be a robotic goldfinch in the garden, Dan could reasonably be understood to be saying this. But suppose the circumstances are such that we wouldn't ordinarily expect either of them to take seriously the suggestion that it might have been a robotic imitation. In that case Dan couldn't reasonably be understood as saying that either he or Ollie has done anything to specifically rule

out this alternative. Were Elle to understand Dan in this way in these circumstances, she would seriously have miscomprehended what Dan was telling her in saying that Ollie knows. The same point can be made in relation to other conceivable alternatives.

What Dan is reasonably to be understood as telling Elle in these circumstances is that Ollie has done *enough* in this situation to establish that there is a goldfinch in the garden. There are different things Ollie might have done which would count as enough in these circumstances. He might, for example, have seen the bird himself, or he might have been informed of the bird's presence by Dan. If Elle were interested in which of these positions Ollie is in, she might ask Dan *how* Ollie knows that the goldfinch is in the garden. The point is that we recognise more or less what kind of position Ollie needs to be in if things are as Dan can reasonably be understood as having said them to be. Given that what Dan says in this situation is that Ollie *knows*, Ollie's being in one of these recognisable positions is plausibly part of our understanding of what it takes to know in such circumstances.

Stroud, however, thinks that this last move is unwarranted. Not that he wants to deny what I have said about how Dan can reasonably be understood in the circumstances described.<sup>22</sup> He recognises that Dan couldn't reasonably be understood as saying that either he or Ollie had done anything to establish that what they saw wasn't a robotic goldfinch, at least not in circumstances where there was no reason to seriously suppose it might have been a robotic bird. Nonetheless Stroud thinks it at least possible that this is how the *words* Dan uses (as opposed to *Dan*) ought to be understood. And it is, he thinks, our understanding of what the *words* say, as distinct from what Dan might reasonably be understood to be saying in using the words, that we should be attending to if we want to understand what it actually takes to count as knowing something.

But why should we suppose that the words Dan uses are to be understood as saying *anything* apart from what a speaker in Dan's circumstances could reasonably be understood as having said in using them? And if we understand *Dan* to be saying something true in such circumstances, why suppose that there is any further question to be asked about whether the *words* Dan uses 'say' something that is true or false? When Dan tells Elle that Ollie already knows there is a goldfinch in the garden, we have a reasonable idea of what kind of position Dan is saying Ollie to be in. That is, we have a reasonably good idea of what kind of position Ollie needs to be in if *Dan* is to have said something true. In the absence of there being any reason to seriously suppose that the

---

<sup>22</sup> Stroud (1984, p. 74) admits that our everyday ascriptions of knowledge may be "marking a real difference" in a person's position, he just doesn't think that this difference we are marking need amount to the difference between knowing and not knowing in particular situations.

bird might have been a robotic imitation, we don't think that either Dan or Ollie need to have done anything specifically to rule out this alternative in order for *Dan* to have said something true of Ollie. The same is true of other situations in which we might ordinarily think that someone knows something. In each case we can make reasonably good sense of what a person in such circumstances says to be so in saying that someone knows. The question is why we should suppose that there is any further question to be asked about whether the words themselves 'say' something true or false, when contemplated in a 'detached' manner.

Stroud suggests that our speaking 'loosely' ordinarily might account for the kind of distinction he has in mind between how *Dan* can reasonably be understood and how his *words* are to be understood. The thought would be that Dan *ought* not to say that Ollie knows if he was properly considering the matter, and this because he should in that case realise that the standards for knowledge are actually much 'stricter' than he is prepared to countenance at other times. The question is whether this is at all a plausible suggestion. Generally when we talk about a person's having spoken loosely, or having said that someone knows in a situation where they 'strictly' ought not to have, this will be because there is some possibility of error we think the speaker *ought* to have taken into account. It is not clear that the example involving Dan is any such case.

Ought Dan strictly to have taken into consideration the suggestion that Ollie might in fact have seen a robotic imitation when judging that Ollie knew there was a goldfinch in the garden? Reflecting on the example, it is not clear why this should be our reaction. After all, as the story has been told, there is no reason at all to seriously suppose that there might have been a robotic goldfinch in the garden. So why should we think that *anyone* (including us considering the example) ought to take seriously the suggestion that Dan and Ollie might have seen a robotic bird? Nothing in our epistemic practices compels us in this direction. Austin reminds us that ours is a practice within which we only think it apt to take certain alternatives seriously where we have some special reason to do so. And in the example as presented, there is nothing to suggest that anyone ought to take seriously the suggestion that Ollie might in fact have seen a robotic imitation.

Suppose we imagined the example so that, unbeknownst to Dan, the military have recently introduced a number of robotic goldfinches into the area.<sup>23</sup> Now we might feel that Dan says something false when he says that Ollie knows there is a goldfinch in the garden. We might think that, given Dan is unaware of the situation with the robotic birds, Dan has still spoken reasonably and appropriately in saying that Ollie knows. Nonetheless, we are liable to feel that what Dan says

---

<sup>23</sup> This example is deliberately meant to echo Stroud's (1984, pp. 67-69) famous example (borrowed from Thompson Clarke) involving the plane-spotters.

here is false. The presence of robotic birds in the region means that there is reason to seriously suppose that Dan and Ollie might have been mistaken about the supposed presence of the goldfinch in the garden. But this makes perfect sense by the lights of our everyday epistemic practice as described by Austin.<sup>24</sup> After all, we are aware of a special reason for thinking that Ollie and Dan might in fact have seen a robotic imitation. If Ollie and Dan were aware of this reason, we would expect them to likewise agree that neither was in a position to know that there was a goldfinch in the garden. This, however, hardly gives us reason to think that a sceptical conception of knowledge might be correct. It at best illustrates the fact that we can sometimes turn out to be mistaken about which alternatives we have reason to seriously suppose might affect our claims to know. It doesn't suggest that we would be wrong to dismiss the suggestion that Ollie might have seen a robotic imitation as irrelevant to the question of whether Ollie knows he has seen a goldfinch in other circumstances.

Imagine instead that both Dan and Elle are aware that the military has released robotic goldfinches into the region. However, given the difficulty there is in distinguishing these robotic birds from actual goldfinches, and the relative unimportance of being mistaken on the matter, Dan and Elle continue for their purposes to ignore the possibility of robotic birds when judging whether there is a goldfinch in the garden.<sup>25</sup> In that case Elle might take Dan to be speaking reasonably and appropriately when he says that Ollie 'knows' that he has seen a goldfinch, even though both will on reflection agree that Ollie 'strictly' isn't in a position to know whether he has seen a goldfinch unless he can rule out its having seen a robotic bird. Now it looks like we have something closer to what Stroud has in mind. Dan and Elle in this example are happy for practical purposes to employ standards for saying that Ollie 'knows' that are looser than what they (and we considering the example) would consider to be the actual requirements for knowledge.<sup>26</sup>

Stroud suggests that such cases give us reason to be wary about taking Austinian considerations at face value, but in fact it is hard to see why we should think this. After all, it is a key feature of this example that Dan and Elle have a special reason to suppose that what appears to be a goldfinch might in fact be a robotic imitation. This accounts for their (and our) inclination to

---

<sup>24</sup> I am indebted to Mark Kaplan (2000) for this general line of response to examples like this and Stroud's own plane-spotter example.

<sup>25</sup> In Stroud's structurally similar example, there is a type of plane (Gs) that is difficult to distinguish from one the plane-spotters have been trained to identify (Fs), and of relatively no importance for their purposes. So the plane-spotters go on calling any such planes 'Fs', and think it appropriate for their purposes to say that they know when they have seen an F, even though it might well have been a G.

<sup>26</sup> I am following Stroud in his interpretation of such examples here, but it isn't clear that we should interpret such examples as being cases in which 'knows' is used 'loosely'. We could just as easily interpret such cases as being ones in which 'knows' is used as it usually is, but in which the term 'goldfinch' has come to take on a different significance for practical purposes in such cases (i.e. to include robotic goldfinches).

hold that Dan is not strictly saying something true when he says that Ollie knows, even if they can understand what *Dan* would mean for practical purposes. But this is nothing like the original case we considered, one in which neither Dan nor Elle (nor us considering the example) have any reason to seriously suppose that there might be robotic goldfinches about and that Ollie might thus have been mistaken in just this way. So while we can make sense of Stroud's suggestion that there might be cases in which practical exigency leads us to speak 'loosely' when it comes to ascribing knowledge, Stroud has not given us any reason to seriously suppose that *every* case in which we say that someone knows something should be understood in this way. The kinds of cases that would illustrate Stroud's idea are ones that Austin's observations allow us to make sense of. They are cases in which we would take the people in question to have some special reason to seriously suppose that a certain alternative might in fact obtain, but in which they are happy to turn a blind eye to the alternative for practical purposes. That we can imagine such cases does not yet give us any reason to take seriously the suggestion that *every* case in which we say that a person knows is one in which we speak 'loosely'. The kinds of cases Austin would have us attend to are not ones where we would think the speakers involved (or us attending to the example) have any reason to take seriously certain alternatives for the purposes of ascribing knowledge. Regarding such cases, Stroud has given us no good reason to suppose that our understanding of what a speaker could reasonably be taken to be saying in saying that a person knows doesn't reflect our understanding of what it takes to count as knowing in such circumstances.

In short, Stroud hasn't given us any good reason to think that we shouldn't take the kinds of considerations Austin adduces at face-value as illuminating what it takes in different circumstances to count as knowing something. Austin furnishes us with examples reminding us what we would in different circumstances understand ourselves to be saying in saying that a person knows something. Stroud's suggestion that Austin is interested merely in when it is appropriate to say someone knows obscures this point. The point of calling to mind everyday situations in which we might naturally say that someone knows something is to get us to appreciate what a speaker in such circumstances could reasonably be understood to be saying. Insofar as we can make sense of what a speaker could in such circumstances reasonably be understood to be saying, *how* he is saying things to be, such considerations can plausibly be taken to be indicative of our understanding of what it takes to count as *knowing* in such circumstances.



## **1.5 Summary**

In this chapter I introduced Austin's approach to reflecting on the concept of knowledge, an approach that involves reflecting on how the concept functions in everyday epistemic and discursive practice. Austin encourages us to call to mind particular circumstances in which we might normally say that a person knows something, and to consider what a person might reasonably be understood to be saying in such circumstances. Such considerations serve to illuminate our understanding of what kind of position a person needs to be in if he is to count as knowing in particular situations. Austin's observations regarding everyday practice suggest that in order to count as knowing that  $p$ , a person needn't be in a position to rule out any conceivable alternative to  $p$ . Compelling though such a proposal might be in the abstract, it would have the absurd implication that no one is ever in a position to know anything. Attending to the kinds of considerations regarding everyday practice that Austin adduces allows us to see where such a proposal goes wrong. As Austin reminds us, a person's inability to rule out a certain alternative is only taken to count against a person's claim to know if there is some special reason to take seriously that alternative in the particular circumstances.

I considered Barry Stroud's influential criticism of Austin to the effect that his observations fail to provide any direct illumination regarding what it takes to know something. Stroud argues that Austin's approach merely furnishes us with considerations regarding the conditions under which we might ordinarily think it appropriate to say that someone knows something, and that such considerations cannot be taken to directly establish anything about what it takes to actually know something. Stroud suggests that for all Austin's observations show, a person might only count as knowing that  $p$  if he is in a position to rule out any conceivable alternative to  $p$ . Stroud's suggestion that Austin is merely interested in 'appropriateness-conditions' for ascriptions of knowledge as opposed to 'truth-conditions' obscures the point of Austin's approach. Austin is interested in getting us to recall what a speaker might reasonably be understood to be saying in saying that someone knows in particular situations. And nothing Stroud says gives us any reason to suppose that our understanding of what a speaker could reasonably be taken to be saying in such circumstances does not reflect our understanding of what it takes to know in such circumstances. If we want to come to better terms with what the distinction between knowing and not knowing comes to in any particular situation, careful attention to what we can reasonably be understood to be saying in saying that a person knows remains a valuable source of insight.

## Chapter 2. Austin, relevant alternatives, and a contextualist account of knowledge

In this chapter I argue that Austinian considerations speak in favour of a contextualist account of knowledge. According to a contextualist account of knowledge, knowledge-ascribing sentences can support a variety of understandings. The same person might count as knowing that there is a goldfinch in the garden on one understanding of what knowing that thing amounts to, while failing to count as knowing on another understanding of what knowing that thing amounts to. Whether or not a speaker says something true in saying that Ollie ‘knows’ that there is a goldfinch in the garden thus depends not merely on Ollie’s position and what he is said to know, but also on how the speaker is reasonably to be understood given the particular circumstances in which he ascribes knowledge to Ollie. A speaker in one set of circumstances might speak truth in saying that Ollie knows, while another speaker in a different set of circumstances might say something false of Ollie in saying that he knows, and this with no change in *Ollie’s* position. To illustrate this point I will be making use of the notion of a ‘relevant alternative’. A relevant alternative to  $p$  is one that the person must be in a position to rule out if he is to count as knowing that  $p$ . An alternative that is not relevant on one understanding of what it is to know that  $p$  may nonetheless be a relevant alternative on another understanding of what it is to know that  $p$ . Whether an alternative is to be considered relevant or not depends on the particular circumstances in which a speaker ascribes knowledge to a person.

I begin by introducing the notion of a relevant alternative, drawing on Austin’s observations regarding everyday ascriptions of knowledge. A relevant alternative to  $p$  is an alternative a person would need to be in a position to rule out in order to count as knowing that  $p$ . The notion of a relevant alternative can be used to articulate the kind of position a person could reasonably be taken to be in, having been said to know something by someone. As Austin notes, there are everyday occasions on which we wouldn’t think it reasonable to expect that a person was in a position to rule out his having seen a stuffed goldfinch instead of a goldfinch in being said to know that he had seen a goldfinch. This is in part because we couldn’t reasonably expect someone in those circumstances to seriously suppose that what the person had seen might in fact have been a stuffed goldfinch, not in the absence of some special reason to take that suggestion seriously. Nonetheless we can imagine circumstances in which we would take a speaker to be saying that the person is in a position to rule out his having seen a stuffed goldfinch in saying that

he knows that what he sees is a goldfinch. These will be circumstances in which the speaker could reasonably be expected to take seriously the suggestion that the person *might* have seen a stuffed goldfinch instead of a goldfinch (even if what the person saw was in fact a goldfinch).

Next I present two pairs of contrasting cases. In each pair of cases the same person in the same position is imagined to have been ascribed knowledge in two different sets of circumstances. In the one case it looks as though the speaker says something true in saying that the person knows, whereas in the other case it looks as though the speaker says something false in saying that this same person knows. The key difference between the two cases is that an alternative which appears to be irrelevant in one of the cases looks to be relevant in the other case. Put otherwise, it looks as though the speaker in each case can be reasonably understood as presenting the person as being in a different position. But in each case what the speaker says is that the person *knows*. I argue that this speaks in favour of a contextualist account of knowledge according to which there can be different understandings of what it would take for a person to count as knowing something. On one understanding of what it is to know the thing in question a person might count as knowing, even though he wouldn't count as knowing on another understanding of what it is to know that thing. Two speakers in different circumstances might thus be understood as saying different things of a person in saying that he 'knows'. In other words a contextualist can hold that a speaker in one set of circumstances says something true when he ascribes knowledge to a person, even if a speaker in another set of circumstances says something false in ascribing knowledge to this same person.

Having presented a contextualist interpretation of the examples, I then consider the prospects for what is generally referred to as an 'invariantist' interpretation of the cases. An invariantist is someone who holds that if a speaker says something true in saying that a person knows something, then it is not possible for another speaker to say something false in saying this same person to know. Likewise, an invariantist will hold that if a speaker says something false in saying that a person knows something, it is not possible for another speaker to say something true in saying that the person knows. This is because an invariantist does not allow that there can be different understandings of what a person's knowing something might come to in a situation, and that a person might count as knowing on one understanding of what it is to know, while not so counting on another understanding. I argue that the invariantist faces a number of hurdles in making sense of the contrasting cases, and that the most immediate strategies for getting around these hurdles lead to some problematic commitments. The contextualist's ability to make better sense of the examples discussed provides *prima facie* reason for adopting a contextualist stance on

the matter. I end by suggesting that the intimate link between our concept of knowledge and our testimonial practices provides an avenue for further investigating the context-sensitivity of ascriptions of knowledge.

## **2.1 Knowledge ascriptions and relevant alternatives**

In the previous chapter I introduced Austin's methodological approach to reflecting on what it takes to count as knowing something in a situation. Austin encourages us to attend to circumstances in which we might ordinarily say that a person knows something. The point of this exercise is to get us to consider how we could reasonably expect to be understood in such situations, what kind of position the person might reasonably be taken to be in. Our intuitions about this can plausibly be taken to reflect our understanding of what it takes to count as knowing something in those circumstances. In this section I focus on one of Austin's key observations regarding everyday epistemic and discursive practice, namely that in saying that a person knows that  $p$  a speaker cannot in general be reasonably understood as saying that the person is in a position to rule out just any conceivable alternative to  $p$ . Rather, we will only expect a person to be in a position to rule out an alternative if we have some reason to seriously suppose that the alternative might in fact obtain. I will refer to alternatives a person needs to be in a position to rule out if he is to count as knowing something in a situation as *relevant alternatives*.<sup>27</sup> In this section I will illustrate how the idea of a relevant alternative can be used to articulate our understanding in different situations of what kind of position a person who is said to know can reasonably be taken to be in.

Suppose Dan comes in from the garden to tell Ollie that he has just seen a goldfinch. Ollie responds that he already knows that there is a goldfinch in the garden. What kind of position could we reasonably expect Ollie to be in, given that he has said that he knows there is a goldfinch in the garden in these circumstances? We might suppose that he has already been in the garden and seen the bird for himself, or that someone else has already told him about the goldfinch (perhaps Elle, who was also in the garden). These are both recognisable ways of coming to know that there is a goldfinch in the garden in this kind of situation; in both cases Ollie would generally be taken to have done enough for present purposes to establish that there was a goldfinch in the garden.

---

<sup>27</sup> The term 'relevant alternative' comes from Fred Dretske's (1970) 'Epistemic Operators'. I am not the first person to use the notion of a 'relevant alternative' to articulate Austin's views about knowledge. For example, Krista Lawlor (2013) presents Austin as having provided us with a 'reasonable alternatives' account of knowledge.

Could Dan in these circumstances reasonably understand Ollie to be saying that he has done something to show that the ‘goldfinch’ wasn’t in fact a robotic imitation, created by the military for clandestine purposes? Such a scenario is at least conceivable; nothing in our understanding of the world makes it strictly impossible that someone could make a robotic goldfinch that was indistinguishable by sight from an actual goldfinch. If Ollie simply went into the garden and saw a goldfinch sitting on a branch, nothing in his experience would enable him to distinguish it from one in which he was instead looking at a robotic imitation.<sup>28</sup> In this sense he wouldn’t be in a position to rule out his having seen a robotic imitation instead of a goldfinch. On the other hand we can suppose that he would be in a position to rule out its being a robotic imitation if he had, for example, tapped the bird to make sure its exterior wasn’t concealing a metallic casing.<sup>29</sup> The question is, could we reasonably understand Ollie to be saying that he has done any such thing in saying that he knows there is a goldfinch in the garden? In many ordinary situations, no.

Austin (1961, p. 52) reminds us that ordinarily when a person claims to know that there is a goldfinch in the garden we will understand him to be saying that he has done

enough to show that (within reason, and for present intents and purposes) it ‘can’t’ be anything else, there is no room for an alternative, competing, description of it.

Our understanding of what ‘enough’ amounts to in any particular situation is rough to be sure. Nonetheless Austin thinks that this much can be said. In many situations we won’t understand a person who claims to know that there is a goldfinch in the garden to be saying that he has done “enough to show it isn’t a *stuffed* goldfinch.” The same presumably goes for a great many other conceivable alternatives to its being a goldfinch. In many situations we couldn’t reasonably understand a person who claims to know that there is a goldfinch in the garden to be saying that he has done anything to show that it isn’t, for example, a hallucination, or a hologram, or a robotic imitation, though each of these alternatives seems to be at least conceivable. The reason we couldn’t reasonably understand the speaker to be saying that the person is in a position to rule out these alternatives is that we wouldn’t generally be prepared to take such alternatives seriously. Where we have no reason to seriously suppose that the goldfinch might be a robotic imitation, it

---

<sup>28</sup> It is worth stressing that the fact that Ollie is not in a position to distinguish his experience of seeing a goldfinch from his experience of seeing a robotic imitation does not entail that Ollie is having the *same* experience in each case. Nor is there any good reason to suppose that Ollie’s experience is the same in each case, as Austin (1962) teaches us. More recent ‘disjunctive’ accounts of experience have echoed Austin’s thinking on this front. For one contemporary formulation of a disjunctive account of experience, see Paul Snowdon (2005).

<sup>29</sup> We can imagine another alternative on which the robotic bird has been so constructed as to be indistinguishable from an ordinary bird even when tapped. But this would be a different alternative to the one being considered.

will generally be uncomprehending to understand a person who claims to know that it is a goldfinch to be saying that he is in a position to rule out that particular alternative.

As Austin reminds us, however, we can imagine circumstances in which it would be reasonable to understand Ollie to be saying he was in a position to rule out one or another of these alternatives in claiming to know. These will be circumstances in which we and Ollie have some reason to seriously suppose that the 'goldfinch' might, for example, have been a robotic imitation. Suppose it has been all over the news that the military have admitted to releasing a large number of robotic goldfinches into the region. In that case if Ollie were to say he knows that something is a goldfinch, we could reasonably understand him to be saying that he has done at least something to rule out the possibility that what he was presented with was a robotic imitation. If we were to ask Ollie how he knows that there is a goldfinch in the garden in these circumstances, we could reasonably expect part of Ollie's response to speak to what he has done to rule out the possibility that it is a robotic imitation. For example, 'Because I've seen it' would no longer cut it as a response to that question, unless robotic imitations were distinguishable by sight from actual goldfinches. What might have satisfied us in the previous case as Ollie's being in a position to know that there is a goldfinch in the garden will no longer satisfy us in the situation where we have reason to take seriously the possibility of robotic imitations in the region. For Ollie to say that he knows he has seen a goldfinch even though he has done nothing to rule out its being a robotic imitation would, in these circumstances, be for Ollie to misrepresent his position to us.

So we have two broad kinds of circumstances in which Ollie might claim to know that there is a goldfinch in the garden. In one situation we might reasonably understand Ollie to be saying that he has done something to show that the goldfinch isn't a robotic imitation, whereas in the other situation it is not reasonable to understand Ollie to be saying this. It is important that the difference between these two cases is not that in one case (but not the other) it is conceivable that Ollie might be presented with a robotic imitation instead of a goldfinch. In both cases we can at least make up a coherent narrative about how something like this *could* happen. What differs between the two cases is what alternatives we could reasonably expect Ollie to be taking seriously into account.

When Austin suggests that a 'concrete' reason is needed to seriously suppose we might be mistaken in some particular way, he can be understood as suggesting a partitioning of the domain of conceivable alternatives in any particular situation. On the one side of the divide are alternatives we couldn't reasonably expect a person in his circumstances to be taking seriously into account; on the other side are alternatives we could reasonably expect him to be taking seriously. Insofar

as an alternative to  $p$  sits on the former side of the divide, we couldn't reasonably understand the person to be saying that he has done anything to rule that alternative out in saying that he knows that  $p$ . But if the person is aware of some reason for taking seriously a certain alternative, we *could* reasonably expect him to be in a position to rule this alternative out if he has said he knows that  $p$ .

The distinction Austin is drawing attention to here resembles in many ways Charles Peirce's distinction between real doubts and paper doubts.<sup>30</sup> Peirce (1868) coins the distinction in the course of criticising Descartes's suggestion that we should begin philosophical reflection by doubting everything we take to be so. He argues that what Descartes in fact does in the course of his meditations is *pretend* to doubt that he is, for example, seated by the fire. This in turn leads to Peirce's (1868, p. 141) famous edict that we should not "pretend to doubt in philosophy what we do not doubt in our hearts." Whether or not we think Peirce is right to criticise Descartes's methodological use of 'pretend' doubts, he is surely right to suggest that there is a distinction to be made here between dreaming up alternatives to what we take to be so and having reason to seriously suppose that such alternatives might in fact obtain. In general, he says, we require a 'positive reason' to take such alternatives seriously, and this presumably means more than simply noting that such alternatives are conceivable.<sup>31</sup>

To say that there is a distinction between alternatives that are relevant and those that are irrelevant in any particular situation is not to explain what in general determines whether an alternative should be considered relevant or not. Nor do I intend to offer any general explanation in that direction. But to eschew giving any explanation of why we consider an alternative relevant or irrelevant on an occasion should not, I think, lead us to suppose that the distinction is of no use. What matters is that we be able to recognise in any particular situation whether an alternative could reasonably be supposed to be relevant or not. This is something that we seem in general capable of doing. Reflecting on everyday situations in which a person might claim to know something, we are able to recognise which alternatives might reasonably be taken to be relevant in those circumstances and which not. If we didn't in general share a sense of which alternatives are relevant and which are not in particular circumstances, our practices of claiming and ascribing knowledge would be liable to break down; we would be liable to misunderstand what position a person was claiming to be in in claiming to know in particular cases. That our practices proceed

---

<sup>30</sup> A similar distinction is also at play in G.E. Moore's work on knowledge. See in particular 'A Defence of Common Sense' and 'Proof of an External World', both to be found in Moore (1993). For an illuminating discussion of Moore's views on this topic see Travis (1989, pp. 147-152).

<sup>31</sup> Hilary Putnam (2012, p. 524) takes Peirce's insight to be "that doubt requires justification as well as belief".

more or less smoothly where they do<sup>32</sup> suggests that we share a sense for what can and cannot reasonably be considered a relevant alternative in a situation.

What I am suggesting is that our understanding of what a speaker says in claiming to know something in particular circumstances can in part be articulated in terms of what alternatives we might reasonably expect him to be in a position to rule out, given what he has said. This is *not* the same as suggesting that in claiming to know something, a speaker is always to be understood as claiming to be in a position to rule out some more or less determinate range of alternatives. Any attempt to construe things in this way is liable to lead to a rather procrustean and strained account of what is going on when we claim to know things.<sup>33</sup> As Austin notes, saying that a person knows something amounts to saying that he has done enough to establish the truth of the matter. In many cases we would consider it enough if the person has consulted another person on the matter. As long as there is no reason to seriously suppose that the information received might have been mistaken, we generally think a person entitled to claim to know on the basis of the testimony in such circumstances. It seems a rather strained way of putting things to talk about the recipient of the testimony having ‘ruled out’ certain alternatives in this situation.

For that reason, I do not mean to endorse an account according to which saying that a person knows that  $p$  is always to be understood as saying that the person is in a position to rule out ‘ $x, y, z, \dots$ ’ alternatives to  $p$ . Nonetheless I *am* saying that the notion of a relevant alternative can in many cases be used to helpfully articulate our understanding of the kind of position a person is said to be in in being ascribed knowledge on an occasion. The role I envisage for the notion of a relevant alternative is largely negative. It can be illuminating to note that a person who is said to know on some occasion could not reasonably be understood as being in a position to rule out a particular alternative. It can be illuminating to note that when Ollie claims to know that there is a goldfinch in the garden, he cannot reasonably be understood as saying that he is in a position to rule out the possibility that it is a robotic imitation on that occasion. And it can be helpful to contrast such occasions with occasions on which we *could* reasonably take Ollie to be in a position to rule out that alternative, having claimed to know that there is a goldfinch in the garden.

---

<sup>32</sup> The reason for the qualification here is that we needn’t suppose that our practices of claiming and attributing knowledge *do* in fact always proceed smoothly. Nothing *guarantees* that we will in every situation share a sense of what is and isn’t reasonable, or that our practices won’t break down in some circumstances. Thanks to David Macarthur for pointing out this out to me. As Putnam (2012, p. 521) notes, in interpreting another person’s knowledge-claim we simply “have to use good judgment and assume that our conversational partners are attuned to us”. For an illuminating discussion of this topic see Stanley Cavell’s (1999) *The Claim of Reason*, particular Part 1.

<sup>33</sup> Thanks to David Macarthur for pointing this out to me.



Now consider the testimony case. It can be illuminating to note that in claiming to know something via testimony in particular situations, a speaker could not reasonably be understood to be saying that he is in a position to rule out the possibility that he was being lied to, or that his informant was mistaken. To say as much is to say that we wouldn't consider such alternatives to be relevant alternatives in those circumstances. This negative role for the notion of a relevant alternative still serves to illuminate our understanding of what kind of position a person who claims to know something in particular circumstances can reasonably be taken to be in.

In this section I have outlined how the notion of a 'relevant alternative' can be used to articulate our understanding of what a speaker says in ascribing knowledge in a particular situation. Austin observes that when we say that a person knows that  $p$ , we are saying that he has done enough for present intents and purposes to establish that  $p$ . The notion of a relevant alternative can be used to articulate our understanding of what 'enough' amounts to in any particular situation. As Austin notes, in many situations we won't understand a person who claims to know that there is a goldfinch in the garden to be saying that he is in a position to rule out his having seen a stuffed goldfinch (or a robotic imitation, or a hallucination, etc.). Even if we grant that such alternatives are conceivable, this in itself is not enough to make the alternative relevant to the question of whether or not the person knows. For an alternative to be relevant we need to have some reason to take that alternative seriously. In many situations we won't think it reasonable to seriously suppose that what a person takes to be a goldfinch might in fact be a robotic imitation. That said, we can imagine other circumstances in which we could reasonably expect a person to be taking seriously the possibility that the 'goldfinch' might in fact be a robotic imitation. In those circumstances we could reasonably understand a person who claims to know that there is a goldfinch in the garden as saying that he is in a position to rule out that particular possibility. In the next section I will consider an aspect of everyday epistemic and discursive practice that is directly relevant to this point, one that has come in for a great amount of attention among epistemologists in recent years.

## **2.2 Contrasting cases**<sup>34</sup>

What are we to be understood as saying when we say that someone knows something? In the previous section I outlined how the notion of a 'relevant alternative' can be used to articulate our

---

<sup>34</sup> It has become common in the literature to use contrasting cases of the kind I present in this section in order to motivate a contextualist view of knowledge. My own construction of these cases is directly influenced by Charles Travis (1989). Other contextualists like Keith DeRose (1992) and Stewart Cohen (1999) tend to present things so

understanding of what a speaker can reasonably be understood to be saying in saying that a person knows something. To say that a person knows something is to say that he has done enough in present circumstances to establish the truth of the matter. One way of articulating our understanding of what ‘enough’ amounts to in a situation is in terms of which alternatives we could reasonably suppose the person to be in a position to rule out, having been said to know. And we can contrast occasions on which we could reasonably expect a person to be in a position to rule out a certain alternative with other occasions on which we couldn’t reasonably expect this. I have referred to alternatives we could reasonably expect a person to be in a position to rule out as *relevant alternatives*. In this section I present two pairs of contrasting cases further illustrating this point. In the next section I detail how these cases can be seen to motivate a contextualist account of knowledge.

Suppose Dan is speaking with Elle, a bird enthusiast, who tells Dan that her favourite type of bird is the goldfinch. ‘You missed out,’ says Dan, ‘there was a goldfinch in our garden just yesterday.’ When asked how he knows it was a goldfinch, Dan replies that Ollie, who was also there to see the bird, told him so, ‘and Ollie would know – he’s completed a course in ornithology.’ This is a situation where we might ordinarily suppose that Dan has said something true of Ollie. Could Elle in this situation reasonably suppose that Ollie has done something to show that the bird wasn’t in fact a robotic imitation? It seems not. In this situation we wouldn’t take the suggestion that Ollie was in no position to show that the bird wasn’t a robotic imitation to count against what Dan says in saying that Ollie knows. In fact we can suppose that in this case Ollie wasn’t in any position to distinguish what he saw from a robotic imitation – we wouldn’t take Dan to have misrepresented Ollie’s position to Elle. In saying that Ollie knows, Dan is not to be understood as saying that Ollie has done anything to show that the bird wasn’t a robotic imitation. This reflects the fact that we couldn’t reasonably expect someone in Dan and Elle’s circumstances to seriously suppose that it might in fact have been a robotic imitation.

Now suppose that instead of speaking with Elle, Dan finds himself speaking with Blaise. Blaise is investigating rumours that the military has recently released a flock of robotic goldfinches into the neighbourhood, indistinguishable by sight from actual goldfinches. Suppose Dan says, ‘In any case I know that there was a goldfinch in the garden yesterday, Ollie told me and he would know.’ In this situation what Dan says about Ollie rings false. On a reasonable understanding of what he says to Blaise, Ollie could be expected to have done something specifically to show that what he

---

that the practical stakes for the conversational participants differ between the cases. While I have no immediate issue with focusing on such cases, I prefer Travis’s approach because it allows us to see that the relevant variation in our understanding of ascriptions of knowledge needn’t result solely from variations in practical stakes.

saw wasn't a robotic imitation. Supposing, as we did in the last case, that Ollie has done no such thing, it seems that Dan misrepresents Ollie's position to Blaise.

In both of these cases Dan is referring to Ollie considered at the same point in time. There is thus no change in Ollie's position between the two cases. Nonetheless it looks as though in one case Dan misrepresents Ollie's position in saying that he 'knows' (says something false of Ollie), whereas in the other case he does not misrepresent Ollie's position in saying that he 'knows'. What the examples thus appear to illustrate is that an ascription of knowledge can be used to accurately convey a person's position in one set of circumstances, and be used to *misrepresent* that same person's position in a different set of circumstances. 'Circumstances' here means circumstances for *ascribing* knowledge. The circumstances in which Dan speaks to Elle are different from the circumstances in which Dan speaks to Blaise. But Ollie's position remains *the same* across the two circumstances. In the next section I will suggest that this speaks in favour of a contextualism about knowledge. For now I just want to present the linguistic data.

Another case. Alistair is pretending he thinks Virginia Woolf wrote *Heart of Darkness*. Phoebe is concerned. Jamie assures her that Alistair is joking: 'He knows it was Joseph Conrad, we've been studying the book in class.' What kind of position can Phoebe reasonably take Alistair to be in in light of what Jamie says here? One thing we can say is that it wouldn't be reasonable for Phoebe to suppose that Alistair had done anything specifically to show that the book was not in fact secretly written by a friend of Conrad's. Jamie's remark is not to be understood in *this* way. We can suppose that *if* the book had been secretly written by someone other than Conrad, Alistair wouldn't be in a position to realise this. In that case the book's having secretly been written by someone other than Conrad is an alternative that Alistair is in no position to rule out, albeit one we would expect someone in Jamie and Phoebe's circumstances to treat as *irrelevant* to the question of whether or not Alistair knows (not something in need of ruling out).

Now suppose Molly is researching *Heart of Darkness* and is interested in following up a rumour that the book might have been secretly written in the main by one of Conrad's close friends. Jamie tells Molly that there is no need to look any further into the matter: 'Alistair already knows that it was Conrad who wrote the book'. Given Alistair's position, what Jamie says to Molly in these circumstances rings false. Given what Jamie says here, Molly could reasonably understand her to be saying that Alistair is in a position to show that the book was not secretly written by someone other than Conrad. But for Alistair to be in such a position he would need to have done more than we supposed he had done in the previous scenario. We can imagine what sorts of tests it might take to rule out the alternative in question. For example, Molly might be planning on

running a comparison between *Heart of Darkness* and other works known to have been written by Conrad.<sup>35</sup> Alistair has done no such thing. So in saying that Alistair already knows that the novel was written by Conrad, it seems that Jamie misrepresents Alistair's position to Molly.

Alistair's position remains the same in the two examples. However, in one case Jamie appears to accurately convey Alistair's position in saying that he knows, whereas in the other case Jamie appears to misrepresent Alistair's position in saying that he knows. The salient difference between the two cases is that in one, but not the other, there is a relevant alternative Alistair is in no position to rule out. There is no change between the cases in which alternatives Alistair is in a position to rule out, and which not.

It might seem strange to talk about Alistair's position being misrepresented in being said to know in one case, but not in the other (the same goes for the cases involving Ollie). The thought (not one I want to endorse) might be articulated as follows. Alistair either knows or he doesn't; there is no room for him to be characterised both as knowing that Conrad wrote the book and not knowing this. If he knows that Conrad wrote the book, then it is an accurate representation of his position, *true*, to say that he knows. Likewise, if he doesn't know that Conrad wrote the book, then it is an inaccurate representation of his position, *false*, to say that he knows. It might be that Jamie is for some reason prevented from seeing which it is in one or the other situation. Nonetheless it is the task of the philosopher to work out exactly what position Alistair needs to be in if he is to count as knowing or otherwise, to work out what position Alistair is in if it is *true* to say he knows. Taking this view of the matter, it is liable to look as though the specifics of the different circumstances in which Jamie *says* that Alistair knows are *irrelevant* to the question of whether he in fact knows the thing in question. All we need to do, we might think, is reflect on *Alistair* and his position. This line of thought has its attractions. It is plausibly something like this idea that underlies Stroud's criticism of Austin to the effect that Austin reminds us only of facts regarding when it is appropriate to say that someone knows something, and that he fails to tell us anything illuminating about what it 'actually' takes to know something.<sup>36</sup> But it is a line of thought that Austin encourages us to challenge. Before considering what is problematic about

---

<sup>35</sup> We can suppose that while Molly is prepared to seriously entertain the possibility that one of Conrad's books was secretly written by someone else, she is not prepared to seriously entertain the possibility that *all* the works attributed to Conrad were written by someone else. So if the writing in *Heart of Darkness* matches the writing in Conrad's other works closely enough, this will be grounds for concluding that the novel was indeed written by Conrad.

<sup>36</sup> At the end of 2.5 I suggest that there is very likely a connection between Stroud's 'invariantist' approach to thinking about knowledge, and his finding something like the sceptical conception of knowledge from 1.1 compelling.

this line of thought I first want to present an alternative way of viewing the two cases, one which sees them as motivating a *contextualist* understanding of knowledge and knowledge ascriptions.

### **2.3 A contextualist interpretation of the cases<sup>37</sup>**

In this section I present a contextualist interpretation of the contrasting pairs of cases considered in the previous section. Contextualism is the view that knowledge-ascribing sentences admit of a variety of understandings, even once we fix who is said to know and what he is said to know. What varies from one understanding to another is what kind of position the person needs to be in if he is to count as knowing the thing in question. This means that while it might be true to say that some person S knows something *p* on one understanding of what it is to know that *p*, there can also be other understandings of what it is to know that *p* on which S wouldn't count as knowing that *p*. So one speaker might speak truth in saying that S 'knows' that *p*, while another speaker would be saying something false if he were to say that this same S 'knows' that *p*. One way of articulating the variety of understandings knowledge-ascribing sentences can support is in terms of the previously discussed notion of a relevant alternative. An alternative that is relevant on one understanding of what it is to know that *p* may not be relevant on another understanding of what it is to know that *p*. But it seems plausible that for any alternative to *p* we can imagine, there will be understandings of what it is to know that *p* for which that alternative counts as a relevant alternative. Given the (perhaps indefinitely) many understandings that a knowledge-ascribing sentence can support, an appreciation of the speaker's circumstances is required in order to determine which understanding is being invoked.

Consider again the pair of cases discussed in the previous section concerning Ollie and the goldfinch. In the case where Dan tells Elle that Ollie knew there was a goldfinch in the garden, we wouldn't think it reasonable to understand Dan as saying that Ollie was in a position to distinguish his seeing a goldfinch from his seeing a robotic imitation. We would not expect either Dan or Elle to treat Ollie's having seen a robotic imitation as a relevant alternative on this occasion. As such, unless we had some other reason to seriously suppose that Ollie might have been mistaken in judging that what he saw was a goldfinch, we would ordinarily take Dan to have said something true in saying that Ollie knew (to have accurately presented Ollie's position in the

---

<sup>37</sup> There has been a proliferation of contextualist accounts of knowledge in recent years, and different theorists have presented differing forms of contextualism. My own thinking on the subject has been particularly influenced by Charles Travis (1989, 1991, 2005). Travis sees his contextualism – what he calls an 'occasion-sensitive' view of language – as arising directly out of the work of J.L. Austin and the later Wittgenstein. Other prominent contextualists include Keith DeRose (1992, 2009), Stewart Cohen (1999), and David Lewis (1996).

situation). However, in the circumstances in which Dan tells Blaise that Ollie knew that there was a goldfinch in the garden, Dan could reasonably be understood as saying that Ollie was in a position to distinguish between his seeing a goldfinch and his seeing a robotic imitation. As such, it appears as though Dan misrepresents Ollie's position in these circumstances – says something *false* in saying that Ollie knew.

A contextualist takes this data at face value, and holds that Dan can be saying something true in saying to Elle that Ollie knew, and saying something false in saying to Blaise that Ollie knew this same thing. The reason this doesn't amount to a contradiction is that what Dan says in each case is that Ollie knows on a *certain understanding* of what it is for Ollie to know the thing in question. On one understanding of what it would take for Ollie to know that he has seen a goldfinch, Ollie needn't be in a position to distinguish his having been presented with a goldfinch from his having been presented with a robotic imitation. But on another understanding of what it would take for Ollie to know this same thing, Ollie does need to be in a position to distinguish his having been presented with a goldfinch from his having been presented with a robotic imitation.

The same interpretation is available to a contextualist in relation to the second pair of cases as well. A contextualist can hold that when Jamie says to Phoebe that Alistair knows Conrad wrote *Heart of Darkness*, she says something that can be true even if Alistair is in no position to show that *Heart of Darkness* wasn't secretly written by someone else. However, in the case where Jamie says to Molly that Alistair knows Conrad wrote *Heart of Darkness*, a reasonable person in Molly's position would take Jamie to be saying that Alistair is in a position to rule out the possibility that the book was secretly written by someone else. Alistair is not in a position to do this, so Jamie would be saying something false of Alistair. In other words Alistair's position alone doesn't tell us whether a person would be saying something true or false in saying that he knows that Conrad wrote *Heart of Darkness*. In order to determine this we need some understanding of the circumstances in which the speaker says Alistair to know this.

The contextualist view is that an appreciation of the particular circumstances in which a sentence is employed is required if we are to understand what the speaker is saying to be so. This is one way of understanding what Austin is saying in the following passage:

It seems to be fairly generally realised nowadays that, if you just take a bunch of sentences [...] impeccably formulated in some language or other, there can be no question of sorting them out into those that are true and those that are false; for [...] the question of truth and falsehood does not turn only on what a sentence *is*, nor yet on what it *means*, but on, speaking very broadly, the

circumstances in which it is uttered. Sentences are not *as such* either true or false. (Austin, 1962, p. 111)

Austin can be understood here as saying that a knowledge of what the *sentence* 'Alistair knows who wrote *Heart of Darkness*' means does not yet determine what it would take for that sentence to be true or false. In one sense this is trivial. Obviously we need to determine who 'Alistair' is and what exactly he is being said to know. A contextualist goes further than this, holding that even once we fix who 'Alistair' is and what he is said to know, there are a variety of distinct understandings of what Alistair's 'knowing' could come to.<sup>38</sup> Depending on which of these understandings is invoked on an occasion for ascribing knowledge to Alistair, it could be either true or false to say in those circumstances that Alistair knows that Conrad wrote *Heart of Darkness*. In order to determine what kind of position Alistair is being said to be in in being said to 'know' something, we need some appreciation of the particular circumstances in which the speaker uses the sentence. Different speakers in different circumstances can be understood to be saying different things of Alistair in saying that he knows, depending on the particular circumstances in which they speak. This can be so even if we fix Alistair's position and what he is said to know.

In summary, a contextualist holds that there are a variety of distinct understandings of what a person's knowing might come to in a particular situation. If we take some person S and some fact *p*, there remains a variety of distinct things to be said of S in saying that he 'knows' that *p*. In each case S is said to know that *p* on a certain understanding of what it is to know the thing in question. On some understandings of what it would be for S to know that *p*, S might count as knowing that *p*. That still leaves open the possibility that there are other understandings of what it would be for S to know that *p* on which S would *not* count as knowing that *p*. One way of articulating the difference between these understandings of what it would be for S to know that *p* is to consider what alternatives S would need to be in a position to rule out in order to count as knowing on a particular understanding. In order to determine which of these (perhaps indefinitely many) understandings is being invoked on an occasion, some appreciation of the particular circumstances in which S is said to know is required.

---

<sup>38</sup> Charles Travis defends an interpretation of Austin along these lines, an interpretation that has also been echoed by Hilary Putnam (2001, 2012) and Alice Crary (2002, 2007) among others.

## **2.4 'Knowing' and 'being green'**

It might help to make clearer the contextualist position if we briefly compare a contextualist account of knowledge with a contextualist account of another attribute: the attribute of 'being green'.<sup>39</sup> There is a piece of confectionary known as a 'clinker'. A clinker is a hardened piece of coloured fondant, coated in a layer of chocolate. The fondant-inner comes in different colours: generally green, pink, or yellow. Consider the sentence 'This clinker is green'. Taken at face-value, such a sentence is to be used for ascribing the property of *being green* to a clinker. Suppose we are presented with a green-centred clinker. We can imagine circumstances in which a person would intuitively say something true in saying that the clinker is green. Consider the following example:

Ashleigh and Brianna are working in a factory where clinkers are produced. Clinkers are stored in different containers depending on the colour of the fondant-inner. It is Brianna and Ashleigh's job to place an even distribution of each type of clinker into bags to be sent to the shops. Brianna says to Ashleigh, 'I need one more green clinker.' Ashleigh fetches her a clinker from one of the containers and says, 'This clinker is green'.

In this situation, provided the clinker Ashleigh presents is one with a green fondant-inner, we would take Ashleigh to have said something true. Now imagine a different scenario:

A child has been tasked with sorting a group of objects by the colour of their exterior. She has a pile for red objects, green objects, brown objects, and so on. Suppose the child places the clinker in the pile for green objects. When asked why she placed the clinker in this pile the child responds, 'This clinker is green'.

In this situation what the child says is intuitively false. The clinker is not to be counted as green in this situation, but is instead to be counted as *brown*. For the child to describe the clinker as green in this situation would suggest that she was not yet clear on the meaning of the word 'green'. Coming to understand what the word means would in this case plausibly involve coming to realise that what she said was mistaken (false) when she said that the clinker was green.

The moral that emerges from these two stories is that, even if we fix the object being spoken about here (i.e. a particular clinker), the sentence 'This clinker is green' can be used to say something true or false of this same object, depending on the circumstances. The reason we can acknowledge this without contradiction is that we see that what is said in each case is that the object is green, on a *certain understanding* of something's being green. Suppose a philosopher (or

---

<sup>39</sup> The strategy I use in this section to illustrate the context-sensitivity of 'being green' is directly modelled on the strategy employed by Charles Travis, particularly in his paper 'Pragmatics' (2008).



anyone else) were to set himself the task of working out whether the clinker is ‘really’ green or not. It is difficult to see what this philosopher could imagine himself to be doing. After all, we can see what would qualify the clinker for being green on the one understanding of something’s being green, and we can also see what would qualify the clinker for being green on the other understanding of something’s being green.<sup>40</sup> What further question is there to be asked regarding whether the clinker is ‘really’ green or not? A contextualist would hold that such a question has no motivated answer.<sup>41</sup>

Similarly a philosopher considering the two situations in which Dan says that Ollie ‘knew’ that what he saw was a goldfinch might want to ask what it would take for Ollie to know, *punkt*. A contextualist about knowledge takes such a question to be as fruitless as the question about whether the clinker is ‘really’ green or otherwise. When Dan says to Elle that Ollie knew, he says so on one particular understanding of what Ollie’s knowing that what he sees is a goldfinch amounts to. When Dan says to Blaise that Ollie knew, he says so on a different understanding of what Ollie’s knowing that what he sees is a goldfinch amounts to. In the latter case, but not necessarily in the former, Ollie only counts as knowing if he is in a position to distinguish between his being presented with a goldfinch and his being presented with a robotic imitation. For a contextualist, this is similar to the fact that when the child says that the clinker is green, it counts against what she has said if the clinker’s exterior is brown, whereas this doesn’t count against what Ashleigh says to Brianna. Just as a contextualist about ‘being green’ doesn’t think that there is any conflict between our intuitions in relation to each case involving the clinker, a contextualist about ‘knowing’ doesn’t think that there is any conflict between our intuitions regarding what Dan says about Ollie in each case.

There remains the fact that more philosophers are inclined to pursue a non-contextualist account of knowledge than are inclined to pursue a non-contextualist account of ‘being green’. In the next section I consider the prospects for an ‘invariantist’ account of knowledge in light of the examples presented in 2.2. I will argue that each of the interpretations of these examples available to an invariantist is less plausible than the contextualist interpretation. That a contextualist is better-placed to make sense of the examples in 2.2 speaks in favour of adopting a contextualist account of knowledge over an invariantist account.

---

<sup>40</sup> The child would intuitively have said something true if the clinker had been painted green.

<sup>41</sup> This is hardly the end of the matter. For a further discussion of the implications of taking such a contextualist account of properties like ‘being green’ seriously, see Travis (2001).

## **2.5 Issues with an ‘invariantist’ account of knowledge**

In this section I consider the prospects for what has come to be referred to in the literature as an ‘invariantist’ treatment of examples like the ones considered in section 2.2. Contextualists hold that there is no single understanding of what it would take for some person *S* to count as ‘knowing’ some fact *p*, and that the word ‘knows’ and its cognates<sup>42</sup> can thus be used to say different things<sup>43</sup> of a person depending on the circumstances in which they are used. In contrast, an invariantist<sup>44</sup> is someone who holds that once we have determined who is said to ‘know’ and what he is said to ‘know’, there is only one truth-evaluable content to be expressed in saying that the person ‘knows’ the thing in question. Put otherwise, if one speaker says something false in saying that *S* ‘knows’ that *p*, then any other speaker who says that *S* ‘knows’ that *p* must also be saying something false. I consider some invariantist interpretations of the examples presented in 2.2, and argue that each of these interpretations is less plausible than the contextualist interpretation. The ability of contextualists to make better sense of our epistemic and discursive practices in this area speaks in favour of adopting a contextualist stance over an invariantist stance.

In 2.2 I considered two contrasting pairs of cases. Consider again the examples involving Ollie and the goldfinch. In the case where Dan tells Elle that Ollie knew it was a goldfinch he saw in the garden, he couldn’t reasonably be understood to be saying that Ollie was in a position to distinguish his being presented with a goldfinch from his being presented with a robotic imitation. To understand Dan in this way would be uncomprehending. We wouldn’t ordinarily think that Ollie needed to be in such a position in order for Dan to have said something true of him. Contrast this with the case in which Dan tells Blaise that Ollie knew that what he saw was a goldfinch. In the circumstances in which Dan says this to Blaise, Blaise could reasonably take Dan to be saying that Ollie was in a position to distinguish between his being presented with a goldfinch and his being presented with a robotic imitation. Construed in this way, Dan would appear to be misrepresenting Ollie’s position. That is, Dan would appear to be saying something false of Ollie. On a contextualist interpretation of these cases, Dan can be saying something true when he says to Elle that Ollie knew, even if he says something false in saying this to Blaise. The reason that our reactions to the two cases needn’t be thought to contradict each other is that in

---

<sup>42</sup> E.g. ‘sees (that)’, ‘remembers’, etc.

<sup>43</sup> In a sense. What a speaker says in each case is that the person *knows*. What varies is what it would take for the person in question to count as knowing.

<sup>44</sup> As with contextualism there has been a proliferation of invariantist accounts of knowledge in recent years. For a contemporary representative of the kind of traditional invariantism implicit in much of the epistemological literature on knowledge, see Timothy Williamson (2005). For presentations of what has come to be known as ‘interest-relative’ or ‘subject-sensitive’ invariantism, see John Hawthorne (2004) and Jason Stanley (2005). While there are important differences between these types of invariantism, the general worries I will raise apply to invariantisms across the board.

each case Dan says that Ollie knew on a different understanding of what it is to know the thing in question.

An ‘invariantist’ for my purposes is anyone who is committed to the view that there is only one thing to be said of Ollie in saying that he knew that what he saw was a goldfinch, irrespective of who is describing Ollie and in what circumstances. In other words, an invariantist is committed to the view that if Dan says something true when he says to Elle that Ollie knew, then he also says something true when he says this to Blaise. Conversely, an invariantist will hold that if Dan says something false when he says to Blaise that Ollie knew, he also says something false in saying this to Elle. There is no room on an invariantist interpretation for the kind of variation in truth-evaluable content that is a fixture of the contextualist interpretation. Once we have fixed who we are talking about and what it is they are to be said to know, there is only one thing to be said of that person in saying that he ‘knows’ the thing in question. The person’s position will either be such that he knows the thing in question, or doesn’t know, regardless of who is considering the matter.

When Dan tells Elle that Ollie knew, he appears to present Ollie as being in a different position than the position he presents him as being in when he tells Blaise that Ollie knew. In the latter case, but not in the former case, Dan appears to present Ollie as having been in a position to distinguish between his being presented with a goldfinch and his being presented with a robotic imitation. A contextualist takes this at face-value as showing that Dan is to be understood as saying that Ollie was in a different position in each case. That is, a contextualist holds that when Dan says to Blaise that Ollie knew, he is to be understood as saying that Ollie was in a position to distinguish between his being presented with a goldfinch and his being presented with a robotic imitation. Given that Ollie was in no such position in the case described, Dan is thus to be understood as having said something false to Blaise. Such is a natural response to the case as described. However, interpreting the case in this way can lead to problems if we adopt an invariantist view of knowledge.

If we are invariantists and hold that Dan says something false when he says to Blaise that Ollie knew that what he saw was a goldfinch, we will also need to hold that Dan says something false when he says to Elle that Ollie knew. This clashes with the intuition that Ollie needn’t have been in a position to distinguish the goldfinch from a robotic imitation in order to be in the position Dan credits him as having been in to Elle. One way of explaining away this intuition is to hold that Dan and Elle are in this particular case mistaken about what does and doesn’t count as a relevant alternative for the purposes of considering whether Ollie was in a position to know. Dan

and Elle are not aware of any special reason to suppose that there might be robotic imitations buzzing about, and we wouldn't expect a person in their position to consider robotic imitations a relevant alternative in need of ruling out. Nonetheless, we are sometimes mistaken about what is and isn't a relevant alternative for the purposes of judging a person's knowledge.<sup>45</sup> If Elle and Dan were to find out that there was a flock of robotic birds in the area at the time, they would likely come to revise their earlier judgement about whether Ollie was in a position to know, and this because they had come to realise that there was a relevant alternative they were previously unaware of. So Dan says something false of Ollie in both cases, and this because Ollie was in no position to rule out his having seen a robotic imitation instead of a goldfinch.

I don't want to dispute the point that Elle and Dan might come to realise that an alternative they hadn't thought relevant on an occasion was in fact relevant to the question of whether Ollie was in a position to know. But there is reason to suppose that an invariantist still faces difficulties in making sense of our reactions to the two cases. Suppose that the rumours Blaise is looking into about robotic imitations in the region prove in the end to be unfounded. Given that the rumours were unfounded, it looks as though Dan wasn't mistaken in saying to Elle that Ollie knew that what he saw was a goldfinch. After all, they had no reason to seriously suppose that what Ollie saw was a robotic imitation, so no reason to think that this was something Ollie should be in a position to rule out if he was to count as knowing that what he saw was a goldfinch. Suppose an invariantist were to share this view of the matter. In that case he would need to hold that Dan says something true when he says to Blaise that Ollie was in a position to know. This looks implausible.

When Dan says to Blaise that Ollie was in a position to know, he presents Ollie as having been in a position to distinguish between his being presented with a goldfinch and his being presented with a robotic imitation. Taken at face-value, it looks to be part of what Dan says to Blaise that Ollie was in such a position. So it looks like Dan has said something *false* to Blaise, has misrepresented Ollie's position in saying that he was in a position to know. This would remain the case even if it later turned out that the rumours Blaise was chasing up proved to be unfounded. Dan would in that case still have presented Ollie as having been in a position he was not to Blaise. Given what Dan said, Blaise could reasonably suppose that Ollie was in a position to rule out his having seen a robotic imitation. Ollie was not in such a position. So it certainly seems as though Dan conveys false information to Blaise in saying that Ollie knew.

---

<sup>45</sup> Travis (1989, p. 161) provides a nice illustration of this point.

An invariantist might respond that, although Dan does indeed convey false information about Ollie to Blaise in this situation, he nonetheless says something true in saying that Ollie knew. To maintain this line the invariantist would need to hold that the position Dan literally says Ollie to have been in when he says that he knew is otherwise than the position that he *suggests* to Blaise that Ollie was in. This is a rather *ad hoc* way of accounting for the intuition that Dan says something that is plain and simply false to Blaise. We are familiar enough with cases in which a person says something true while suggesting something false. Suppose Ashleigh asks Brianna where she can get petrol at this time. Brianna, wanting to get rid of her in a hurry, tells her that there is a petrol station around the corner. Brianna is well aware that the petrol station is closed, and that Ashleigh will understand her to mean that the station is open. Nonetheless it is plausible to suppose that what Brianna says here is true, even if what she thus *implies*<sup>46</sup> is that the station is open. So we can plausibly imagine cases where a person says something true while implying something false. It does not, however, look plausible that the case in which Dan tells Blaise that Ollie knew might be understood in this way.

In the case where Brianna tells Ashleigh that there is a petrol station around the corner, we can readily make sense of the idea that Brianna says something true even while she misrepresents the situation to Ashleigh.<sup>47</sup> Part of what Brianna can be understood as meaning to convey to Ashleigh is that there is a petrol station (open or closed) around the corner, and both Brianna and Ashleigh and anyone else in these circumstances could be expected to have a reasonable understanding of how things would need to be in order for this to be true in this case. It is not clear that anything plays this same role in the case where Dan tells Blaise that Ollie was in a position to know. Dan says that Ollie knew and thus presents Ollie as having been in a position to rule out his having been presented with a robotic imitation. In what sense can Dan be reasonably understood as having presented *anything* that is true to Blaise? It is not as though there is in these circumstances a way of taking Dan's words on which Dan could be reasonably understood to be conveying something accurate about Ollie's position. Unless someone wants to insist that Dan (albeit unwittingly) at least accurately presents Ollie as *knowing* to Blaise. But what is at issue is our understanding of what knowing is to come to in the situation! And in Dan and Blaise's

---

<sup>46</sup> The notion of implicature came into philosophical prominence as a result of Paul Grice's work in the philosophy of language. See Grice (1989). Grice's notion of implicature has been thought by Grice and others to neutralise some of the main criticisms of traditional philosophy by Austin and Wittgenstein. Charles Travis (1991) provides a forceful criticism of such a view.

<sup>47</sup> I'm assuming for present purposes that Brianna says something true while implying something false, but it is important to note that this reading of the situation isn't uncontroversial. It is at least an open question whether Brianna hasn't simply said something to Ashleigh that is false, even if some philosophers would like to suggest that it is not at all an open question. As this is tangential to the present point, I will assume the more orthodox interpretation of the situation.

circumstances it seems that the most reasonable way of understanding what Ollie's *knowing* comes to involves him being in a position to rule out his having been presented with a robotic imitation. It's not as though Blaise could be expected to have any idea of what Ollie's knowing is meant to come to, independently of the position Dan presents him as being in in these circumstances.

The idea that Dan says something true to Blaise while conveying something false seems to be an overly stretched interpretation of what goes on in this situation. And it seems the only reason one might have for accepting such a strained interpretation of the situation is a prior commitment to an invariantist account of knowledge, exactly what this manoeuvre was supposed to make plausible.

This makes it look as though the most plausible route for an invariantist is to say that Dan says something false in both cases after all. But this strategy is liable to lead us in a thoroughly unpalatable direction. To opt for this strategy means holding that Ollie's having seen a robotic imitation was in fact a relevant alternative in Dan and Elle's situation, in spite of our not generally taking it to be so. Does that mean that any time a person knows that something is a goldfinch, he must be in a position to distinguish his being presented with a goldfinch from his being presented with a robotic imitation? Blaise is following up a *rumour*, a rumour that might prove to be unfounded in the end. Are we to suppose that the mere existence of such a rumour is enough to make it so that people in Elle and Dan's circumstances have reason to seriously suppose that there might be robotic imitations about (even if they are happily unaware of the rumour)? To take this line on the matter is to begin a quick walk down the path to something like the sceptical conception of knowledge considered at the beginning of the previous chapter. After all, we can conceive of all kinds of situations in which people might come to take seriously a certain alternative to something *p*, even temporarily, that we wouldn't in other circumstances take ourselves to have any reason to seriously entertain. And for these people in these circumstances, to say that someone knows that *p* will amount to saying that the person has done something to rule out that particular alternative. But if we take an invariantist line we will need to suppose that ascriptions of knowledge made in other circumstances likewise require that the person be in a position to rule out that alternative. In that case we are effectively opening the door to the idea that just about any conceivable alternative is a relevant alternative in every situation. This suggests something of a link between taking an invariantist view of knowledge and finding oneself driven in the direction of a sceptical conception of knowledge like Stroud's.<sup>48</sup>

---

<sup>48</sup> So Stroud's real stand-off with Austin might in fact be understood as being over whether we are to adopt a contextualist or an invariantist conception of knowledge. This seems to be Putnam's (2001) understanding of the

This is hardly the end of the matter and an invariantist might come upon other strategies for making sense of the linguistic data in the examples described. That said, it looks as though the contextualist has the advantage here in being able to readily make sense of our epistemic and discursive practices in this area. This speaks strongly in favour of adopting a contextualist view of knowledge over an invariantist position. It also looks as though the contextualist can make some sense of why we might in some circumstances be compelled in the direction of a sceptical conception of knowledge. We are liable to be compelled in such a direction precisely when we attempt to speak ‘absolutely’ about knowledge, when we imagine that there is only one understanding of what a person’s knowing might come to in a particular situation. And where we find ourselves compelled in such a direction, Austin’s approach furnishes us with an antidote. We are to remind ourselves once again of the various kinds of circumstances in which a speaker might say that someone knows, and to consider how the speaker could reasonably be understood in each case. We will then be in a position to see that the idea that there is just one thing a person’s knowing might come to is not a compulsory one.

## **2.6 Assurance-giving and contextualism about knowledge**

I have argued that the examples adduced in 2.2 speak in favour of a contextualist account of knowledge. On a contextualist account of knowledge there are variety of different understandings of what a person’s knowing might come to, even once we fix the person and the thing he is to be said to know. In order to determine what it would take for the person to count as knowing the thing in question, we require some understanding of the particular circumstances in which a speaker ascribes knowledge to him. When we survey the different circumstances in which a speaker might judge that someone knows something, we can see that alternatives which are irrelevant in one set of circumstances might nonetheless be relevant on other occasions for judging the matter. In this section I link this contextualist understanding of knowledge with another important insight that emerges from Austin’s work, namely that there is an intimate connection between our understanding of what it is to know something and our understanding of who is in a position to act as an informant on which matters. Taking seriously this connection provides us with further means for making sense of the context-sensitivity of ascriptions of knowledge.

---

disagreement between Stroud and Austin. Travis (1991, pp. 245-246) similarly suggests that an invariantist (or what he calls ‘Gricean’) interpretation of the linguistic data is liable to compel us in the direction of an implausible sceptical conception of knowledge.

Austin (in)famously suggests that there is an interesting comparison to be made between the act of saying ‘I know’ and the act of saying ‘I promise’:

[W]hen I say ‘I promise’, a new plunge is taken: I have not merely announced my intention, but, by using this formula (performing this ritual), I have bound myself to others, and staked my reputation, in a new way. Similarly, saying ‘I know’ is taking a new plunge. But it is *not* saying ‘I have performed a specially striking feat of cognition, superior, in the same scale as believing and being sure, even to being merely quite sure’: for there *is* nothing in that scale superior to being quite sure. Just as promising is not something superior, in the same scale as hoping and intending, even to merely fully intending: for there *is* nothing in that scale superior to fully intending. When I say ‘I know’, I *give others my word*: I *give others my authority for saying* that ‘S is P’. (1961, p. 67)

This passage has proved to be obscure, with many commentators unsure what exactly Austin means to be suggesting in asking us to compare the expressions ‘I know’ and ‘I promise’.<sup>49</sup> One plausible reading of this passage comes from Benjamin McMyler (2011), who suggests that Austin is inviting us to compare the acts of promising and telling. Just as ‘I promise I’ll be at the airport’ can be used to promise a person that I will be at the airport, it seems plausible that ‘I know that’s a goldfinch’ can be used to tell another person that something is a goldfinch. Telling is a normatively rich act; when a speaker tells another person something, he assumes a distinctive responsibility for the person’s belief on the matter.<sup>50</sup> Similarly, when a speaker promises a person to do something, he assumes a distinctive responsibility in relation to the person for doing that thing. Furthermore, there is a difference between telling another person something and merely voicing one’s own opinion on a matter. When we want to convey to another person that we are not prepared to assume responsibility for his belief on a matter, we will tend to qualify what we say by saying ‘I think’ or ‘I believe’. So the expressions ‘I know’ and ‘I believe’ are fitted for use in very different kinds of discursive acts.<sup>51</sup> This certainly seems to be part of what Austin takes his comparative remarks to be illustrating. The thought that there is an interesting comparison to

---

<sup>49</sup> A number of commentators have thought to ascribe to Austin the patently absurd view that ‘I know’ is used to perform the ‘act’ of knowing. ‘I promise’ can be used to perform the act of promising, but ‘knowing’ is not an act that can be performed in speaking. For an example of a reading of Austin tending in this direction, see Chisholm (1964, p. 10).

<sup>50</sup> See Richard Moran (2006, 2013) for an excellent discussion of the act of telling another person something.

<sup>51</sup> Note that Austin needn’t be understood as suggesting that the *only* way of performing the act of telling is by using the expression ‘I know’, any more than he need be understood as suggesting that the only way of making a promise is by saying ‘I promise’. It is plausibly only in certain circumstances where we need to make clear what kind of commitment we mean to be undertaking with respect to our audience that we will have recourse to explicitly *say* ‘I know’.



be made between the acts of telling and promising is one that continues to inspire a great amount of philosophical reflection today.<sup>52</sup>

While McMyler is surely right to suggest that Austin is interested in comparing the acts of promising and telling, there remains the question of how Austin's remarks might be thought to bear on our understanding of what knowing amounts to. What exactly is the link between the expression 'I know' and the act of telling? One plausible way of understanding Austin's point here is as follows. Austin's point is that in saying 'I know' in the context of telling another person something, a speaker presents himself as being in a position to settle the matter in question on behalf of his audience. Put otherwise, in saying 'I know' a speaker presents himself to his audience as having already done enough (for present intents and purposes) to establish the truth of the matter, and so is in a position to absolve them from having to go through the same steps. If this is right, this suggests a more general moral to be taken from Austin's remarks on this topic. The thought is this: taking a person to be in a position to give his word on a matter means taking him to know the thing in question.

The thought that there is an intimate connection between knowing and being in a position to provide sound assurances on a matter has been developed in recent years by Edward Craig and Michael Welbourne, among others. Craig (1990, p. 11) begins his reflections on the concept of knowledge with the hypothesis that the core function of the concept is to flag persons as approved sources of information. And Welbourne (1986) defends the view that a person's being in a position to authoritatively communicate the truth of a matter to others is what knowing amounts to in our epistemic and discursive practices. I want to suggest that the idea that there is an intimate connection between knowing and being in a position to provide sound assurances on a matter can be used to further investigate the context-sensitivity of ascriptions of knowledge.

Consider the following two occasions for judging whether or not Alistair knows that Conrad wrote *Heart of Darkness*. On the first occasion Jamie is compiling a reading list for a book group. She has included *Heart of Darkness* on the list but has no idea who the author is. In this context she asks Alistair if he knows who wrote the book. On the hypothesis in question, in asking Alistair whether he *knows* who wrote the book Jamie is interested in whether Alistair has done enough to settle the matter for her purposes, and is thus in a position to issue her with a sound assurance on the matter. What kind of position would Alistair need to be in to settle the matter for Jamie's purposes? Ordinarily we would think that Alistair had done enough to settle the matter for Jamie's

---

<sup>52</sup> In addition to Moran and McMyler, see Hinchman (2005) and Faulkner (2007).

purposes if he had, for example, recently read the book and taken note of the author, or if he had performed a quick Google search. So reflecting on what kind of position Alistair would be presenting himself as being in in telling Jamie who wrote the book in these circumstances provides us with a lens through which to articulate our understanding of what it would take to know this in such circumstances.

Now consider a different occasion for judging whether or not Alistair knows who wrote *Heart of Darkness*. Suppose Jamie is talking with her friend Molly who is conducting research to see whether certain novels might in fact have been secretly written by someone other than the accredited author. *Heart of Darkness* is on the list of novels to be investigated. Were Jamie to suggest to Molly in these circumstances that Alistair knows that Conrad wrote the novel, she could reasonably be taken to be saying that Alistair has already done the requisite work for Molly's purposes establishing who wrote the novel. That is Jamie would be presenting Alistair to Molly as being in a position to shoulder the burden for having established the truth of the matter in this case. But for Molly's purposes, establishing the truth of the matter requires showing that the book wasn't secretly written by someone else, something that Alistair has not done. So Jamie misrepresents Alistair's position to Molly in saying that he knows the thing in question in these circumstances, presents him as being in a position to shoulder a responsibility he cannot.

The above pair of contrasting cases are not intended to furnish us with any fresh intuitions beyond what has already been presented in section 2.2. Nonetheless these examples do provide some further illumination regarding the kinds of considerations that might be shaping our understanding of who is to be counted as knowing in different circumstances. The suggestion here is that in many cases where we consider what someone says in saying that a person 'knows', our understanding of what is said is shaped by our understanding of what kind of position the person would need to be in in order to authoritatively settle the matter in those circumstances. So reflecting on what kind of position we would expect a person to be in if he is to authoritatively settle the truth of a matter on behalf of others in a situation is one way (not the only way) of further exploring our understanding of what it takes to know something. Given that what doubts we might require an assurance to be given against is liable to vary depending on the circumstances in which we find ourselves, this provides us with a way of making further sense of the context-sensitivity of ascriptions of knowledge.

## **2.7 Summary**

In this chapter I argued that Austinian considerations speak in favour of a contextualist account of knowledge. On a contextualist account of knowledge it is possible for one speaker to say something true in saying that a person knows something, even if it would be false for another speaker in different circumstances to say that this same person knows this same thing. Such a situation is possible because on a contextualist account of knowledge, the word 'knows' admits of (perhaps indefinitely many) different understandings. A person might count as knowing on one understanding of what it is to know the thing in question, while not so counting on a different understanding. To determine whether a speaker says something true in saying that a person knows something, it is not enough to fix who the person is and what he is said to know. What we need is some idea of how the speaker is to be understood in saying that the person *knows*, and in order to determine this we need some appreciation of the particular circumstances in which the speaker is ascribing knowledge.

In spelling out this contextualist view of knowledge I have made use of the notion of a relevant alternative. A relevant alternative to  $p$  is one that a person would need to be in a position to rule out if he is to count as knowing on that particular understanding of what it is to know that  $p$ . An alternative which is irrelevant on one occasion for ascribing knowledge to a person might nonetheless be relevant in other circumstances for ascribing knowledge to that same person.

I have contrasted a contextualist account of knowledge with an invariantist account, and have argued that a contextualist account is better placed to make sense of certain aspects of our epistemic and discursive practices. I have also argued that adopting an invariantist account of knowledge is liable to compel us in the direction of an implausible conception of knowledge according to which a person will only count as knowing that  $p$  if he is in a position to rule out any conceivable alternative to  $p$ . Further reflection might turn up alternative strategies for an invariantist wanting to make sense of the examples discussed. However, in the absence of such an invariantist account it seems that the contextualist's ability to better make sense of these examples provide *prima facie* support for a contextualist account of knowledge.

## Chapter 3. An Austin-inspired response to the 'argument from ignorance'

In this chapter I set out an Austin-inspired response to a sceptical puzzle occasioned by what has come to be known as 'the argument from ignorance'.<sup>53</sup> The argument from ignorance is a form of argument that can appear to expose something of an incoherence within our epistemic and discursive practices. Beginning with some candidate fact  $p$  to be known, and some incompatible alternative  $q$ , an argument from ignorance generally takes the following form:

- P1. S doesn't know that not- $q$ .
- P2. If S doesn't know that not- $q$ , S doesn't know that  $p$ .
- C. Therefore, S doesn't know that  $p$ .

The reason that such an argumentative form has been thought to occasion a sceptical puzzle about knowledge<sup>54</sup> is that it seems that for any situation in which we might ordinarily say that a person knows something  $p$ , we can conceive of alternatives to  $p$  that the person doesn't appear to be in a position to know not to obtain. If the reasoning employed in the argument from ignorance is correct, this would suggest that in each of these cases the person isn't in a position to know that  $p$  after all. That we nonetheless feel that these are cases in which the person might count as knowing that  $p$  would suggest that there is some incoherence in our epistemic and discursive practices surrounding ascriptions of knowledge.

My aim in this chapter is to put forward an Austin-inspired strategy for diagnosing the intuitive appeal of such arguments while showing that there is in fact no incoherence within our epistemic and discursive practices. The strategy I outline begins with Austin's suggestion that in saying that a person knows something  $p$ , we are saying that the person has done 'enough' for present purposes to establish that  $p$ . In the previous chapter I argued that what counts as 'enough' can

---

<sup>53</sup> The label 'argument from ignorance' comes from Keith DeRose (1995), though the argument itself has been around for much longer. The argument from ignorance should not be confused with the argument from illusion. While there are some commonalities between the two arguments, the argument from illusion has traditionally been used to argue for the view that our experience is never 'of the world' but always of something else (e.g. sense-data). Austin's *Sense and Sensibilia* (1962) remains one of the best critiques of such uses of the argument from illusion. For a more contemporary defence of Austin's thoughts in this area, see Part 1 of Putnam (1999).

<sup>54</sup> I am indebted to Kaplan (2011) for first bringing to my attention how distinctive the sceptical puzzle posed by arguments from ignorance is from other sceptical puzzles.

vary depending on the circumstances for considering the matter. In many cases in which someone is said to know that there is a goldfinch in the garden, we couldn't reasonably take the speaker to be saying that the person was in a position to distinguish his being presented with a goldfinch from his being presented with a robotic imitation. Unless we have some reason to seriously suppose the person might in fact have been presented with a robotic imitation, we won't take this to be a relevant alternative for the purposes of considering whether or not the person knows that he has seen a goldfinch. So far it seems as though there needn't be any incoherence within the practice. But an argument from ignorance might be constructed as follows:

- P1. S doesn't know that what he sees isn't a robotic imitation.
- P2. If S doesn't know that what he sees isn't a robotic imitation, S doesn't know that what he sees is a goldfinch.
- C. Therefore, S doesn't know that what he sees is a goldfinch.

Such an argument occasions a sceptical puzzle regarding our epistemic and discursive practices because it appears to establish *invariantly* that the person is in no position to know that what he sees is a goldfinch. Regardless of whether S's having seen a robotic imitation is considered to be a relevant alternative, it can look as though S is in no position to *know* that what he sees isn't a robotic imitation. But if we hold that S *doesn't* know that what he sees isn't a robotic imitation, it sounds contradictory to hold that he knows that what he sees is a goldfinch. In that case we seem to be compelled in the direction of admitting that S is invariantly in no position to know that what he sees is a goldfinch.

I suggest that Austin furnishes us with the means of making sense of how the first premise could fail to be true without needing to insist that the S does in fact know that what he sees isn't a robotic imitation. The key again is Austin's suggestion that in saying that a person knows that *p* we are saying that he has done 'enough' for present purposes to establish that *p*. If the circumstances are such that we wouldn't think it apt to seriously suppose that S might in fact have seen a robotic goldfinch, this amounts to saying that its not being a robotic goldfinch isn't something that needs establishing for present purposes. In that case it is unclear what (determinate) sense we might make of S's having done or having failed to do 'enough' for present purposes to establish that he hasn't seen a robotic goldfinch. So in that case it mightn't make any determinate sense to say either that S 'knows' or that S 'doesn't know' that what he sees isn't a robotic imitation. In that case there can be circumstances in which the first premise of the argument from ignorance needn't be true or false.

The rest of the chapter is spent accounting for the intuitive appeal of the sceptical puzzle and comparing the Austin-inspired response I have outlined with two other responses to the puzzle. The intuitive appeal of the argument comes from the fact that, while we can readily imagine circumstances in which it would be true to say that someone in S's position *doesn't know* that what he sees isn't a robotic imitation, we cannot readily imagine circumstances in which it would be true to say that someone in S's position *knows* it isn't a robotic imitation. Forced to judge one way or the other we are liable to feel compelled to say that S *doesn't* know that what he sees isn't a robotic imitation. What we fail to recognise is that such a judgment only makes determinate sense in certain circumstances.

I then consider Mark Kaplan's Austin-inspired response to the sceptical puzzle occasioned by arguments from ignorance, a response that resembles my own in certain respects. Both of us think that Austin furnishes us with means for making sense of the idea that it mightn't be true in certain circumstances to say either that a person 'knows' or that he 'doesn't know' that, for example, what he sees isn't a robotic imitation. But Kaplan holds that this is because Austin makes plausible the idea that there might be circumstances in which it would be *false* to say either that the person 'knows' or that he 'doesn't know' that what he sees isn't a robotic imitation. I on the other hand am suggesting that it might be *neither* true nor false to say that the person 'doesn't know' in certain circumstances. This is a substantial difference between mine and Kaplan's respective Austin-inspired accounts, one that calls for further consideration than I have been able to devote to the matter here.

I end by comparing the Austin-inspired account I have presented with David Lewis's contextualist response to arguments from ignorance. While Lewis's contextualist response resembles in very broad terms the response I have outlined, there are two aspects of his response in particular that I find problematic. First, Lewis's 'Rule of Attention' appears to get the mechanics of our epistemic and discursive practices wrong. As Austin suggests, it is not enough to 'make' an alternative relevant simply to draw attention to it. Rather, we need to have some reason to seriously suppose that the alternative might in fact obtain. Second, Lewis thinks it an advantage of his account that it allows him to retain the view that knowledge is closed under known entailment. I argue that Lewis holds onto the idea that knowledge is closed under known entailment only at the cost of accepting a strange view of how the concept of knowledge functions in practice.

### **3.1 The argument from ignorance**

In the previous chapter I argued that Austinian considerations speak in favour of a contextualist conception of knowledge according to which a person needn't be in a position to rule out any conceivable alternative in order to count as knowing something. Instead a person need only be in a position to rule out an alternative if it is a relevant alternative on that occasion for judging whether or not he knows, where what counts as a relevant alternative is liable to vary depending on the circumstances in which the ascription of knowledge is made. In other words, even if it is conceivable that Ollie could be presented with a robotic imitation instead of a goldfinch, and Ollie is in no position to distinguish between his being presented with a goldfinch and his being presented with a robotic imitation, this does not mean that Ollie cannot count as knowing that what he sees is a goldfinch on some occasions for judging the matter. These will be occasions on which Ollie's looking at a robotic imitation is not a relevant alternative for the purposes of judging whether or not he knows that what he sees is a goldfinch.

While I think this is a plausible account of what it is to know something, there remains the question of why we are liable to find the following line of reasoning compelling: *Say what you like about some alternatives being 'relevant' and others not, the fact of the matter is that Ollie is in no position to know that what he sees is not a robotic imitation. And insofar as he doesn't know this, he doesn't know that what he sees is a goldfinch. One can't know that something is a goldfinch if one doesn't know it isn't a robotic imitation.* Two claims are made here. The first is that Ollie is in no position to know that what he takes to be a goldfinch is not instead a robotic imitation (at least not without investigating further). The second claim is that if Ollie doesn't know that what he sees is not a robotic imitation, then he is in no position to know that what he sees is a goldfinch. These two claims are together taken to show that Ollie is in no position to know that what he sees is a goldfinch, regardless of who is considering the matter and in what circumstances.

Why might we find this line of reasoning *prima facie* plausible, and feel that it shows that Ollie is (contra our judgement in some circumstances) in no position at all to know that what he sees is a goldfinch unless he can rule out<sup>55</sup> his seeing a robotic imitation? I will begin with the first claim, namely that Ollie is in no position to know that what he takes to be a goldfinch is not instead a robotic imitation. Many epistemologists have noted that it sounds wrong to say that someone in

---

<sup>55</sup> A reminder from earlier chapters: as I am using the expression, Ollie is in a position to 'rule out' an alternative if he is in a position to distinguish his experience were the alternative to obtain from his experience were things as he supposes them to be. In this particular case Ollie would be in a position to rule out his having seen a robotic imitation instead of a goldfinch if he was in a position to distinguish between the course of his experience in each case.

Ollie's position knows that he isn't looking at a robotic imitation, as opposed to a goldfinch. One plausible explanation for this intuition is that in judging that a person knows that something he takes to be a goldfinch is not in fact a robotic imitation, we would ordinarily be understood to be saying that the person is in a position to in some way show that what he sees is not a robotic imitation (e.g. that he's tapped the bird, or performed some other test). Ollie is in no position to show that what he is looking at is not a robotic imitation. If he were to be presented with a robotic imitation instead of a goldfinch, he would be in no position to distinguish his experience from his experience in a situation where he is presented with a goldfinch. So whatever else we might want to say about Ollie in this situation, it seems wrong to say that he *knows* that what he sees isn't a robotic imitation. Forced to say whether Ollie knows or doesn't know that what he sees is not a robotic imitation, we are liable to feel compelled to say that Ollie does not know.

Suppose we grant for the moment that this intuition is correct, and that Ollie does not know that what he is looking at is not a robotic imitation. Can we perhaps avoid the conclusion that Ollie is in no position to know that what he is looking at is a goldfinch by denying the second claim? The second claim, recall, is that if Ollie doesn't know that what he is looking at is a robotic imitation, then he doesn't know that what he is looking at is a goldfinch. Denying this would appear to be tantamount to saying that Ollie might simultaneously count as knowing that he is looking at a goldfinch, while not knowing that what he is looking at is not a robotic imitation. Such a view seems implausible. One reason for thinking this is that it just *sounds* contradictory to judge both that Ollie knows that he is looking at a goldfinch, and that he doesn't know that what he is looking at is not a robotic imitation. The latter judgement just *sounds* as though it clashes with the former judgement.<sup>56</sup> Furthermore, it certainly seems that in any ordinary situation in which a person is judged not to know that what he is looking at is not a robotic imitation, this will be taken to mean that he doesn't know that what he is looking at is a goldfinch (or a bird at all). Such considerations provide at least some support for the view that Ollie is in no position to know that he is looking at a goldfinch if he doesn't know that what he is looking at is not a robotic imitation. If this view is mistaken it will at least take some showing.

---

<sup>56</sup> Some philosophers will want to say that the reason this *sounds* wrong is that knowledge is 'closed under known entailment'. The view that knowledge is closed under known entailment can be expressed in the following principle: If S knows that *p*, and S knows that *p* entails *q*, then S must also know that *q*. Applying this to the case in question: If Ollie knows that what he sees is a goldfinch, and Ollie knows that its being a goldfinch entails its not being a robotic imitation, then Ollie must also know that what he sees is not a robotic imitation. My reason for not relying on this explanation is that I don't agree that knowledge is closed under known implication. Nonetheless I *do* agree that if Ollie doesn't know that what he sees isn't a robotic imitation, then he doesn't know that what he sees is a goldfinch. This might seem to put me in an impossible bind. The burden of the next section is to suggest that Austin provides us with a way of avoiding this bind. I directly address the issue of whether knowledge is closed under known entailment in section 3.4.



If, on the other hand, we feel that both of the claims I have considered are correct, then it seems we have the resources to construct an argument showing that Ollie is in no position to know that what he sees is a goldfinch:

- P1. Ollie does not know that what he sees is not a robotic imitation.
- P2. If Ollie doesn't know that what he sees is not a robotic imitation, then he doesn't know that it is a goldfinch.
- C. Therefore, Ollie does not know that what he sees is a goldfinch.

Following Keith DeRose (1995) I will refer to arguments of this form as 'arguments from ignorance'. The kind of considerations adduced in constructing an argument of this form can be deployed in relation to just about any candidate case for knowledge. To construct such an argument we simply need to locate a conceivable alternative that the person is not in a position to rule out at that point in time. And in just about any situation in which a person takes something to be so there will be conceivable alternatives he is in no position to rule out. The reasoning used in the above argument can then be deployed to suggest that in each case the person is not in a position to know the thing in question. If this is so, then the above argument appears to have sceptical implications. If successful, this argument appears to show that we are not actually in a position to know much if anything of what we ordinarily take ourselves to.

There are different approaches we might take in responding to an argument like this, even if we accept something like the contextualist account I presented in the previous chapter.<sup>57</sup> In the next section I present an Austin-inspired strategy for resisting such arguments. In addition to showing how we might resist an argument from ignorance, it is an advantage of the approach I will present that it can at the same time account for some of the intuitions we are liable to have when presented with such arguments.<sup>58</sup> I will argue that such arguments appear to compel us in the direction of a sceptical conception of knowledge only because we fail to properly appreciate the context-sensitivity of questions about knowledge.

---

<sup>57</sup> Later in this chapter I will compare my preferred contextualist response to another contextualist response given by David Lewis.

<sup>58</sup> DeRose (1995, p. 3) likewise thinks it incumbent on responses to sceptical arguments that they enable us to account for the puzzlement/conflicting intuitions they are liable to generate.

### **3.2 An Austin-inspired response to the argument from ignorance**

In this section I suggest how Austinian considerations might be deployed to show that arguments from ignorance like the one presented in the previous section need not compel us in the direction of a sceptical conception of knowledge. Austin observes that in saying that a person knows that  $p$ , we are saying that he has done enough for present intents and purposes to establish that  $p$ . Likewise in saying that a person doesn't know that  $p$ , we are say that he hasn't done enough for present intents and purposes to establish that  $p$ . Suppose, however, that we are in circumstances in which  $p$  is not something we have any reason to seriously suppose *needs* establishing for present purposes. In that case it is not clear what determinate sense we might make of the suggestion that a person either has or hasn't done 'enough' for present purposes to establish that  $p$ . And by extension it is not clear what determinate sense we might make of the suggestion that the person does or doesn't 'know' that  $p$ . I take this idea and argue that such might be the case with respect to the question of whether a person 'knows' or 'doesn't know' that a particular alternative does not obtain, at least in circumstances in which we have no reason to seriously suppose that it might in fact obtain. Where we don't think that it needs to be established for present purposes that an alternative does not obtain, it is not clear what determinate sense we might make of a person's having done or having failed to do 'enough' *for present purposes* to establish that the alternative does not obtain. In that case there mightn't be any determinate sense to the suggestion that the person 'knows' or 'doesn't know' that the alternative does not obtain. If so, it won't be true in such circumstances to say that he 'doesn't know' that the alternative does not obtain. If it isn't true to say that Ollie 'doesn't know' that what he sees isn't a robotic imitation, the suggestion that he 'doesn't know' poses no obstacle to his knowing that what he sees is a goldfinch.

In the previous chapter<sup>59</sup> I employed the notion of a 'relevant alternative' to articulate one aspect of Austin's observations regarding everyday epistemic and discursive practice as it relates to the concept of knowledge. I argued that there is a distinction implicit in our practices (as detailed by Austin) between *merely* conceivable alternatives, and alternatives we think ought seriously to be treated as ways things might in fact be.<sup>60</sup> And I presented an account of knowledge according to which a person's not being in a position to rule out a particular alternative will only count against what a speaker says in saying that the person knows if the alternative is considered to be a relevant alternative in those circumstances.

---

<sup>59</sup> See 2.1.

<sup>60</sup> Austin talks about having 'concrete' or 'special' reason to think that a particular alternative might obtain.

Suppose Jamie and Alistair are preparing lunch and Jamie asks Alistair whether he knows if there is any cold water in the fridge. Alistair replies that there is cold water in the fridge; that he has just looked. In this case we can imagine Jamie pushing Alistair on the matter: 'Did you *check* that it was cold?' If Alistair has not done enough to show that the water is cold (e.g. put his hand to the bottle), then we would think it apt for Jamie to say that Alistair doesn't know there is cold water in the fridge. There is a relevant alternative, namely that the water is warm, that Alistair is in no position to rule out. Suppose instead that Alistair has put his hand to the bottle and tells Jamie so. This would generally be the end of the matter and we would think that both Alistair and Jamie were right in taking themselves to know that there is cold water in the fridge. Alistair, however, is in no position to show that it isn't vodka in the bottle. It certainly isn't inconceivable that the contents of the bottle could be vodka instead of water. So there is a conceivable alternative to there being cold water in the fridge that Alistair is in no position to rule out. Nonetheless we would think Alistair had accurately presented his position to Jamie in telling her that he knows there is cold water in the fridge, unless there was some special reason to think it might in fact be vodka in the bottle.

In the example described, Jamie might reasonably take Alistair to be saying that he has done enough to establish that the water is not warm in saying that he knows there is cold water in the fridge. However, we would think Jamie unreasonable if she were to take Alistair to be saying that he had done anything to establish that it is not vodka in the bottle, unless they had some special reason to think this might be the case.<sup>61</sup> Why can Jamie reasonably expect that Alistair be able to show one alternative not to obtain, but not the other? One plausible explanation: we think it apt for someone in Jamie and Alistair's circumstances to seriously treat warm water in the bottle as a way things might in fact be, whereas we don't think it apt for someone in their circumstances to seriously treat vodka in the bottle as a way things might in fact be. That is, we might think a person in Alistair's situation somewhat rash in supposing that there is cold water in the fridge without checking that the water is in fact cold, but we wouldn't think it rash of someone in Alistair's situation to suppose that there is cold water in the fridge without checking that it isn't vodka.

On Austin's account, to say that a person knows that  $p$  is to say that he has done enough (for present purposes) to establish that  $p$ . In other words to say that Alistair knows there is cold water in the fridge is to say that he has done enough to establish that there is cold water in the fridge.

---

<sup>61</sup> They might have a special reason to suppose there might be vodka in the bottle if their housemates were in the habit of storing vodka in the water bottles.

In Jamie and Alistair's circumstances we would generally think it enough to show this if Alistair has looked in the fridge and seen that the water bottle is full, and perhaps held his hand to it to check that it is cold. Suppose now that someone were to say that Alistair knows it is not vodka in the water bottle. Such a person would generally be taken to be saying that Alistair has done something to establish that it is not vodka in the bottle. Given that Alistair has done no such thing, it doesn't seem apt to say that Alistair knows it is not vodka in the bottle. On this point Austinian considerations appear to support the intuition that it is wrong to say that Alistair knows it is not vodka in the bottle. But in that case it seems we can construct an argument from ignorance as follows:

- P1. Alistair doesn't know that it is not vodka in the bottle.
- P2. If Alistair doesn't know that it is not vodka in the bottle, he doesn't know that it is cold water in the bottle.
- C. Therefore, Alistair is in no position to know that it is cold water in the bottle.

If we think there are occasions (e.g. Jamie and Alistair's) on which it would be true to say that Alistair knows there is cold water in the fridge, how are we to respond to such an argument? The key, I suggest, is making plausible the idea that there can be circumstances in which it would neither be true to say that Alistair knows it is not vodka in the bottle *nor* that he doesn't know this.<sup>62</sup>

Why might we think that neither statement is true in Jamie and Alistair's circumstances? Building on the relevant alternatives account I have gleaned from Austin's work, I want to investigate the prospects for a view according to which judgements regarding whether or not a person knows an alternative *not* to obtain only make determinate sense (say something true or false) in circumstances in which it is apt to treat this alternative as something in need of ruling out for present intents and purposes. Apart from such circumstances there is, on this view, no motivated answer to the question of whether or not the person knows the alternative not to obtain. Suppose, as Austin's remarks suggest, there are ordinary circumstances in which we would not think it apt to seriously treat vodka in the bottle as a way things might be, one in need of ruling out for present intents and purposes. On the view to be investigated these would be circumstances in which it is

---

<sup>62</sup> Such an idea is not unprecedented. Putnam (2001, p. 9) has likewise suggested that the most promising route for a contextualist is to deny that it makes sense in many circumstances to say either that we know *or* that we don't know that the sceptical alternatives do not obtain. Later in this chapter I compare my own suggestion for how to understand this idea with a similar proposal from Mark Kaplan (who also takes Austin's work as his starting point).

neither true to say that a person knows nor that he doesn't know that there isn't vodka in the bottle.

Here is the view I want to put forward in brief. To say that S knows that not-*p* is to say that S has done what is needed to establish that not-*p* for present purposes (Austin's observation). Conversely, to say that S does not know that not-*p* is to say that S has not done what is needed to establish that not-*p* for present purposes. This leaves room for a third category. Suppose we don't think that not-*p* *needs* establishing for present purposes, and this because we take ourselves to have no reason to seriously suppose that maybe-*p* for present intents and purposes. In that case it wouldn't make sense to suggest that S has not done enough to establish that not-*p* for present purposes. After all, in this situation we don't think that *p* is to be seriously treated as a way things might be for present purposes, and so it makes no sense to talk about doing 'enough' to establish that not-*p* for present purposes. Insofar as judging that S does not know that not-*p* is tantamount to saying that S has not done enough to establish that not-*p* for present purposes, it is not clear what could be meant in saying that S doesn't know that not-*p* in such circumstances. What more exactly could we be saying that S *ought* to have done for present purposes in order to show that not-*p* if we don't think that not-*p* needs to be shown for present purposes? But nor will it make sense to say in such circumstances that S knows that not-*p*. Insofar as we don't think that not-*p* needs showing for present purposes, it makes no sense to talk about S having done 'enough' for present purposes to establish that not-*p*.

We can apply this suggestion in relation to the above example. In the circumstances described Jamie and Alistair have no reason to think that there might in fact be vodka in the water bottle. We can easily contrast this with circumstances in which we *would* think that Jamie and Alistair have reason to think there might be vodka in the water bottle. Austin suggests that saying that Alistair knows it is not vodka in the bottle amounts to saying that he has done *enough* for present purposes to establish that it is not vodka in the bottle. But that it is not vodka in the bottle is not something we think needs to be established for Jamie and Alistair's purposes. In that case it's not clear what someone in Jamie and Alistair's circumstances could mean in saying that Alistair has done 'enough' to establish this for their purposes. So it makes no determinate sense in these circumstances to say that Alistair knows it is not vodka in the bottle. For much the same reasons, however, it makes no determinate sense in these circumstances to say that Alistair *doesn't* know it is not vodka in the bottle. To say as much would be to say that Alistair hasn't done enough for his and Jamie's purposes to establish that it is not vodka in the bottle. But given that it doesn't for their purposes need to be established that it isn't vodka in the bottle, it makes no sense to talk

about failing to do ‘enough’ to show this for their purposes. So it makes no sense to say that Alistair doesn’t know that it isn’t vodka in the bottle.

Suppose for a moment that this is right, and that it makes no sense in Alistair and Jamie’s circumstances to say either that Alistair knows it isn’t vodka in the bottle *or* that he doesn’t know this. In that case we have the means for showing that the above argument from ignorance does not threaten Alistair’s claim to know that there is water in the bottle. If it makes no sense in Jamie and Alistair’s circumstances to say that Alistair doesn’t know it isn’t vodka in the bottle, then someone in their circumstances wouldn’t be saying something true in saying that Alistair doesn’t know this. And that means the first premise of the argument from ignorance isn’t true. However, it is important that in saying that the first premise isn’t true I am not saying that it is *false*. The suggestion is that if it makes no sense in Alistair and Jamie’s circumstances to say that Alistair doesn’t know that it isn’t vodka in the bottle, then it is neither true nor false in Jamie and Alistair’s circumstances to say as much.

If this strategy works in relation to the above example, then the same strategy can be employed to resist other instances of the argument from ignorance. Consider again the example of Ollie and the goldfinch. If there are circumstances in which we ordinarily don’t think it needs to be shown that what Ollie takes to be a goldfinch is not in fact a robotic imitation, then in such circumstances it won’t make sense to say either that Ollie knows it isn’t a robotic imitation or that he doesn’t know this. If we don’t think that it needs to be shown for present purposes that the bird isn’t a robotic imitation, then it makes no sense to talk about Ollie having done ‘enough’ for present purposes to show that the bird isn’t a robotic imitation. For that reason it makes no sense, on my Austin-inspired account, to say that Ollie knows it is not a robotic imitation. Nor does it make any sense in these circumstances to talk about Ollie having *failed* to do ‘enough’ for present purposes to show that the bird isn’t a robotic imitation. So it makes no sense to say that he doesn’t know it isn’t a robotic imitation. In a situation where it makes no sense to say that Ollie doesn’t know that what he sees isn’t a robotic imitation, saying this won’t undermine Ollie’s claim to know that what he sees is a goldfinch.

If this suggestion is on the right track, it makes sense that in many ordinary situations we wouldn’t think it apt to respond to the question ‘Do you know the ‘goldfinch’ isn’t a robotic imitation?’ by saying either that we do know or that we don’t know. In many situations we would think it reasonable for a person who was asked this question to respond with a question of his own: ‘Why should I suppose that it’s a robotic imitation?’ If the would-be challenger was unable to provide a good reason for taking seriously in present circumstances the suggestion that it might have been

a robotic imitation that was seen, we would generally take his failure to do so as grounds for dismissing his question, along with the suggestion that the bird ‘might’ have been a robotic imitation. This is very different to attempting to provide an answer to the question. On the present account of ascriptions of knowledge it makes sense that we wouldn’t in such situations try to answer the question. If the situation is one in which we do not think it apt to treat the possibility that what was seen was a robotic imitation as something that for present purposes needs to be shown not to be so, then on the present account there mightn’t be any determinate sense to the question ‘Do you know it isn’t a robotic imitation?’

In summary, the suggestion here is that it only makes sense to say that someone knows or doesn’t know that  $p$  in circumstances in which we think it apt to treat  $p$  as something in need of establishing for present intents and purposes. On the Austinian assumption that there are in any situation conceivable possibilities we do not think it apt to treat as things we need to establish not to obtain, it won’t in such circumstances make sense to say either that we know or that we don’t know these possibilities not to obtain. This provides us with a response to the sceptical puzzle occasioned by certain uses of the argument from ignorance. Suppose that  $p$  and  $q$  are alternatives. The argument from ignorance generates a sceptical puzzle only on the assumption that on any occasion for judging that S knows that  $p$ , there will be some  $q$  such that S doesn’t know that not- $q$ . However, we needn’t think that this is so. If the circumstances are such that not- $q$  is not something we think needs to be established for present purposes, then it won’t be true to say either that S knows or that S doesn’t know that not- $q$ . In that case S’s claim to know that  $p$  cannot in such circumstances be undermined by the suggestion that S doesn’t know that not- $q$ . The final piece in the response: in many of the situations where we would ordinarily think it apt to judge that some person S knows something  $p$ , there is no  $q$  that we think needs to be established not to obtain for present purposes. If this is so, arguments from ignorance pose no threat to our everyday claims to know.

### **3.3 Why do arguments from ignorance seem compelling?**

In this section I will consider why arguments from ignorance so much as seem to compel us in the direction of a sceptical conception of knowledge. The diagnosis I will suggest is that we feel that arguments from ignorance pose a sceptical threat to our everyday ascriptions of knowledge in large part because we overlook the context-sensitivity of questions about knowledge. Insofar as we overlook the context-sensitivity of ascriptions of knowledge, it is not at all surprising that we should find ourselves having the intuitions many of us are liable to have when presented with

an argument from ignorance. The Austin-inspired account I have outlined can help us make sense of why, when we suppose that there must be an answer to the question of whether or not a person 'knows' some alternative not to obtain, we are liable to feel that the person is not in a position to know. The reason is that, while we can readily imagine circumstances in which it *would* be true to say that the person doesn't know the alternative not to obtain, we cannot readily imagine circumstances in which it would be true to say that person knows the alternative not to obtain. However, the Austin-inspired account I have outlined allows us to make sense of the idea that, although there are imaginable circumstances in which it would be true to say that the person doesn't know that some alternative doesn't obtain, there might also be other circumstances in which it is neither true nor false to say that the person 'doesn't know' that the alternative doesn't obtain.

Consider again the example from 2.2 involving Ollie and the goldfinch. In the circumstances in which Dan tells Elle that Ollie knew that what he saw was a goldfinch, it would be incomprehending of Elle if she were to understand Dan to be saying that Ollie was in a position to rule out his having seen a robotic imitation. In such circumstances we wouldn't think that Ollie's having seen a robotic imitation was a relevant alternative, or that Ollie's not having been in a position to rule out the alternative counts against the truth of what Dan says in saying that Ollie knew. However, the following argument from ignorance purports to show that Dan couldn't have said something true in saying that Ollie knows:

- P1. Ollie doesn't know that what he sees isn't a robotic imitation.
- P2. If Ollie doesn't know that what he sees isn't a robotic imitation, then he doesn't know that what he sees is a goldfinch.
- C. Therefore, Ollie doesn't know that what he sees is a goldfinch.

Arguments of this kind have struck epistemologists like Stroud and others as providing intuitively compelling grounds for thinking that it could never be true to say of Ollie that he knows he is looking at a goldfinch. Forced to choose between saying either that Ollie knows he isn't looking at a robotic imitation or that Ollie doesn't know this, many epistemologists feel compelled to agree that Ollie *doesn't* know. And having been brought to judge that Ollie doesn't know that what he sees isn't a robotic imitation, many epistemologists also feel that it is wrong to say that Ollie knows that what he sees is a goldfinch. Instead it seems that having admitted that Ollie doesn't know that what he sees isn't a robotic imitation, we should agree with the conclusion that Ollie doesn't know that what he sees is a goldfinch. But we started off thinking that Ollie's position



was such that he might correctly be said to know that he was looking at a goldfinch ordinarily! So these intuitively plausible steps have led us to a seemingly paradoxical conclusion in relation to this particular case. What's more, it is clear that the same reasoning can be deployed to the same effect in relation to just about any candidate example of someone's being in a position to know something.

The Austin-inspired strategy outlined in the previous section enables us to make sense of our intuitions here while allowing us to see why such an argument need not compel us in the direction of a sceptical conception of knowledge.

Starting with the second premise of the argument. We are liable to feel that if we judge that Ollie doesn't know that what he sees isn't a robotic imitation, then we should judge that Ollie doesn't know that what he sees is a goldfinch. The strategy outlined in the previous section does not require us to revise this intuition. In fact on my Austin-inspired account this intuition is correct. In judging that Ollie doesn't know that what he sees is not a robotic imitation, we are judging that Ollie has not done enough for present intents and purposes to establish that what he sees is not a robotic imitation. In such a case we are treating Ollie's having seen a robotic imitation (instead of a goldfinch) as something in need of ruling out for present intents and purposes. In other words Ollie's having seen a robotic imitation is a *relevant alternative* to Ollie's having seen a goldfinch. Given that there is a relevant alternative that Ollie is in no position to show not to obtain, Ollie doesn't know that what he sees is a goldfinch (he hasn't done enough for *present intents and purposes* to show that it is a goldfinch). In other words, on any occasion on which it is true to say that Ollie doesn't know that what he sees is not a robotic imitation, it will also be true to say that Ollie doesn't know that what he sees is a goldfinch.

This means that if there are occasions on which it is true to say that Ollie knows that what he sees is a goldfinch, there must be occasions on which it is not true to say that he doesn't know that what he sees is not a robotic imitation. This is an ugly way of putting the point, but necessary if the Austin-inspired strategy I have outlined is correct. Key to that strategy is the idea that we can imagine circumstances in which it would be neither true nor false to say 'Ollie doesn't know that what he sees isn't a robotic imitation'. The advantage of such a strategy is that it allows us to do justice to what are many people's initial intuitions with respect to the first premise in the above argument.

There are two common intuitions in relation to the first premise which need accounting for. The first is that it is wrong to say that Ollie knows that what he sees isn't a robotic imitation. The

second intuition is that, when forced to decide one way or the other, it seems we need to say that Ollie *doesn't* know this. If the Austin-inspired strategy I have sketched is correct, we can make sense of why we are liable to have each of these intuitions. Suppose Ollie is looking at a goldfinch in broad daylight. As things stand, we cannot readily imagine an occasion for aptly saying that Ollie knows that what he sees is not a robotic imitation. For this to be correct the circumstances would need to be such that it was apt to treat Ollie's seeing a robotic imitation as something in need of ruling out for present purposes, and in which Ollie had done enough for those purposes to establish that it wasn't a robotic imitation. But it doesn't look as though Ollie is in a position to establish that what he sees is not a robotic imitation (he hasn't tapped the 'bird' or performed any similar test). So it is difficult to imagine a situation in which we would intuitively be saying something true in saying that Ollie knows that what he sees isn't a robotic imitation. On the other hand we can imagine circumstances in which it would be apt to say that Ollie *doesn't* know that what he sees isn't a robotic imitation. All we need to do is imagine circumstances in which we have a reason to seriously treat robotic imitations as something in need of ruling out for present purposes. Such circumstances will likely be outlandish, but are not inconceivable. In such circumstances it would be true to say that Ollie doesn't know that what he sees is not a robotic imitation, and this because he hasn't done enough (or anything) for present purposes to show that this is so.

We can now see why, when compelled to say either that Ollie knows or doesn't know that what he sees isn't a robotic imitation, we are liable to feel that it is closer to the truth to say that Ollie *doesn't* know this. On the present account we are liable to have this intuition because we can readily imagine occasions on which it would be true to say that Ollie doesn't know that what he sees isn't a robotic imitation, but we cannot readily imagine occasions on which it would be true to say that Ollie knows that what he sees isn't a robotic imitation. One of the chief advantages of the strategy I have outlined is that it allows us to make sense of these intuitions while at the same time seeing that they need not compel us in the direction of a sceptical conception of knowledge. We are liable to feel as though these intuitions push us in the direction of sceptical conception of knowledge only when we disregard the context-sensitivity of questions about knowledge, and in turn disregard the particular circumstances in which a person might be said to know something.

When we disregard the context-sensitivity of ascriptions of knowledge, we are liable to think that if it is true in some circumstances to say that Ollie doesn't know that what he sees isn't a robotic imitation, then it will be true to say this regardless of the circumstances for considering the matter. And it will then seem as though our only means of salvaging the idea that a person could speak

truth in saying that Ollie knows that what he sees is a goldfinch is to adopt the seemingly contradictory view that Ollie can know while not knowing that what he sees isn't a robotic imitation.

If we take seriously the context-sensitivity of ascriptions of knowledge, another route opens up. When we attend to the different circumstances in which someone might have an interest in judging the state of Ollie's knowledge, we see that there can be circumstances in which it is not true to say that Ollie doesn't know that what he sees is not a robotic imitation. But not because we can imagine circumstances in which it would be true to say that Ollie (as he stands) *does* know this. Rather it is because in attending to particular occasions for considering the state of Ollie's knowledge we can see how there could be circumstances in which it isn't true to say either that Ollie knows *or* that he doesn't know that what he sees isn't a robotic imitation. These will be circumstances in which we do not think it apt to seriously treat robotic imitations as a possibility in need of ruling out for present purposes, and in which it thus makes no sense to talk about a person having done enough for present purposes to establish that the alternative does not obtain. This allows us to preserve the intuition that there can be circumstances in which it is true to say that Ollie knows that what he sees is a goldfinch.

In summary, the Austin-inspired account of ascriptions of knowledge outlined in the previous section allows us to make sense of the intuitions we are liable to have when confronted with the argument from ignorance, intuitions which seem to compel us in the direction of a sceptical conception of knowledge. Austin also provides us with the means for showing why these intuitions needn't compel us in the direction of a sceptical conception of knowledge. The key is attending to the particular circumstances in which a person might be said to know or not know something, as Austin encourages us to do. When we do this it becomes plausible that there can be circumstances in which it would be neither true nor false to say that someone knows something. This in turn furnishes us with the means for avoiding the sceptical implications of arguments from ignorance.

### **3.4 What about closure?**

If the strategy I have outlined is right, then we will need to reject the view that knowledge is closed under known implication. For my purposes the closure principle can be formulated as

follows: if S knows that  $p$ , and S knows that  $p$  entails  $q$ , then S must also know that  $q$ .<sup>63</sup> We can suppose that Ollie knows that if what he sees is a goldfinch, then it is not a robotic imitation. The Austin-inspired account I have presented suggests there can be situations in which it is true to say that Ollie knows that what he sees is a goldfinch, but in which it is not true to say that Ollie knows that what he sees isn't a robotic imitation. But nor will it, on my account, be true in such circumstances to say that Ollie *doesn't* know that what he sees isn't a robotic imitation. It is often suggested that rejecting the closure principle commits us to the seemingly absurd view that statements like 'Ollie knows that what he sees is a goldfinch, but he doesn't know that it isn't a robotic imitation' might in some circumstances be true. My Austin-inspired proposal does not commit us to any such view. Nor are we required to give up the following intuitively plausible idea: if Ollie takes himself to be seeing a goldfinch, then he is committed to its not being a robotic imitation. We thus needn't suppose that the present rejection of the closure principle saddles us with any obviously absurd commitments.

It is sometimes suggested that it is our commitment to a closure principle on knowledge that accounts for the intuitive appeal of the second premise in arguments from ignorance. Suppose Ollie knows that if what he sees is a goldfinch, then it is not a robotic imitation. The second premise in an argument from ignorance could take the following form: If Ollie doesn't know that what he sees is not a robotic imitation, then he doesn't know that what he sees is a goldfinch. I have said that I take the principle stated here to be correct. Any occasion on which it is true to say that Ollie doesn't know that what he sees is a robotic imitation will be one on which it is true to say that he doesn't know that what he sees is a goldfinch. Many epistemologists think it absurd to judge otherwise. Some epistemologists, however, take this intuition to show that we are committed to the view that it is only true to say that Ollie knows that what he sees is a goldfinch on occasions on which it is also true to say that he knows that what he sees isn't a robotic imitation. As we've seen, this occasions a sceptical puzzle because it seems difficult to imagine circumstances in which it would be true to say of Ollie (as he stands) that he knows that what he sees isn't a robotic imitation.

Why think that we are committed to the view that if Ollie knows that what he sees is a goldfinch, he must also know that what he sees isn't a robotic imitation? The thought seems to be that if we are not committed to this view, then we must be prepared to hold that there can be situations in

---

<sup>63</sup> The closure principle can be formulated in different ways. For example, Stroud (1984, p. 29) puts forward a particularly strong version of the principle: "if somebody knows something,  $p$ , he must know the falsity of all those things incompatible with his knowing that  $p$  (or perhaps all those things he knows to be incompatible with his knowing that  $p$ )." What I have to say about the closure principle I discuss also applies to a closure principle like Stroud's.

which Ollie knows that what he sees is a goldfinch, even though he doesn't know that what he sees isn't a robotic imitation. If these were the only two alternatives available, it is hard to see how a satisfying response to the argument from ignorance might be had. As it turns out, however, the Austin-inspired account I have sketched furnishes us with the means for rejecting the closure principle without saddling ourselves with the absurd view that Ollie might know that he has seen a goldfinch even if he doesn't know that it wasn't a robotic imitation. On the account I have outlined there can be occasions on which it is true to say that Ollie knows he is looking at a goldfinch, and where it is neither true nor false to say that Ollie knows it isn't a robotic imitation. Nonetheless my account allows that on any occasion on which it is true to say that Ollie doesn't know that what he sees isn't a robotic imitation, this will imply that Ollie doesn't know that what he sees is a goldfinch. My Austin-inspired account gives us a way of accepting the second premise in arguments from ignorance without taking this to commit us to a closure principle on knowledge.

Is there any other reason we should feel compelled to adopt a closure principle for knowledge? Some epistemologists simply feel that the closure principle is intuitively compelling in its own right. I grant that the principle does have some intuitive appeal, but this can be accounted for within my Austin-inspired view. Why might it seem plausible to say that if Ollie is to know that what he sees is a goldfinch, he must know that it is not a robotic imitation? The following explanation is available to a proponent of my Austin-inspired account: this seems plausible because we can readily imagine occasions on which it would be true to say this. These would be occasions on which we thought it apt to treat the possibility that Ollie was looking at a robotic imitation as something that needed to be *shown* not to obtain. But if Austin's observations regarding everyday epistemic practice are correct, then not every occasion will be one on which we think it apt to treat this alternative as something that needs to be shown not to obtain for present purposes. My account allows that there can be situations on which it is true to say that if Ollie is to know that what he sees is a goldfinch, he must know that it is not a robotic imitation. But we needn't think that it will in every situation be true to say this.

There is, I suspect, another reason why we are inclined to find a closure principle for knowledge intuitively appealing. The reason is that in *claiming* to know that what he sees is a goldfinch, Ollie expresses a truth commitment. In other words Ollie not only presents himself as having done enough to establish for present purposes that what he sees is a goldfinch, but he also commits himself to its being a goldfinch. Insofar as a robotic imitation is *not* a goldfinch, Ollie likewise commits himself to its not being a robotic imitation. If it should turn out that what he takes to

be a goldfinch is in fact a robotic imitation, this will mean that things are not as he has committed himself to their being. But we need to separate the commitment Ollie undertakes in *judging* that what he sees is a goldfinch from the commitment he undertakes in claiming to *know* this. In judging that what he sees is a goldfinch, Ollie plausibly commits himself both to its being a goldfinch *and* its not being a robotic imitation. Furthermore, in claiming to *know* that what he sees is a goldfinch, Ollie expresses his commitment to the fact that he has done enough (for present intents and purposes) to establish that the bird is a goldfinch. But if this is a situation in which there is no reason to think that Ollie might in fact have seen a robotic imitation, then Ollie needn't be understood as claiming to have done enough to establish that what he sees isn't a robotic imitation. We mistakenly think that because Ollie's commitment to his having seen a goldfinch also amounts to a commitment to his not having seen a robotic imitation, Ollie's commitment to his knowing that he has seen a goldfinch similarly commits him to his knowing that it is not a robotic imitation. The confusion arises because knowledge-claims are used to simultaneously express these two distinct commitments.

While there is no doubt more to be said about the appeal of a closure principle for knowledge, and the implications of rejecting this principle, my Austin-inspired account avoids at least one absurd view commonly thought to follow from rejecting closure. On my Austin-inspired account, rejecting closure does not commit us to the view that there might be occasions on which it is true to say, for example, that Ollie knows that what he sees is a goldfinch even though he doesn't know that it is not a robotic imitation. Apart from a desire to avoid such a view, there doesn't appear to be any other immediate and compelling motivation for adopting a closure principle for knowledge.

### **3.5 Comparison with Mark Kaplan's Austin-inspired strategy for dealing with arguments from ignorance**

In this section I consider Mark Kaplan's<sup>64</sup> Austin-inspired response to the sceptical puzzle posed by arguments from ignorance, a response that resembles in many ways my own Austin-inspired response to the puzzle. Like me, Kaplan suggests that Austin furnishes us with the means of making sense of the idea that there can be circumstances in which it is not true to say either that a person 'knows' or that he 'doesn't know' something. On both of our accounts there can be

---

<sup>64</sup> Kaplan has written a number of illuminating papers on Austin's views on knowledge in recent years. See Kaplan (2000, 2006, 2008, 2011).

circumstances in which it is not true to say either that Ollie ‘knows’ or that he ‘doesn’t know’ that what he sees is not a robotic imitation. In these circumstances the suggestion that Ollie ‘doesn’t know’ that what he sees isn’t a robotic imitation won’t serve to undermine his claim to know that what he sees is a goldfinch. However, Kaplan’s account relies on the idea that there can be situations in which it would be *false* to say either that Ollie ‘knows’ or that he ‘doesn’t know’ that what he sees isn’t a robotic imitation. This strikes me as a counterintuitive idea, and one that we needn’t commit ourselves to on my own Austin-inspired account. On the account I have presented there can be circumstances in which it is *neither* true or false to say that Ollie ‘doesn’t know’ that what he sees isn’t a robotic imitation. As I have argued, the key to making sense of this idea lies in Austin’s suggestion that knowing is a matter of having done ‘enough’ for present purposes to establish the truth of the matter.

Like me, Kaplan is interested in exploring what resources Austin furnishes us with for resisting arguments from ignorance and the sceptical conception of knowledge they appear to press upon us. And like me, he places a strong emphasis on Austin’s suggestion that in claiming to know something we are claiming to have done enough (for present intents and purposes) to establish that something is so (e.g. that there is a goldfinch in the garden). Kaplan notes that

what counts as enough is rough to be sure. But doing enough doesn’t require that we do enough to prove false every imaginable hypothesis whose truth is incompatible with our knowing that *p*. (2011, p. 55)

As Austin observes, in many situations where a person claims to know that there is a goldfinch in the garden, we wouldn’t take him to be saying that they are in a position to establish that it is not a *stuffed* goldfinch. Only in certain circumstances where we have a special reason to take seriously the suggestion that the ‘goldfinch’ might be stuffed will we expect a person to be able to show that this is not so before granting his claim to know that what he sees is a goldfinch.

While Kaplan thinks this takes us some way toward showing what is wrong with the sceptical conception of knowledge, he notes that it doesn’t yet tell us how exactly to respond to arguments from ignorance like the following:

- P1. I don’t know that it’s not a robotic goldfinch.
- P2. If I don’t know that it’s not a robotic goldfinch, then I don’t know it’s a goldfinch.
- C. Therefore, I don’t know that it’s a goldfinch.

If we assume that there are occasions on which I might speak truth in saying that I know that there is a goldfinch in the garden, even though I'm not in a position to *show* that it is not a robotic goldfinch, which of the three premises in the above argument should be rejected? Like me, Kaplan doesn't think that Austin should respond by saying that I *do* know that what I see isn't a stuffed goldfinch. To say that I know that it is not a stuffed goldfinch would be to present myself as having done something to establish that it isn't a stuffed goldfinch, and in these circumstances I am imagined to have done no such thing. But Austin thinks there can nonetheless be occasions where I speak truth in saying that I know that what I see is a goldfinch, even though I have done nothing to establish that it isn't a stuffed goldfinch. This leads Kaplan to try to eke out an alternative response on Austin's behalf.

It can, however, seem as though Austin has nowhere left to turn. If Austin doesn't think it true for me to say that I know that what I see isn't a stuffed goldfinch, then it seems as though his only other option is to admit that it would be true for me to say that I *don't* know this. In that case the first premise of the above argument will be true. This would seem to leave Austin only with the option of denying the second premise. But then it looks like Austin would need to admit that there can be situations in which it is true to say 'I know it's a goldfinch, but I don't know it isn't a stuffed goldfinch'. Kaplan agrees that it would be worrying if Austin were committed to such statements being true.

Kaplan argues that the most promising strategy Austin has available to him for avoiding the sceptical conclusion is to make plausible the idea that there can be circumstances in which it isn't true for me to say either 'I know it's not a robotic goldfinch' *or* 'I don't know it's not a robotic goldfinch'. On this point Kaplan and I are in agreement. Kaplan, however, seems to think that the best way of making sense of this idea is to show that there can be circumstances in which it will be *false* to say either 'I know it's not a robotic goldfinch' or 'I don't know it's not a robotic goldfinch'. This strikes me as a counterintuitive idea, and one that we needn't commit ourselves to if we adopt my own Austin-inspired account.

Kaplan begins with Austin's suggestion that to claim to know that  $p$  is to claim to have done enough for present intents and purposes to establish that  $p$ . To say that I know that what I've seen isn't a robotic imitation is to say that I have done enough to establish that it is not a robotic imitation, and to say that I don't know that it isn't a robotic imitation is to say that I have not done enough (perhaps anything) to establish that it isn't. Kaplan thinks that by *supplementing* this idea of Austin's he can make plausible the idea that there can be situations in which it is false to say either that I know that it isn't a robotic imitation *or* that I don't know this.



Kaplan supplements the idea as follows. He begins by considering what our interest might be in particular situations in judging that we do or don't know something to be so, and makes the following suggestion:

The answer, I think, is that we want to distinguish questions that (at least, for the purposes we have at hand) we can consider closed from those we cannot – to distinguish propositions whose truth (at least, for the purposes we have at hand) we can act from those on which we cannot. When we claim to know that  $p$ , we are claiming a certain entitlement, for present purposes, to act on  $p$ ; when we claim not to know that  $p$ , we are claiming that we have no such entitlement for present purposes – that we cannot act on  $p$  for present purposes. (Kaplan, 2011, p. 64)

Kaplan admits that the idea of being 'entitled' to act on the understanding that  $p$  is one that requires some working out. That said, the suggestion is plausible enough. If I take myself to know that there is a goldfinch in the garden, it does seem that in doing so I take myself to be entitled to act on the understanding that there is a goldfinch in the garden for present purposes. Kaplan likewise supposes that in taking myself *not* to know that there is a goldfinch in the garden, I am taking myself not to be entitled to act on the understanding that there is a goldfinch in the garden for present purposes. There is room to interrogate this point but doing so would take the discussion unnecessarily far afield. For now I will take Kaplan's suggestion for granted.

Having made this suggestion Kaplan takes himself to have the tools required for making sense of the idea that there could be circumstances in which would not be true to say either that I know that what I see isn't a robotic imitation, or that I don't know this. He suggests that the concept of knowledge is used for placing propositions in one of two categories. In the one category are propositions I can be said to know, propositions I have done enough for present purposes to establish as being true and whose truth I am thus entitled to take for granted for the purposes of acting. In the other category are propositions I do not know, propositions I have not done enough for present purposes to establish as being true and whose truth I am thus not entitled to take for granted for the purposes of acting.

Kaplan argues that this schematisation leaves room for a third category in any situation. In the third category will be propositions whose truth I have not established but am nonetheless entitled to take for granted for the purposes of acting. Going off the schematisation that Kaplan has provided, I won't count either as knowing or as not knowing propositions in this category. I won't count as knowing such propositions to be true because I won't have done anything to establish their truth. That said, I won't count as not knowing such propositions to be true either. This is because to count a proposition as one I do not know to be true is to count it as one whose truth

I cannot take for granted for the purposes of acting. If I am entitled to take it for granted that a proposition is true for the purposes of acting, I won't count as not knowing the truth of this proposition (on Kaplan's account).

What kind of propositions might fit into this third category on Kaplan's account? Kaplan suggests that the kinds of propositions typically adduced in creating sceptical arguments from ignorance plausibly belong in this third category. As Austin notes, it is only in certain circumstances that I will be taken to be in a position to show that what I see is not a robotic goldfinch in claiming that I know that what I see is a goldfinch. Where we have no special reason to suppose that I might in fact be presented with a robotic goldfinch in this case, we won't generally take the suggestion that it might be a robotic goldfinch to count against my claim to know. This, according to Kaplan, suggests that we take me to be entitled in such circumstances to take it for granted that I am not presented with a robotic imitation, even though I have done nothing to establish that this is the case. So in such circumstances I won't (on Kaplan's view) count either as knowing that it isn't a robotic imitation *or* as not knowing this. Neither fits my standing with respect to the suggestion that what I see might be a robotic goldfinch.

Suppose that 'I am not being presented with a robotic goldfinch' belongs in this category as something that I neither count as knowing nor as not knowing. In that case, says Kaplan, we are in a position to diagnose what goes wrong in the argument from ignorance. The argument was as follows:

- P1. I don't know that what I see isn't a robotic goldfinch.
- P2. If I don't know that what I see isn't a robotic goldfinch, I don't know that it is a goldfinch.
- C. Therefore, I don't know that what I see is a goldfinch.

Kaplan argues that such an argument seems compelling because the expression 'I don't know that what I see isn't a robotic goldfinch' is crucially ambiguous. Taken one way, 'I don't know that what I see isn't a robotic goldfinch' says that the proposition 'What I see isn't a robotic goldfinch' belongs in the category of propositions whose truth I have not done enough to establish and am not entitled to take for granted for the purposes of acting. Taken in this way the second premise will be true. If I am not entitled to take it for granted that what I see isn't a robotic

goldfinch, then I am not entitled to take it for granted that what I see is a goldfinch.<sup>65</sup> In that case ‘I am seeing a goldfinch’ will likewise belong in the category of things I am not entitled to take for granted, and thus count as not knowing. However, Kaplan argues that on this understanding of ‘I don’t know that what I see isn’t a robotic goldfinch’, the first premise is one we have reason to reject, at least ordinarily. Insofar as we would ordinarily think that I am entitled to take it for granted that what I see isn’t a robotic goldfinch in such situations, we needn’t hold that I *don’t know* that what I see isn’t a robotic goldfinch. Kaplan takes this to suggest that, understood as saying that I am not entitled to take it for granted that what I see isn’t a robotic goldfinch, the first premise is false.

Kaplan goes on to suggest that the appeal of the argument comes from the fact that, read in another way, the first premise of the argument looks to be true. If all the first premise means is that the proposition ‘What I see isn’t a robotic imitation’ does not belong in the category of things I know, then the first premise will be true. But Kaplan thinks that on this understanding of what ‘I don’t know that what I see isn’t a robotic imitation’ comes to, it no longer conflicts with the claim that I *know* that what I see is a goldfinch. It only conflicts with this claim on the assumption that ‘not knowing that what I see isn’t a robotic imitation’ amounts to not being entitled to take it for granted that what I see isn’t a robotic imitation. In that case the second premise looks to be false and the conclusion is again avoided. Or so goes Kaplan’s Austin-inspired response to the argument from ignorance.

Kaplan’s Austin-inspired account furnishes us with a vantage point from which to appreciate why certain instances of the argument from ignorance are liable to strike us as puzzling while at the same time seeing that such arguments needn’t compel us in the direction of a sceptical conception of knowledge. While Kaplan’s strategy resembles in many ways the Austin-inspired account I have sketched, there is at least one substantial difference between our approaches. Whereas Kaplan is happy to grant that we can in every situation make sense of what knowing and not knowing an alternative will come to, my own approach has been to question whether this is so. Consider the case involving Ollie and the goldfinch once again. In the situation where Ollie has seen the bird, Kaplan is happy to grant that it will be false to say of Ollie that he *knows* it isn’t a robotic imitation. However, Kaplan doesn’t think that this commits him to its being *true* to say that Ollie doesn’t know it isn’t a robotic imitation. Rather, Kaplan holds that it may also in such circumstances be *false* to say that Ollie doesn’t know it isn’t a robotic imitation. This will be the

---

<sup>65</sup> While Kaplan’s (2011, p. 72 fn. 17) position is that “knowledge is not closed under known entailment,” he holds that “the entitlement to act on a proposition is so closed.”

case if there is no reason for Ollie to seriously suppose he is being presented with a robotic imitation, and so no grounds for thinking that Ollie is not entitled to take it for granted that it is not a robotic imitation.

My own view is that this concedes more than is needed in the face of such arguments from ignorance. Like Kaplan, I begin with Austin's suggestion that in saying that person knows that *p*, we are saying that he has done enough for present purposes to establish that *p*. If the circumstances are such that we have no reason to seriously suppose that an alternative might in fact obtain, and so do not see the alternative's not obtaining as something in need of establishing for present purposes, it isn't at all clear what sense to make of the suggestion either that a person has or has failed to do 'enough' for present purposes to establish that the alternative doesn't obtain. In such circumstances it thus won't make any sense to say either that the person 'knows' or 'doesn't know' that the alternative doesn't obtain. This is not the same as saying that it will be *false* to say either of these things (Kaplan's suggestion).

Kaplan looks to supplement Austin's core suggestion with his own observations regarding the link between knowing that *p* and being entitled to act on the understanding that *p*. While I think the link Kaplan draws attention to here is very real and worth exploring, I don't think the idea should or needs be employed in the way Kaplan employs it. Kaplan employs the idea so as to make the case that it is *false* to say that a person doesn't know that *p* in circumstances where we would take the person to nonetheless be entitled to take it for granted that *p*. My own view is that Kaplan is too quick to grant that the suggestion that a person 'doesn't know' some far-fetched alternative not to obtain always makes determinate sense. And this goes hand in hand with his being prepared to grant that the suggestion that a person 'knows' some far-fetched alternative not to obtain will always make determinate sense. It seems to me that it is in both cases contestable whether or not a speaker would be making any determinate sense in saying that the person 'knows' or 'doesn't know' these far-fetched alternatives not to obtain. And my view is that Austin's work furnishes us with resources for at least beginning to make out why it might not make determinate sense to say such things in certain circumstances.

While mine and Kaplan's respective Austin-inspired responses to the argument from ignorance resemble each other in many ways, I have highlighted at least one substantial point over which our accounts diverge. Both of us think that Austin furnishes us with resources for showing that it might not be true to say either that a person knows or doesn't know in a situation, and that this is the key to showing that arguments from ignorance needn't compel us in the direction of a sceptical conception of knowledge. That said, we differ in our respective ways of making sense

of this idea. I have suggested that Austin furnishes us with tools for making sense of the idea that it might not be either true or false to say that a person does not know something, and this because saying that they ‘do not know’ in the envisaged circumstances wouldn’t make any determinate sense. In contrast, Kaplan grants that it will in the envisaged circumstances make sense to say that a person doesn’t know that a certain alternative doesn’t obtain, but holds that we can view what is said in such circumstances as false without being committed to the view that the person must in that case *know* that the alternative doesn’t obtain. I find this to be a strange outcome. My own Austin-inspired account does not require us to accept any such suggestion.

### **3.6 Comparison with David Lewis’s contextualist response to the sceptical puzzle**

I now want to contrast my Austin-inspired strategy for resisting arguments from ignorance with the contextualist strategy presented by David Lewis.<sup>66</sup> Lewis shares the idea that in order to count as knowing something, a person needn’t be able to rule out any conceivable alternative to what he is taken to know. Rather, in order to count as knowing something a person need only be in a position to rule out any alternatives which are relevant on that occasion for ascribing knowledge. In addition to this, Lewis is committed to the contextualist view that what counts as a relevant alternative on one occasion for considering whether a person knows needn’t count as a relevant alternative on another occasion for considering whether this same person knows this same thing. For example, Ollie’s having seen a robotic imitation instead of a goldfinch may count as a relevant alternative on one occasion for considering whether Ollie knows that he has seen a goldfinch, while not counting as a relevant alternative on another occasion for considering whether Ollie knows this. In spite of the similarities between my Austin-inspired contextualism and Lewis’s contextualism, there are important differences. One is that Lewis seems to think it sufficient for an alternative to be relevant on an occasion if someone merely draws attention to that alternative. This means that Ollie might count as knowing that what he sees is a goldfinch at one moment, but then *cease* to know this simply as a result of our having turned our attention to the possibility that what he sees is a robotic imitation. I argue that this gets our epistemic practices wrong and leads to some highly counterintuitive commitments, commitments not entailed by my Austin-inspired contextualism. Another point of difference is that Lewis remains committed to the closure principle. In other words he thinks that Ollie *must* know that what he sees isn’t a robotic

---

<sup>66</sup> Aspects of Lewis’s contextualist approach are echoed in the respective approaches of Keith DeRose (1995) and Stewart Cohen (1999).

imitation if he is to count as knowing that what he sees is a goldfinch. This leads Lewis to propose the following response to the sceptical paradox. When we are not attending to the possibility that what Ollie sees is a robotic imitation, Ollie might count as knowing that what he sees isn't a robotic imitation. But as soon as we attend to this possibility, Ollie no longer counts as knowing that what he sees isn't a robotic imitation. I argue that we do better to follow the Austin-inspired strategy I have presented and reject the closure principle.

### **3.6.1 Lewis's contextualism and the Rule of Attention**

Lewis puts forward the following contextualist thesis regarding knowledge:

*S* knows that *P* iff *S*'s evidence eliminates every possibility in which not-*P* – Psst! – except for those possibilities that we are properly ignoring. (1996, p. 554)

What makes this a contextualist thesis is that Lewis thinks the range of possibilities we can properly ignore varies depending on the occasion for ascribing the knowledge in question. If we understand this as meaning that a possibility which counts as a relevant alternative (in need of ruling out) on one occasion for ascribing knowledge may not count as a relevant alternative on another, and that a person need only be in a position to rule out the relevant alternatives if he is to count as knowing, then Lewis's account shares this much in common with the Austin-inspired account I have presented.<sup>67</sup> In spite of this similarity, Lewis's response to the sceptical puzzle occasioned by arguments from ignorance differs substantially from the response I have sketched. One way of unpacking the differences in our respective responses is to focus on Lewis's Rule of Attention.

Lewis's (1996, p. 559) Rule of Attention states that "a possibility not ignored at all is *ipso facto* not properly ignored." This rule is supposed to explain both the context-sensitivity of ascriptions of knowledge, and also the appeal of arguments from ignorance. Lewis suggests that no matter how properly we *might* have ignored an alternative in a particular context, we are no longer properly ignoring the alternative once our attention has been directed towards it. If it is an alternative to *p* that the candidate knower is in no position to rule out, then we will be forced to conclude that this person does not know that *p*. This means that on Lewis's picture, the truth-value of an ascription of knowledge can be altered simply by drawing attention to a possibility that was previously being ignored. For example, we might in a particular situation be properly ignoring the

---

<sup>67</sup> There are some substantial differences between mine and Lewis's views on the matter, but these are mostly tangential to the points I want to consider in this section.

possibility that what Ollie sees is a robotic imitation (instead of a goldfinch). To say that we are ‘properly ignoring’ this possibility is, on Lewis’s account, to say that it is not a relevant alternative for the purposes of judging whether Ollie knows that what he sees is a goldfinch. But if someone were to draw our attention to this alternative, it would now be a relevant alternative for the purposes of considering whether or not Ollie knows.

Lewis thinks that the Rule of Attention can explain the appeal of sceptical arguments like the argument from ignorance. Arguments from ignorance proceed by drawing our attention to alternatives we were previously not attending to. In drawing attention to these possibilities, they become relevant alternatives for the purposes of judging whether the person knows the thing in question. Provided they are alternatives which the candidate knower is not in a position to show not to obtain, we will be forced to admit that the person does not know the thing in question. In other words sceptical arguments like the argument from ignorance are appealing precisely because they are *sound* in any context in which they are rehearsed. There are occasions on which we would ordinarily think it true to say that Ollie knows that what he sees is a goldfinch, even though he isn’t able to show that what he sees isn’t a robotic imitation. Lewis’s explanation for this is that we are on such occasions *ignoring* (turning a blind eye to) the possibility that what Ollie sees is a robotic imitation. But as soon as we attend to this possibility we see that Ollie is in no position to rule it out, and is thus in no position to know that what he sees is a goldfinch. While it might have been true to say that Ollie knows that what he sees is a goldfinch when we were properly ignoring the possibility that it is a robotic imitation, it will now be false to say this of Ollie. This, at least, is Lewis’s take on the matter.

My own view is that Lewis’s explanation is implausible as an account of the appeal of sceptical arguments, and that it leads to some highly counterintuitive claims about the concept of knowledge. First, our everyday epistemic practice does not suggest that it is sufficient for making an alternative a relevant alternative if we simply ‘attend’ to it. As Austin’s observations regarding everyday epistemic and discursive practice suggest, there is a difference between admitting that an alternative is conceivable and having reason to seriously treat it as something that might be so. Suppose Ollie and Dan are bird-watching and Ollie says that there is a goldfinch in a nearby tree. Dan asks him how he knows, to which Ollie responds by pointing out some of the features of the bird that have allowed him to identify it as a goldfinch in these circumstances. Ordinarily this might be enough to establish Ollie’s claim to know that what he sees is a goldfinch. But suppose Dan were now to say that Ollie isn’t in a position to know that what he sees is a goldfinch, and this because it might (for all he can tell) be a robotic imitation. Lewis suggests that in turning

Ollie's attention to this possibility, Dan creates a context in which it would now be false if Ollie were to say he knows that what he sees is a goldfinch. But it's not clear that we would treat Dan's 'suggestion' as having this implication ordinarily. Rather, we would ordinarily expect Dan to give us a reason for thinking that this particular bird might in fact be a robotic imitation. If Dan was unable to provide any such reason, and simply insisted that it was nonetheless possible, we wouldn't take Dan to have shown that Ollie is in no position to say he knows that the bird is a goldfinch. And this is plausible because we don't see the point in someone with Ollie and Dan's present purposes seriously treating robotic imitations as something to guard against.

In order for an alternative to be considered a relevant alternative, it needs to make sense to us why someone with our intents and purposes should treat the alternative as such, why it should be seriously treated as a way things might be. For that reason I don't think Lewis's Rule of Attention, at least as he presents it, accurately captures our everyday epistemic and discursive practices. It might be that if we adjust our understanding of what it means to 'attend' to a possibility the rule would be more plausible. If by 'attend' we mean seriously treating an alternative as a way things might be for present purposes, this might bring us closer to our actual epistemic and discursive practices. Dan might succeed in making out to Ollie that they have a reason (for their bird-watching purposes) to seriously treat robotic imitations as something to be on the look-out for. But in that case we wouldn't usually think that Dan had created a context in which it was only now false for Ollie to say he knows that what he sees is a goldfinch. Rather, it would be more natural to say that Ollie had been made aware of a relevant alternative he was previously unaware of. In that case Ollie would have reason to think that what he said previously was false.

It is a common criticism of contextualist accounts of knowledge that they imply that conversations like the following make perfect sense:

Ollie: I know that's a goldfinch.

Dan: But you don't know it isn't a robotic imitation, so you don't know it's a goldfinch.

Ollie: When I said 'I know that's a goldfinch' I was saying something true, but now that you've mentioned the possibility that it is a robotic imitation it is no longer true for me to say this.

Lewis's Rule of Attention seems to commit him to the view that, even if it would be inappropriate (for pragmatic reasons) for Ollie to speak in this way, he might nonetheless be saying something true here. And this is because Lewis allows that merely drawing attention to a previously ignored



alternative is enough to shift the conversational context so that Ollie goes from saying something true to saying something false in saying ‘I know that’s a goldfinch’. If, however, we reject Lewis’s Rule of Attention we needn’t be committed to the idea that a person in Ollie’s position might be saying something true here. Dan’s suggestion is directed as a *challenge* to Ollie’s claim to know in this situation. If it is successful as a challenge in these circumstances, if Dan has adduced an alternative that a reasonable person in his and Ollie’s circumstances should treat as a relevant alternative, then what Dan has shown is that Ollie said something false in saying that he knows, and Ollie should recognise this.

All in all I think we should reject Lewis’s Rule of Attention and the response to the argument from ignorance he develops off the back of it. One needn’t endorse the Rule of Attention in order to hold that which alternatives are to be treated as relevant alternatives is something that can vary depending on the present intents and purposes of the people judging the matter. As long as we still retain this contextualist insight, we still have the means of making sense of why arguments from ignorance should so much as seem compelling. The reason is that we can imagine occasions on which we would think that the alternatives to  $p$  adduced in the argument were relevant to the question of whether some person (considered on that occasion) knows that  $p$ . On such occasions the reasoning employed in the argument could well be sound, supposing that the person was in no position to show the alternative not to obtain.

### **3.6.2 Closure and Lewis’s response to the sceptical puzzle**

Lewis’s endorsement of the Rule of Attention is not the only point at which his account differs substantially from the Austin-inspired account I have presented. Whereas the account I have sketched involves giving up a closure principle for knowledge, Lewis (like a number of other contextualists<sup>68</sup>) thinks it an advantage of his account that it allows us to retain a closure principle for knowledge. (Lewis, 1996, p. 564) The closure principle, recall, states that if S knows that  $p$ , and S knows that  $p$  implies  $q$ , then S must also know that  $q$ . This means that Lewis is committed to the idea that if a person is to know that what he sees is a goldfinch, he must know that what he sees isn’t a robotic imitation. But it was just such a commitment that seemed to compel us in the direction of a sceptical conception of knowledge. In many ordinary situations where we might think it right to say that a person knows that what he sees is a goldfinch, it seems that it would be wrong to say that the person knows that what he sees isn’t a robotic imitation. Lewis proposes

---

<sup>68</sup> DeRose and Cohen being the two most prominent examples.

an account according to which a person can know that what he sees is not a robotic imitation, even if he has done nothing to establish that this is so. The appeal of the sceptical argument, according to Lewis, comes from the fact that any time we turn our attention to this knowledge we create a context in which the person no longer counts as knowing. My own view is that this response to the sceptical puzzle again has counterintuitive implications, implications that do not follow from the Austin-inspired account I have sketched.

Lewis's response to the sceptical puzzle involves rejecting the same premise I have found to be problematic in arguments from ignorance, namely the first premise. But whereas I have suggested that there can be circumstances in which the first premise is neither true nor false, Lewis maintains that the first premise must be false if the conclusion is also to be avoided. Consider again the following argument from ignorance:

- P1. Ollie doesn't know that what he sees isn't a robotic imitation.
- P2. If Ollie doesn't know that what he sees isn't a robotic imitation, then he doesn't know that what he sees is a goldfinch.
- C. Ollie doesn't know that what he sees is a goldfinch.

Like me, Lewis does not think it a plausible strategy to try rejecting the second premise in this argument. But unlike me he thinks that the only way for us to avoid the conclusion of such an argument is showing that the first premise is false. In other words he thinks that Ollie can only count as knowing that what he sees is a goldfinch if he knows that what he sees isn't a robotic imitation. Given that Lewis wants to preserve the idea that we can speak truth in saying that Ollie knows that what he sees is a goldfinch even in situations where Ollie is not in a position to show that what he sees isn't a robotic imitation, Lewis is thus committed to the idea that Ollie can know that what he sees isn't a robotic imitation even though he has done nothing to establish that it isn't one.

Suppose there is an occasion on which we judge that Ollie knows that what he sees is a goldfinch and on which his having seen a robotic imitation is not a relevant alternative. In judging that Ollie knows that what he sees is a goldfinch, Lewis thinks we are implicitly committed to Ollie's knowing that what he sees isn't a robotic imitation. However, Lewis is sensitive to the fact that it is liable to strike us as wrong to say Ollie that *knows* that what he sees isn't a robotic imitation in a situation where he has done nothing to establish that this is so. This leads Lewis to make the following concession. While Ollie might count as knowing that what he sees isn't a robotic

imitation even if he has done nothing to establish that this is so, any attempt to say (or think) that Ollie knows this will render what we say (or think) false. This is because in attending to the possibility that Ollie might instead be looking at a robotic imitation, the context changes in such a way that Ollie needs to have done something specifically to rule out its being a robotic imitation if he is to count as knowing that it isn't one.

One strange consequence of Lewis's strategy is that it suggests there is a whole galaxy of things we count as knowing, but that we can never speak (or think) truth in saying (or thinking) that we know. This is partly a result of Lewis's commitment to the view that in merely attending to a possibility we shift the context so that we are no longer entitled to take it for granted that the possibility does not obtain. Just as we needn't share Lewis's commitment to the Rule of Attention, we needn't share Lewis's commitment on this front. We might instead take a Moorean line and insist that as long as we haven't been given any reason to seriously suppose that a particular doubt might obtain, we can dig our heels in and say we know that the doubt does not obtain. But many of us feel that there is something wrong with saying, for example, that Ollie *knows* that what he sees isn't a robotic imitation, in a situation where Ollie has done nothing at all to establish that this is so. For those of us who feel this way, the Moorean response is bound to sit uncomfortably.

An advantage of the Austin-inspired account I have sketched is that it provides us with a way of rejecting the first premise in the argument from ignorance without needing to insist, for example, that Ollie knows that what he sees isn't a robotic imitation in a situation where he has done nothing to establish that this is so. On the account I have outlined, it mightn't be true to say in particular circumstances either that Ollie knows or that he doesn't know this. These will be circumstances in which we have no reason to seriously treat the possibility that Ollie has seen a robotic imitation as something which needs to be shown not to obtain for present purposes. In that case it doesn't make sense to say either that Ollie has or hasn't done 'enough' (for present purposes) to *show* that the possibility doesn't obtain. We can thus preserve the Moorean insight that merely drawing attention to a conceivable alternative does not in every case amount to a legitimate challenge to a person's claim to know something. And we can do so without needing to take the further Moorean step of insisting that we 'know' such alternatives to not obtain. Nor do we need to hold (with Lewis) that we can know these alternatives not to obtain so long as we are not attending to them.

### **3.7 Summary**

In this chapter I have outlined an Austin-inspired response to a sceptical puzzle occasioned by arguments from ignorance. Key to the Austin-inspired strategy I have outlined is the idea that there can be circumstances in which it is neither true nor false to say that a person ‘doesn’t know’ something. To make sense of this idea I have focused on Austin’s suggestion that in saying that a person knows something  $p$ , we are saying that he has done ‘enough’ for present purposes to establish that  $p$ . I then suggested that in circumstances in which we don’t think that  $p$  is something in need of establishing, there mightn’t be any determinate sense to the suggestion that a person has done or has failed to do ‘enough’ for present purposes to establish that  $p$ . In that case it is not clear what determinate sense we are to make in such circumstances of the suggestion that a person either ‘knows’ or ‘doesn’t know’ that  $p$ . If there are circumstances in which there is no determinate sense to the suggestion that a person ‘knows’ or ‘doesn’t know’ that  $p$ , the suggestion that the person ‘doesn’t know’ that  $p$  won’t count against the person’s knowledge on some other matter. Such at least is the Austin-inspired account I have set out. Clearly this is quite a radical suggestion, and more work is needed spelling out the idea. Nonetheless I think it is a suggestion that is worth investigating in further detail.

## Concluding remarks

My aim throughout this thesis has been to illustrate the continuing relevance and importance of Austin's approach and observations on the topic of empirical knowledge. In order to gain a clearer understanding of what knowing comes to, Austin encourages us to attend to the particular circumstances in which someone might judge a person to know something, and to consider what we might reasonably understand the speaker to be judging in such circumstances. To articulate our views on such matters is to articulate our understanding of what it takes to count as knowing in such circumstances. When we do this we see that a speaker cannot generally (if ever) be taken to be judging that the person is in a position to 'rule out' any conceivable alternative to that which he is taken to know. In any particular situation we have a (more or less rough) idea of what is 'enough' to count as knowing in those circumstances. Unless we have some reason to seriously suppose that some alternative to  $p$  might in fact obtain, we won't require a person to be in a position to 'rule out' that alternative in order to count him as knowing that  $p$ .

I have argued that when we follow Austin's lead and focus on how ascriptions of knowledge are reasonably to be understood in practice, the considerations adduced speak in favour of a contextualist account of knowledge. Depending on who is ascribing the knowledge and in what particular circumstances, our understanding of what it would take for things to be as the speaker has said can vary. In particular, an alternative that is not relevant in one set of circumstances might nonetheless be relevant in a different set of circumstances. In the latter case, though not necessarily in the former, the person being ascribed knowledge will need to be in a position to rule out that alternative if he is to count as knowing the thing in question. I have argued that one implication of this is that a person might count as knowing that  $p$  on one understanding of what it is to know that  $p$ , and not count as knowing on another understanding of what it is to know that  $p$ . One speaker could speak truth in saying that the person knows that  $p$ , and another person (in different circumstances) could say something false in saying that this same person knows that  $p$ . I contrasted a contextualist account of knowledge with an invariantist account, and argued that a contextualist account allows us to make better sense of the relevant aspects of our epistemic and discursive practices.

I have also argued that Austin's observations regarding knowledge furnish us with the beginnings of at least one novel response to a sceptical puzzle occasioned by 'arguments from ignorance'. More work is needed filling out the Austin-inspired response I have presented, but it has the

advantage of avoiding certain counterintuitive implications of other responses to the puzzle. I suspect that considerations adduced in the second chapter could be used to make further sense of the idea that there might be circumstances in which saying that a person does or doesn't 'know' carries no determinate sense. If, as I have argued in Chapter 2, ascriptions of knowledge admit of various (perhaps indefinitely many) understandings, it is possible that the circumstances in which a speaker says that a person 'knows' might not enable us to determine how he is to be understood.<sup>69</sup> This might prove to be a fruitful way of further spelling out the suggestion I have put forward, namely that there might be circumstances in which it is indeterminate what would count as a person's having done or having failed to do 'enough' for present purposes to establish that something is so.

There are many points I have touched on here that suggest avenues for further investigation. One question I am particularly interested in pursuing relates to the notion of a 'relevant alternative': What factors lead in practice to an alternative that might in other circumstances have been considered irrelevant being taken up as a relevant alternative? Such a question could be read in such a way as to invite psychological and sociological investigation, but there is also room for philosophical reflection on the matter as well.

Another avenue for further investigation would be to connect the contextualist considerations I have discussed with broadly-speaking 'functionalist' considerations regarding the concept of knowledge. By 'functionalist' considerations I mean considerations pertaining to the normative implications in practice of judging that a person does or doesn't know something. Put otherwise I am interested in further exploring the practical upshot (the point) in particular circumstances of judging that a person does or doesn't know something. I am particularly interested in further investigating the connection between the concept of knowledge and our testimonial practices (briefly alluded to in Chapter 2), though there is also no doubt interesting work to be done on the connection between ascriptions of knowledge and (moral) responsibility.<sup>70</sup> Austin's remarks on testimony are worth studying in their own right, if just for the fact that Austin is more sensitive than some to the important differences between knowledge gained through testimony and knowledge gained by other means. Austin's remarks about the relationship between claiming to

---

<sup>69</sup> Charles Travis is one Austin-inspired philosopher who stresses the importance of this idea, arguing that it has implications for the kinds of questions about knowledge we can fruitfully ask in the context of philosophical reflection. See Travis (1991, pp. 242-243).

<sup>70</sup> Was the theme park owner in a position to *know* that the ride was faulty? *Should* he have known? If so, what kind of position are we saying he *ought* to have been in? What *ought* he to have done?

know, testimony, and responsibility to others suggests one place in which to begin getting a grip on the complex second-personal relations testimonial exchange typically involve.

In any case Austin's work and methodology remain a relevant and important resource for reflecting on the topic of empirical knowledge. Remarking on the study of "what we say when", Austin (1961, p. 129) once declared "there is gold in them thar hills". The view animating this thesis, one which I continue to hold, is that much the same can be said of Austin's writings.

# Bibliography

- Austin, J. L. (1961). *Philosophical papers*. Oxford: Clarendon Press.
- Austin, J. L. (1962). *Sense and sensibilia*. Oxford: Clarendon Press.
- Cavell, S. (1999). *The Claim of Reason: Wittgenstein, skepticism, morality, and tragedy*. New York: Oxford University Press.
- Chisholm, R. (1989). *Theory of Knowledge* (3 ed.). Englewood Cliffs, N.J.: Prentice Hall.
- Chisholm, R. M. (1964). J. L. Austin's Philosophical Papers. *Mind*, 73(289), 1-26.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Nous*, 33(s13), 57-89.
- Craig, E. (1990). *Knowledge and the state of nature: an essay in conceptual synthesis*. Oxford;New York;: Clarendon Press.
- Crary, A. (2002). The Happy Truth: J.L. Austin's *How to Do Things With Words*. *Inquiry*, 45(1), 59-80.
- Crary, A. (2007). Objectivity Revisited: A Lesson from the Work of J.L. Austin. In *Beyond Moral Judgement* (pp. 49-95). Cambridge, Massachusetts: Harvard University Press.
- Davidson, D. (1996). The Folly of Trying to Define Truth. *The Journal of Philosophy*, 93(6), 263-278.
- DeRose, K. (1992). Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52(4), 913-929.
- DeRose, K. (1995). Solving the Skeptical Problem. *The Philosophical Review*, 104(1), 1-52.
- DeRose, K. (2009). *The case for contextualism: knowledge, skepticism, and context*. Oxford: Oxford University Press.
- Descartes, R. (1993). *Meditations on First Philosophy: in focus* (S. Tweyman Ed.). London: Routledge.
- Dretske, F. (1970). Epistemic Operators. *The Journal of Philosophy*, 67(24), 1007-1023.
- Faulkner, P. (2007). On Telling and Trusting. *Mind*, 116(464), 875-902.
- Gettier, E. (1966). Is justified true belief knowledge? *Analysis*, 23(6), 121-123.
- Grice, P. (1989). *Studies in the Way of Words*. Cambridge, Massachusetts: Harvard University Press.
- Hawthorne, J. (2004). Sensitive Moderate Invariantism. In *Knowledge and Lotteries* (pp. 157-191). Oxford: Clarendon.
- Hetherington, S. (1996). *Knowledge Puzzles: an introduction to epistemology*. Boulder, Colo.: Westview Press.
- Hinchman, E. (2005). Telling as Inviting to Trust. *Philosophy and Phenomenological Research*, 70(3), 562-587.
- Kaplan, M. (2000). To what must an epistemology be true? *Philosophy and Phenomenological Research*, 61(2), 279-304.
- Kaplan, M. (2006). If You know, You Can't Be Wrong. In S. Hetherington (Ed.), *Epistemology Futures* (pp. 180-198). Oxford: Clarendon Press.
- Kaplan, M. (2008). Austin's Way with Skepticism. In J. Greco (Ed.), *The Oxford Handbook of Skepticism* (pp. 348-366). Oxford: Oxford University Press.
- Kaplan, M. (2011). Tales From The Unknown: Austin and the Argument from Ignorance. In M. Gustafsson & R. Sørli (Eds.), *The Philosophy of J.L. Austin*. Oxford: Oxford University Press.
- Lawlor, K. (2013). *Assurance: An Austinian View of Knowledge and Knowledge Claims*. Oxford: Oxford University Press.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549-567.  
doi:10.1080/00048409612347521
- McDowell, J. (1995). Knowledge and the Internal. *Philosophy and Phenomenological Research*, 55(4), 877-893.
- McMyler, B. (2011). Believing What the Man Says About His Own Feelings. In M. Gustafsson & R. Sørli (Eds.), *The Philosophy of J.L. Austin*. Oxford: Oxford University Press.



- Moore, G. E. (1993). *Selected Writings*. London: Routledge.
- Moran, R. (2006). Getting Told and Being Believed. *Philosopher's Imprint*, 5(5), 1-29.
- Moran, R. (2013). Testimony, illocution and the second person. *Aristotelian Society Supplementary Volume*, 87(1), 115-135.
- Peirce, C. S. (1868). Some Consequences of Four Incapacities. *The Journal of Speculative Philosophy*, 2(3), 140-157.
- Plato. (1973). *Theaetetus* (J. McDowell Ed.). Oxford: Clarendon Press.
- Plato. (2006). *Plato's Meno* (D. Scott Ed.). Cambridge: Cambridge University Press.
- Putnam, H. (1999). *The Threefold Cord: Mind, body and world*. New York: Columbia University Press.
- Putnam, H. (2001). Skepticism, stroud and the contextuality of knowledge. *Philosophical Explorations*, 4(1), 2-16. doi:10.1080/13869790108523339
- Putnam, H. (2012). Skepticism and Occasion-Sensitive Semantics. In D. Macarthur & M. De Caro (Eds.), *Philosophy in an age of science: Physics, mathematics, and skepticism* (pp. 514-534). Cambridge, Massachusetts: Harvard University Press.
- Radford, C. (1966). Knowledge: By examples. *Analysis*, 27(1), 1-11.
- Shope, R. K. (1983). *The analysis of knowing: A decade of research*. Princeton, N.J.: Princeton University Press.
- Snowdon, P. F. (2005). The Formulation of Disjunctivism: A response to Fish. *Proceedings of the Aristotelian Society*, 105, 129-141.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Clarendon.
- Stroud, B. (1984). *The significance of philosophical scepticism*. Oxford: Clarendon.
- Travis, C. (1989). *The Uses of Sense: Wittgenstein's Philosophy of Language*. Oxford: Clarendon.
- Travis, C. (1991). Annals of analysis. *Mind*, 100(398), 237-264. doi:10.1093/mind/C.398.237
- Travis, C. (2001). Mind Dependence. *Revue Internationale de Philosophie*, 218(4), 503-524.
- Travis, C. (2005). A Sense of Occasion. *The Philosophical Quarterly (1950-)*, 55(219), 286-314. doi:10.1111/j.0031-8094.2005.00399.x
- Travis, C. (2008). Pragmatics. In *Occasion-sensitivity: selected essays* (pp. 109-129). New York: Oxford University Press.
- Travis, C., & Kalderon, M. (2013). Oxford Realism. In M. Beaney (Ed.), *The Oxford Handbook of the History of Analytic Philosophy* (pp. 489-515). Oxford: Oxford University Press.
- Welbourne, M. (1986). *The community of knowledge* (Vol. no. 9). Aberdeen;[Atlantic Highlands], N.J.: Aberdeen University Press.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *The Philosophical Quarterly*, 55(219), 213-235.
- Zagzebski, L. (1994). The Inescapability of Gettier Problems. *The Philosophical Quarterly*, 44(174), 65-73.