

[Penultimate version of paper forthcoming in *Erkenntnis*. Please cite the original version of the paper.]

Authoritative Knowledge

Abstract: *This paper investigates ‘authoritative knowledge’, a neglected species of practical knowledge gained on the basis of exercising practical authority. I argue that, like perceptual knowledge, authoritative knowledge is non-inferential. I then present a broadly reliabilist account of the process by which authority yields knowledge, and use this account to address certain objections.*

Keywords: authority, non-inferential knowledge, practical knowledge.

Authoritative Knowledge


Word Count: 11,827 (13,973 with footnotes)

August 4, 2020

You may recognize the following picture:



If you do not recognize it, you might have good guesses as to what it represents: a hat, a slug, or a falling tent, perhaps. Of course, all of these guesses are wrong. As we learn from the opening pages of *The Little Prince* (2001 [1943]), the picture represents a boa constrictor eating an elephant. We are so told by the author of the picture, who draws a second one to help us, unimaginative adults, appreciate the true nature of his drawing:



Suppose that you are one of the unimaginative adults in the story, and that the author shows you the first picture. You say, 'What a nice hat,' and he replies, 'It's not a hat; it's a boa eating an elephant'. The child's assertion is clearly true—but I want to claim something stronger: I want to claim that this assertion expresses knowledge, a knowledge of a special type that the author has

precisely in virtue of his authority as the author of the picture. I shall thus call it ‘authoritative knowledge’. This paper aims to offer an account of authoritative knowledge.

I shall approach the topic from within the Reidean tradition in epistemology. A central tenet of this tradition is that there is a set of privileged sources of information whose deliverances give us direct (non-inferential) knowledge of the world. The traditional list of sources includes perception, memory, testimony, reasoning, and reflection. Let’s call these ‘basic knowledge sources’. A basic knowledge source is not an infallible source of knowledge. For instance, we can employ our perceptual systems and fail to gain knowledge (say, because our vision is blurry). Rather, what is distinctive about basic knowledge sources is that we can gain knowledge directly on their basis: in the right circumstances, engaging such systems is sufficient to gain knowledge. I shall argue that authority is an unrecognized basic knowledge source.

If authority yields knowledge in some way analogous to the direct way in which sources like perception and memory give us knowledge, that is an important epistemological finding.¹ After all, the philosophical study of authority has predominantly fallen to ethicists and political philosophers, receiving little attention by epistemologists, unlike other basic knowledge sources.² Moreover, authoritative knowledge is a species of what Anscombe ([1957] 1963) famously called “practical knowledge”: unlike speculative knowledge, which “is derived from the objects known”, practical knowledge is “the cause of what it understands” (p.87; cf. pp.13, 57). Although there is now a large literature on practical knowledge, the scope of the discussion is in my view excessively

¹ I shall use locutions such as ‘*x* yields direct knowledge’ as shorthand for ‘we can gain knowledge directly on the basis of exercising *x*’.

² More recently, it has received significant attention in philosophy of language (see n5 below). It is important that I’m concerned with what is sometimes called ‘practical authority’, such as a captain has over her soldiers. This is to be distinguished from what is sometimes called ‘epistemic authority’, which is the power to speak and be trusted on a certain subject. A doctorate degree in biology, for instance, might give you epistemic authority, insofar as it entitles you to speak and be trusted about certain biological subjects, without conferring any practical authority, such as the authority to command others. In pioneering work, Zagzebski (2015) appeals to the literature on practical authority to shed light on the notion of epistemic authority.

narrow.³ It is restricted to knowledge of an agent's own actions, whereas I shall be arguing that we can have practical knowledge of others's actions, and even of non-agentive facts. In this way, the investigation advances our understanding of practical knowledge.

The paper will proceed as follows. I start with some further remarks meant to elucidate the nature of authoritative knowledge and its status as practical knowledge, and to explain the scope of the thesis that it is direct (Section 1). Next, I explain in more detail what's at stake in deciding between direct and inferential accounts (Section 2), and offer two arguments in favour of the former (Section 3). I then present a broadly reliabilist account to explain how authority yields direct knowledge (Section 4), and use this account to further clarify my view and respond to objections (Section 5).

1. Authoritative Knowledge Further Explained

The notion of authority that gives rise to authoritative knowledge is practical authority in general, of which an artist's authority is only a species. To help us get a better grip on this phenomenon, let me present two further examples of authoritative knowledge:

Captain: Captain O'Heguemon holds command over her ship's crew. Through a speakerphone in the lower cabin, she orders the sailors in the upper cabin to clean the floors of the main cabin before the day is over. Just on the basis of this order, and without performing any further epistemic checks, the captain comes to believe that the floors will be cleaned by the next morning. The crew is both competent and obedient, and they accomplish the task before the end of the day, as ordered. Here, the captain has authoritative knowledge that the floors of the cabin will be cleaned.

³ See Schwenkler (2015) for a helpful review of the literature on practical knowledge, along with a positive Anscombean account of the nature of the phenomenon.

Single Father: A mother has died at childbirth before the parents could settle on a name for the child. The father is now holding the baby, affected by unnameable emotions, finding himself both a widow and a father. A nurse passes by and asks him, ‘What’s the name of the child?’. The father, who hasn’t given a thought to this until now, considers for a moment and replies, ‘her name is Juliana, like her mother’. This assertion expresses knowledge, a knowledge that he has in virtue of exercising his baptismal authority. It is thus a case of authoritative knowledge.

We now have three examples of the target notion: the author’s knowledge about the picture from *The Little Prince*, the captain’s knowledge about the status of the floors in the morning, and the father’s knowledge about the name of the child. These cases are simple: they are cases where the primary bearer of authority is a single person, and where the exercise of authority is successful. Cases of shared authority and failures of execution raise complex problems of their own. I will mostly ignore these problems in this paper. My aim is to develop a theory that accounts for these simple cases, that can then be used to assess the more complex issues raised by shared authority and failure of execution. Finally, although there are obvious differences between the three examples, I aim to show that they should receive parallel *epistemic* treatment.

Throughout the paper, I shall rely on an intuitive understanding of authority: authority is the power to settle a matter through some conventionally recognized means (such as a decision, an order, or a baptismal act).⁴ I take this understanding to be compatible with the central accounts of authority in the literature, such as Joseph Raz’s account of authority as the source of exclusionary reasons (1985a, 1981, 1985b, 2010), and Stephen Darwall’s account of authority as

⁴ Note that this is not meant as a reductive definition. That would require an independent characterization of the class of conventional means that define authority, which I am skeptical can be done.

the source of second-personal reasons (2010, 2011). In determining whether someone possesses authority we can appeal to one of its characteristic marks: a person in authority is one who is in a position to (correctly) make certain speech acts in imperatival form. Thus, ‘Let there be light!’, ‘Clean your room!’, or ‘Let her name be Juliana’, are all paradigm expressions of authority. When expressed by someone who lacks authority, such illocutions misfire.⁵

I take it as intuitively clear that those in authority possess the relevant knowledge in the three examples I provided. That it is knowledge, rather than (say) merely true belief, can be seen from features of the speech acts expressing the mental states of these subjects. First, the parties can correctly express their mental states by making assertions. Focusing on the case of the captain, it is clear that he can correctly assert that the floors have been cleaned (e.g. if a passenger asked him). Now, it is generally agreed that assertion is governed by a knowledge norm (KNA) to the effect that one can correctly assert only what one knows.⁶ On that assumption, it follows that the captain knows that the floors have been cleaned. Moreover, such assertions can serve as the starting point of a testimonial chain that culminates in knowledge. If the captain tells a passenger that the floors have been cleaned, the passenger will acquire knowledge of the fact, suggesting that the initial assertion also manifests knowledge.

As noted, I also take authoritative knowledge to be practical knowledge in Anscombe’s sense. In holding this thesis, I take myself to be following Anscombe. One of her central cases to illustrate the notion of practical knowledge in *Intention* describes a man directing “the erection of a building which he cannot see and does not get reports on, purely by giving orders”. She takes

⁵ The classical discussion of this issue is in Austin (1975: 28-29). A complicating factor is that one can sometimes establish authority simply by acting as if one had it. See Thomason (1990), Witek (2013) and Langton (2015) for elucidative discussion of this issue.

⁶ Defenders of versions of KNA include Unger (1975), Slote (1979), Williamson (2000), DeRose (2002), Reynolds (2002), and Hawthorne (2004), among many others. Critics of the knowledge norm include Weiner (2005), Douven (2006), Lackey (2007), McKinnon (2013, 2015), Hill and Schechter (2007), and Schechter (2017).

such a man to be in a position to know “what the house is like” (Anscombe [1957] 1963, §45, p.82).⁷ The example plays an important role in the context of *Intention*. It is one of the two central cases in the monograph where someone’s knowledge is grounded in a *purely* practical basis, unaided by ‘speculative’ ones.⁸ This may come as a surprise, since much of the literature on practical knowledge treats it as restricted to the knowledge agents have of their own intentional actions. However, taken at face value, the example is of someone who knows facts such as that the house has three rooms and two bathrooms, that there is a chimney on the first floor, etc. These are facts that don’t concern actions, never mind the director’s own.

Why, then, should we regard this as a case of practical knowledge, of the same type as my knowledge that I am typing these words as I do? First and foremost, I take it, because this knowledge is the cause of what it understands: the builder’s knowledge of the house is the cause of the house, just as my knowledge of what I am doing is the cause of what I am doing. It is so as efficient cause, since the orders cause the workers to build the prescribed house.⁹ But it is also the formal cause, since the orders set the standards of correctness for the product: if the house doesn’t conform to the orders of the director, then there is an error in the house.¹⁰

On the other hand, the formal dependence is weaker in some cases of authoritative knowledge. For the existence of an intentional Φ ing as such is metaphysically dependent on the

⁷ My example of the captain is modelled after this case.

⁸ The other case is of someone writing something on a board without looking, discussed both earlier (p.53) and immediately after (p.82). The centrality of the builder case to the project of *Intention* has not been fully appreciated. It picks up the discussion about means from the previous section, and it is in turn picked up by much of the ensuing discussion, in particular at the end of §49, the section where Anscombe presents Aquinas’s definition of practical knowledge cited in my introduction (pp.88-89). The example is crucial to Anscombe’s further ethical aims of providing a proper moral psychology (on this topic, see Schwenkler’s helpful introduction (2019)). President Truman, after all, did not directly drop the atomic bombs, but he did have practical knowledge that they were dropped, having commanded this. However, as my aim is not Anscombe exegesis, I shall not pursue the point.

⁹ Here, I am disagreeing with Setiya (2016), who take practical knowledge to be merely a formal cause of intentional action. The example of the builder seems to me to make it clear that it is also the efficient cause of the thing known.

¹⁰ The knowledge is also the formal cause of the house by placing it in an instrumental order, making Anscombe’s special sense of the question ‘Why?’ applicable: e.g. ‘Why is the bathroom tiled thus?’ (Answer: ‘This tile is easy to clean’). *This* sense of the question is senseless for mere objects of nature, e.g. ‘Why is marble white?’.

agent's knowledge or practical cognition:¹¹ whether my hand's movements constitute an intentional action (say, a greeting) depends on whether I cognize what I do as a greeting. Some cases of authoritative knowledge share this feature: whether the lines on the page constitute a drawing of a boa eating an elephant depends on the author's regarding them as such. However, in other cases of authoritative knowledge the known object may exist independently of the knowledge: whether the walls of a room are white is independent of what the director thinks.¹² Although this is an important difference, I shall follow Anscombe in thinking that it doesn't cast doubt on the status of authoritative knowledge as practical.¹³

We have, then, grounds to distinguish between two species of practical knowledge: the standardly discussed knowledge that an agent has of her own actions ('agential knowledge'), and the knowledge that one has on the basis of exercising authority ('authoritative knowledge'), which may include knowledge of non-agential facts. Yet, even though authoritative knowledge is wider in scope, it is dependent on agential knowledge. Most obviously, this is because exercising authority is an action which (at least in the normal case) is intentional. Less obviously, this is because the existence of authority of various sorts is dependent on the existence of various practices constituted by intentional actions. We can see this from the fact that expressions of authority ('this is a boa constrictor', 'her name is Juliana', 'Sailors: clean the cabin!') make sense only in a context where people engage in activities like drawing, naming, speaking, obeying, etc.

¹¹ For our purposes, it doesn't much matter at this point whether the cognition at play always has the status of knowledge or a weaker state. I argue against the thesis that intentional action requires knowledge in Piñeros Glasscock (forthcoming-b), though I have become doubtful whether the thesis should be ascribed to Anscombe, as is universally assumed in the scholarship.

¹² Perhaps not completely independent in the given example, since the status of a thing as a *room* or a *wall* is at least somewhat dependent on agents taking it to be such. But that is incidental. Consider the case of someone directing the placement of certain plants in a garden, and knowing, say, where there the roses are on this basis.

¹³ It may call for a distinction along the lines Aristotle draws between craft knowledge (τέχνη) which concerns productions (ποιήσεις), and the knowledge associated with actions-proper (πράξεις) (*EE/NE* 6.5, 1140a24-b30). I discuss important differences between the two knowledges in Piñeros Glasscock (2019).

These are precisely the sorts of actions that Anscombe identifies as requiring practical knowledge to engage in them (§47, p.85).¹⁴

Recognizing the dependence of authoritative knowledge on agential knowledge does not detract from its importance, any more than recognizing the dependence of memory and testimony on perception detracts from the importance of recollective and testimonial knowledge. The epistemic importance of testimonial and recollective knowledge stems partly from their status as forms of direct knowledge. Now, Anscombe's claim that practical knowledge is non-observational is commonly taken to entail that it is non-inferential (i.e. direct).¹⁵ Much of the paper thus ends to defend the thesis that authoritative knowledge is non-inferential and present the outline of an account of how authority gives us such direct knowledge. If successful, the project would constitute a vindication of the claim that authoritative knowledge is an important epistemic category.

My ambitions for the paper, however, are limited. My aim shall be to show that *at least as good a case* can be made for thinking that authoritative knowledge is non-inferential as can be made for the claim that recollective and testimonial knowledge are non-inferential. This will allow me to set aside certain objections to the view that call into question the very possibility of direct knowledge generally. Such objections must be answered; but they must be answered by a general theory of the matter, which I cannot hope to provide.

¹⁴ The actions are denoted by what Jonathan Bennett (1988: 205-6) calls "intention-drenched verbs" and Michael Moore calls "intentionally complex" (2010: 174).

¹⁵ See O'Brien (2007) and Setiya (2008, 2009, 2007), and Piñeros Glasscock (forthcoming-a) for arguments in favour of non-inferentialism.

2. Direct and Indirect Models

The question whether some knowledge is direct or inferential is a question about the epistemic bases of the belief that constitutes such knowledge. Thus:¹⁶

Inferential Knowledge: Knowledge that p is inferential just in case it is warranted in virtue of an inferential relation this knowledge has to some other belief(s) of the subject.

Direct Knowledge: Knowledge that p is direct just in case it is warranted in virtue of something other than an inferential relation this knowledge has to any belief(s) of the subject.

It will help to have an example in mind to explain the key terms of these definitions. Thus, suppose that I am baking a loaf of bread, and I have set an alarm for 30 minutes. Time goes by, and upon hearing the alarm, I come to believe that 30 minutes have elapsed, acquiring knowledge. My knowledge that 30 minutes have elapsed is clearly inferential: it depends on the further knowledge (i) that I set the timer for 30 minutes, (ii) that if the timer is going off, that indicates that the amount of time I set it for has elapsed, and (iii) that a timer is going off. By contrast, the knowledge that a timer is going off is intuitively not inferential. Rather, like much of my perceptual knowledge, I come to have it directly when I perceive the timer going off.¹⁷ With these examples in mind, let us now consider the key parts of the definition.

¹⁶ Apart from the terminology, the definitions are the same as those influentially presented by Alston (1983). The main significant difference is that Alston recognizes two distinct ways in which knowledge might depend on other beliefs: “(a) having adequate evidence for the belief in question and (b) the belief in question having been arrived at by inference” (p.75). However, Bird (2016) has convinced me that a belief is based on evidence just in case it is based on an inference where that evidence is used. So, I think (a) and (b) collapse into a single option.

¹⁷ Here and throughout the paper I shall assume the standard position—which I take to be correct—that we can have direct perceptual knowledge of ordinary material objects and the like (such as that a timer is going off). For an alternative, more restrictive view, see McGrath (2017).

It is warranted...: To say that a knowledge state is warranted by the relation it bears to some other entities is to say that it is knowledge in virtue of these relations.¹⁸ In the examples, the knowledge that 30 minutes have elapsed is knowledge in virtue of the relation it bears to beliefs (i)-(iii); whereas, intuitively, the belief that a timer went off is not dependent on a relation to any belief. I shall restrict the term ‘justification’ (and cognates) to the specific form of warrant gained by inference.¹⁹

The crucial point is that warrant pertains to the specific property of *being knowledge*. For a knowledge state has many additional properties, such as being a mental state, a propositional state, or a belief state; and possession of these additional properties may be grounded in the subject’s having certain beliefs. For instance, my knowledge that a timer is going off is arguably dependent on my knowledge of what a timer is. However, as William Alston notes, recognizing this dependence is consistent with holding that the knowledge is direct because “[u]nless I know what it is to be *P*, I can’t so much as form the belief that *x* is *P*, for I lack the concept of *P*” (1983: 78). That is, knowledge of what a timer is, is a prerequisite to having propositional states of the form ‘ \ulcorner a timer is *F* \urcorner ’. Given this view, my knowledge that a timer is going off is dependent on my knowledge of what a timer is, but only insofar as it is a propositional state; but it is not warranted by this knowledge, since its being knowledge is not thus dependent on the belief. By the same token, we can claim that authoritative knowledge is direct even if, for instance, the child could not know that the picture represents a boa eating an elephant without knowledge of what a boa or an elephant are.

¹⁸ I prefer ‘warranted’ (and cognates) to Alston’s ‘epistemized’ for stylistic reasons. I should also emphasize that my use of the term is meant to be consistent with knowledge-first views (Williamson 2000). For instance, following a suggestion by Williamson, one might hold that a knowledge state is warranted in virtue of being reliable, but hold also that the kind of reliability required for knowledge is defined in terms of knowledge (ib., p.100).

¹⁹ For a similar usage, see Burge (1993, 2003).

...in virtue of (something other than) an inferential relation...: Only a particular form of dependence is relevant for determining whether some knowledge is direct. To see why, suppose that I know that my timer went off because I was standing close enough to the kitchen, and that I was standing close enough because I know that if I stood farther, I would not hear the timer. Then my knowledge that the timer went off is dependent on this latter piece of knowledge. But the dependency is merely causal, not inferential. If we were to state the justificatory grounds for the belief, we should leave that belief out. By contrast, my belief that 30 minutes have elapsed *is* inferentially dependent on other beliefs such as the belief that I set a timer for that amount of time. If we left out that proposition from an account of the justificatory grounds for the belief, the account would be incomplete, and would fail to explain why the belief is justified.

The inferential condition also bears on a standard distinction in epistemology between propositional and doxastic justification. Roughly, an agent has propositional justification iff she has good reasons to hold a certain belief; and she has doxastic justification iff she holds the belief *on the basis* of those reasons. Thus, doxastic justification is strictly stronger than propositional justification. To have inferential knowledge, a belief must be doxastically justified, which means it must be based on the beliefs that constitute the premises of the inference. This means that, at a minimum, the beliefs must be *used* to reach the conclusion.

To see why this is important consider two versions of the case above. Case 1 is just as described, where I form the belief that 30 minutes have elapsed, based on my other beliefs about the case. In case 2, I form the belief that 30 minutes have elapsed simply out of a whim of fancy, even though I possess the same beliefs as in Case 1. Clearly, only the belief in the first case is in good standing, and constitutes knowledge. This is because in the first, but not the second case, I *used* those beliefs as the basis of an inference to the conclusion that 30 minutes have elapsed.

...to...*belief(s)*: Finally, the definitions of direct and non-inferential knowledge are stated in terms of a dependence on *belief*. This is because a piece of knowledge can be direct while being dependent on some other non-doxastic operation of our minds, such as perceptual experiences, reliable processes, perceptual skills, etc. (cf. Alston (1983: 75)). Thus, one can hold that perceptual knowledge is direct, even if it is warranted in virtue of being based on perceptual experiences with a related content (say, the experience of a sound).

Having clarified the distinction between direct and indirect knowledge, let us apply it to authoritative knowledge, beginning with a direct account. According to such an account, someone who exercises authority is in a position to know certain facts—such as what a picture represents, or whether a cabin has been cleaned—and this knowledge is warranted in virtue of something other than an inferential relation to further beliefs. Just as I am in a position to know that a timer went off in virtue of features of my perceptual system and its exercise, the direct account holds that I am in a position to know the relevant facts on the basis of features of my authority and its exercise. Later, I shall try to identify what such features might be (Section 4).

In contrast, according to an inferential account, the knowledge that I come to have on the basis of exercising authority is warranted in virtue of an inference from further beliefs. At first sight, such an account may seem hopeless. One reason is that it seems to contradict our phenomenology. If you imagine yourself in the position of the father or the captain, it does not seem like you must consult your other beliefs to learn the name of your child, or the whereabouts of the sailors. Rather, you come to know as soon as you baptize the child or give the order. Second, there is a specific challenge stemming from the practical nature of authoritative knowledge. Given the dependence of the object known on practical knowledge, there do not appear to be enough facts

antecedent to the knowledge on the basis of which one could infer, for instance, that the name of the child is ‘Juliana’, or that the sailors would clean the cabin. And if there were such facts (e.g. if it were a fact that the father has a clear predilection for the name ‘Juliana’, or that the captain is obsessed with clean floors), the knowledge one could gain on their basis seems different from the kind of knowledge the father and captain have. Thus, only a direct account seems capable of capturing the distinctive character of authoritative knowledge.

Similar problems arise for an inferential account of agential knowledge. To see why they are far from decisive, I shall now turn to what I consider the best inferential account in the literature, due to Sarah Paul (2009). This account will then serve as the basis for an inferential account of authoritative knowledge that can answer these objections.

Paul begins from the well-agreed contention that agents generally have (a fairly) direct knowledge of their intentions.²⁰ Moreover, she argues that intentions generally and reliably lead to actions: generally, if someone intends to Φ , she is either Φ ing, or she will Φ .²¹ And again, most agents know this. Based on this knowledge, therefore, an agent can generally reason as follows:

- (1) I intend to Φ [in circumstances C].
- (2) (Generally) If I intend to Φ [in circumstances C], I am Φ ing/I will Φ .
- (3) Therefore, I am Φ ing/I will Φ .

(1)-(3) represent an inferential schema. On Paul’s account, an agent who reasons according to this schema is in a position to gain knowledge of her actions based on knowledge of her intentions. I have included in square brackets certain conditions that must hold for the agent to gain knowledge

²⁰ Paul (2012) argues that knowledge of intentions is based on knowledge of decisions. Because decisions nearly always result in intentions, we have an entitlement to believe that we intend to do something when we decide to do so.

²¹ There are some complications about the nature of the acts in question and the difference between knowledge of ongoing and future action that, though important for Paul’s account of agential knowledge, are irrelevant for the application to authoritative knowledge.

on the basis of such an inference. According to Paul, these include “knowledge of ability, conduciveness of circumstances, and one’s history as an agent [who reliably does what she intends]” (p.15). For our purposes, however, we may remain neutral with respect to what the circumstances are, whether they are part of the inferential content, and whether they need to be epistemic conditions (on a more externalist account, ability, conduciveness, etc., might suffice even if unknown to the agent). Given the general nature of (2), this is obviously a defeasible inference. Sometimes an agent who knows the relevant premises will reason according to the schema, and yet fail to know simply because the conclusion is false; but she will know when the conclusion is true.

To address the worry that an inferentialist account doesn’t reflect our phenomenology, Paul contends that inferential reasoning “can take place rapidly and automatically at a non-conscious level, without the mindful entertaining of premises or feeling of drawing a conclusion” (p.10). And to address the worry about insufficient evidence, Paul contends that the fact that the inference relies on knowledge of the agent’s intentions explains why she can know what she is doing (or will do), even though no external piece of evidence could give her or another person grounds for such a belief, antecedent of forming the given intention. Moreover, the fact that knowledge of intention is (in some sense) privileged explains why knowledge of actions based on it is special.

We can now give a parallel account of authoritative knowledge. On such a view, someone who exercises authority is in a position to gain knowledge by reasoning according to the following schema (for simplicity, I take the contents to be propositional, but nothing hangs on this):

Authoritative Inferential Schema

(4) I am exercising my authority to the effect that p .

(5) (Generally) If I exercise my authority to the effect that p , then p .

(6) Therefore, *p*.

So, for instance, captain O'Heguemon can reason that since she is ordering the sailors to clean the cabin and since orders of this sort are generally obeyed, the cabin will be cleaned. Exercising authority is an action. Therefore, we may assume that someone in authority will (normally) have agential knowledge of exercising authority. And, presumably, in many cases a person in authority will be in a position to know that her authority is reliable, as per (5). Finally, like Paul, a defender of an inferentialist account can point out that inferences can occur unconsciously to address the concern about phenomenology; and she can point out that knowledge of the exercise of authority seems to provide the required evidential basis for the conclusion, while explaining why this knowledge seems special (since based on agential knowledge).

3. Two Arguments for the Direct Account

We now have two accounts of authoritative knowledge, one according to which it is direct, and one on which it is inferential. Which of these is correct? We have seen that the initial considerations about phenomenology and lack of evidence that appeared to support the direct account are inconclusive. And since the inferentialist account explains the phenomenon of authoritative knowledge economically as simply a type of inferential knowledge, the onus now appears to be on the defender of the direct account. This section aims to discharge this onus by presenting two arguments in favour of such an account.

3.1. Chatty Charlie's Assertion

Imagine a man who cannot keep his thoughts in his head: whatever he thinks, he expresses out loud. Let's call him 'Chatty Charlie'. If you were to run into this peculiar man at the store, you

would hear him voicing out massive amounts of information. For instance, you would hear him voicing out all his perceptual beliefs: ‘there are tomatoes over here; there are apples over there; the floor is white; this tomato feels soft; it’s so cold in here!’ etc. In addition, of course, you would also hear information about his plans: ‘I still need to find eggs; where are the eggs?; I’ll walk over there; should I get a box or two?’ etc. Like us, Chatty Charlie processes information like this at amazing speeds, so you might have to imagine that he speaks really fast, or perhaps that unlike us who speak in a single voice, Charlie speaks in several voices at once, like a chorus (you may also need to imagine that we have a machine to help us disentangle everything that Charlie says).

There are powerful Wittgenstenian objections to the idea that all our thoughts could occur privately. However, who would object to the claim that it could all occur publicly as with Charlie? Hence, however remarkable, I take Charlie to be a possible being, and one very similar to us. The point of focusing on his case is that it allows us to evaluate a person’s thoughts directly and transparently, without worrying that the correctness of our evaluations might depend on features of the agent’s psychology (like unconscious states) that stand outside our purview.

With this in mind, suppose you met Charlie at the hospital, and you knew his situation was that of the single father above: he just lost his wife at childbirth, and he is holding in his hands a yet unnamed child. You see him holding her and you ask him, ‘What’s her name?’. Charlie responds, ‘He asked me what the child’s name is; well, we never settled this; what should I call her?’ and, *without saying anything else* relevant to the case at hand, he replies after these musings, ‘her name is Juliana, like her mother’.

How should we evaluate this assertion? It seems clear that it is in good standing. After all, Charlie has full authority over the name of the child, so why shouldn’t he be able to state what it is as soon as he decides it? If KNA is correct, therefore, we can infer from the correctness of this

assertion that Charlie possesses knowledge. This conclusion is in line with the direct account, according to which an exercise of authority is sufficient to provide knowledge. Charlie's assertion merely constitutes a limit case, where the exercise of authority is self-warranting because it is both an exercise of authority and an assertion: qua exercise of authority, it provides warrant for the belief it expresses, ensuring that it can amount to knowledge. Qua assertion, it is correct (by the lights of KNA) because Charlie possesses the requisite knowledge.

By contrast, an inferentialist account entails that the assertion is incorrect. This is because Charlie has not gone through an inference matching the authoritative inferential schema. If he had, we would have heard him say both:

(7) I am exercising my authority to name the child by naming her 'Juliana'.

(8) If (7), then her name is 'Juliana'.

So, according to the inferentialist account Charlie does not know that his child's name is 'Juliana' when he asserts this. Therefore, the assertion is incorrect. Because it predicts, wrongly, that such an assertion is incorrect, we should reject the inferentialist account.

One could object that the case is too farfetched to draw safe conclusions. But why would the farfetchedness of the case matter? After all, the case is one where there is a clear normative judgment: that Charlie asserts correctly. And clear normative judgments are generally not called into question by the farfetchedness of a case. For instance, consider the judgment that Voldemort acted wrongly when he used an *Avada Kedavra* spell to kill Harry Potter's parents. This judgment concerns a very farfetched situation, one where humans can kill using magic spells (unlike Charlie's case, I take this to be nomologically impossible). However, the fact that such a possibility is farfetched does not call into question the truth of the judgment. Like the judgment that

Voldemort acted wrongly in killing Harry's parents, our judgment that Charlie asserts correctly seems perfectly safe.

Another objection is that Charlie's speech-act is not a genuine assertion, and thus outside the scope of KNA. But what does that mean? It surely *sounds* like an assertion (it is speech in indicative form addressed to a listener with the aim of informing her); and it behaves exactly like one. For instance, the speech-act can be the proper origin of a testimonial chain, arguably the central function of assertion. The nurse can gain knowledge of the name of the child on its basis, and she can share this knowledge with others. Another sign that Charlie's speech act is a genuine assertion is that it gives rise to Moorean paradoxes. There would be a 'clash' of sorts if Charlie had instead asserted, 'Her name is Juliana, but I don't know whether that's her name'.²² KNA is widely regarded as the best explanation of Moorean paradoxes: the clash is explained by the fact that, though consistent, the conjuncts of a Moorean sentence cannot be jointly known. Hence, propositions of that form cannot be asserted without violating KNA. The fact that Charlie's speech-act can give rise to the paradox suggests it is an assertion, governed by KNA.

Finally, although Charlie's case is designed to allow a transparent normative assessment of his state of mind, one might object that it leaves out important features, dispositions in particular.²³ Aren't beliefs widely held to be dispositional states? If so, one might object that we should attribute to Charlie beliefs with content (7) and (8), even if Charlie has not given voice at the time to these dispositions. Alternatively, if Charlie is not even *disposed* to assert and act on and these beliefs, we should revise our judgment that his assertion is correct.²⁴

²² 'Clash' is DeRose's (2002) term for the air of incompatibility associated with these sentences.

²³ I thank a referee for pressing me on this point.

²⁴ I am here taking a fairly standard view that belief is a disposition to judge and act on those judgments.

In reply, it must be first noted that the argument is consistent with both a dispositional account of belief, and with the view that Charlie must have certain beliefs and dispositions to gain knowledge just on the basis of exercising authority.²⁵ So, it can be granted that if Charlie is not even disposed to assert (7) and (8), he lacks knowledge (why possession of those beliefs might be deemed necessary will be examined later). However, the argument *does* depend on the assumption that the beliefs must be occurrently manifested if they are to provide inferential grounds for a belief. We need this assumption because inference yields knowledge by providing doxastic justification, rather than just propositional justification. Since Charlie's case is stipulated to be such that the beliefs with the relevant content ((7) and (8)) are not even used (never mind used as part of an inference), the inferentialist account incorrectly predicts that Charlie's belief is unwarranted.

The considerations in this section apply equally to the case of the single father, the case from *The Little Prince*, and the case of the captain. To see how it applies to the latter, imagine a case where the captain is assigning posts to the sailors for the next day, employing assertions like the following: 'Sailors A and B will be in the main cabin', 'Sailor C will be in the top deck', etc. Suppose further that we know these assertions are direct exercises of authority, so that, in particular, the captain hasn't made an inference to the effect that her orders will be obeyed (to do this, we might imagine that Chatty Charlie is the captain, giving the orders as they occur to him). Given all these assumptions, the assertions still seem clearly correct.

Charlie's case thus forces a hard choice on the defender of the inferential account: either reject KNA or accept the direct account. Because I take there to be strong reasons to preserve KNA, I believe the argument forces us to accept the direct account.

²⁵ We saw this above in connection with knowledge that is required for concept possession (p.11).

3.2 Stipulative Knowledge

The second argument for the direct account is based on one offered by Burge (1993) in response to Chisholm (1987) to show that memory gives us direct warrant. Chisholm argued as follows: Certain mathematical and logical proofs are extremely long and complex, requiring anyone running through them to employ her memory to keep track of all the steps. This employment of memory plays a warranting role: without it, the conclusions reached would fail to constitute knowledge. Hence, any knowledge acquired in this way must be supported by a premise to the effect that memory is a reliable source of information, such as:

(MEMORY): If I (seem to) remember that p , then p .

However, the warrant we have for a claim like (MEMORY) is at least partly empirical. Hence, Chisholm concludes that there are many mathematical truths for which “we cannot be said to have an a priori demonstration of the conclusions” (p.30).

Burge’s response flips this argument on its head. He argues that since it is evident that we have non-empirical knowledge of at least some of the conclusions that we reach as a result of long deductions, it must be that no such step as (MEMORY) is needed for us to gain knowledge on the basis of memory. Instead, he argues, memory is a content-preserving faculty, that gives us direct access to the facts, and on the basis of which we are immediately warranted in forming a belief without the need of further epistemic support.

I want to argue that a parallel argument is available for thinking that authority gives us direct knowledge. To see this, consider stipulative steps, steps of the form “Let n be such and such”, which play a crucial justificatory role in many formal proofs. For instance, consider the following simple example of a proof showing that the area of an ABC isosceles triangle is 12 cm^2 , given that its base AC is 6 cm long and its legs AB and BC are 5 cm long:

1. Let l be a line running perpendicular to AC that intersects B .
2. Let D be the point at which l intersects AC : evidently, D is the half-point of AC .
3. Let db be the segment of l from D to B .
4. Then, the length of $db=ABC$'s height.
5. DBC is a right angle, so by the Pythagorean Theorem, the length of $db=4$.
6. Hence, ABC 's area= $4 \times 6 : 2 = 12\text{cm}^2$.

The proof is incomplete, relying on intuition where explicit grounds would be preferable (as in step 2). But it will do for our purposes. My interest is in steps 1-3, what I'm calling 'stipulative steps'. These steps clearly play a crucial justificatory role in the proof (e.g. without step 3, step 4 would be unjustified, since db would lack a referent).

Stipulative steps seem to be clear instances of exercises of authority: by taking the stipulative step, the person performing the proof settles that certain terms refer to certain entities (such as a line segment). This explains why their characteristic expression: 'Let $db \dots$ ' is in paradigm imperatival form (cf. 'let them eat cake', 'let there be light'). Indeed, it is attractive to think of stipulative steps as baptismal acts, like the one performed by the single father in naming the child. The only difference is that in this case the baptized subject is an abstract entity rather than a person.

Suppose, then, that inferentialism about authoritative knowledge were true. This would mean that to be warranted in believing that the line from D to B is db on the basis of premise 3, one would have to base one's beliefs in the further belief that:

(DECISION): If I decide to name the line db , then its name will be db .

One problem with this is the same one that Burge identified in relation to memory: if (DECISION) is part of the justification of the above proof, the knowledge we would gain on its basis would

seem to be empirical rather than a priori. This is because, like the grounds for (MEMORY), the grounds that a subject has for believing (DECISION) are empirical (since they will include facts about the reliability of one's memory and authority).²⁶ However, there is an even more basic problem, both in this case and in the case of memory: the (propositional) justification we gain by such a proof should not include *any* reference to mental states like memories or decisions, nor to the person performing the proof. The justification should consist entirely of the mathematical facts displayed in the proof. Since the inferential account entails otherwise, it should be rejected in favour of the direct account.

The argument seems to me decisive, both for memory and authority. It shows that when we gain knowledge on the basis of exercising these systems we can gain direct knowledge. The exercise of these systems thus plays an important warranting role, licencing beliefs whose content just is the content displayed in the conclusion of the proof; but beliefs *about these systems* need not play any warranting role whatsoever. This is why when these considerations (i.e. the contents of the beliefs thus involved) are warranted *a priori*, the belief reached on their basis can be *a priori* knowledge.

Notice that it is irrelevant to the argument whether the warranting role played by stipulative reasoning in a proof like the above is irreducible. There may be a proof that reaches the same conclusion that forgoes appeal to stipulative reasoning. By the same token, we may find more economic proofs that forgo the need to employ our memories where we previously needed to use them. However, it is sufficient for the argument if *some* proof that employs memory or authority yields *a priori* knowledge.²⁷

²⁶ Alternatively, if the grounds were a priori, one could argue, following Peacocke (2003), that we are *entitled* to form beliefs on the basis of these facts without the need of an inference.

²⁷ For the record, I doubt that logical proofs that employ the classical universal introduction rule (using stipulative steps that employ so-called 'arbitrary objects' (Fine 1985, 1983)) can be reduced to different inferential forms. King

It is possible, however, to challenge the scope of the conclusion for authoritative knowledge. For one may contend that the authority exercised in stipulative reasoning is a very special one. Hence, the thesis that authoritative knowledge in general is direct cannot rest on it. Thus, we might grant that a direct model holds for ‘Stipulative Authority’, but deny that it holds for other forms of authority.

The concession that at least some forms of authority give us direct knowledge is already significant. However, there are good reasons to defend the stronger position about authority in general. First, note that Stipulative Authority is not so different from other sorts of authority. As noted, it is natural to treat it as a form of baptismal authority. Second, anyone tempted by this line of response should consider whether they would draw the corresponding conclusion for the argument pertaining to memory: should we also say that there is a ‘Proof-Specific Memory’ which should be distinguished from ordinary memory as employed to remember recent moves in a game of chess, or the names of people we just met at a party? No: even if there are important differences between them, they should receive the same epistemic treatment. Of course, this is consistent with thinking that there are important *psychological* differences between kinds of memory (short-term/long-term; episodic/declarative; etc.). What I am claiming is that insofar as they belong to the same genus, memory, they do so in part because they yield direct knowledge.

The same holds for authority: the things over which we exercise authority have varied natures (as my initial cases brought out). As a result, there will be peculiarities to all these cases. In particular, there will be peculiarities with respect to the conditions that must obtain for the exercise of authority to be successful, and hence for the beliefs formed on their basis to be true. However, as far as acquiring knowledge is concerned, they all seem just the same: simply by

(1991) attempts a reduction of the sort, but it suffers from well-known difficulties (Breckenridge and Magidor 2012).

exercising authority one is in a position to gain knowledge of the matter over which one exercised authority. Hence, I am inclined to generalize from the conclusion reached on the basis of application of authority in mathematical reasoning to its application in other domains.

Still, one might accept the claim that the cases of stipulative knowledge and the knowledge of the name of the child are analogous, but dispute whether these should receive the same epistemic treatment as the case of the captain O’Heguemon, or of Anscombe’s builder. For, one might hold, the cases of stipulative and baptismal authority *incorrigibly ensure* that the relevant fact obtains, e.g. that *db* stands for the relevant line, or that the name of the child is ‘Juliana’. And it is their incorrigibility that enable them to yield direct knowledge.

There is, I think, something right in this response: it is partly in virtue of ensuring that the relevant fact obtains that authority can be a source of direct knowledge. However, the demand for incorrigibility is misguided. Therein lies the path to skepticism. Moreover, it does not yield the desired classification, since even baptismal authority is not an incorrigible source of knowledge. To see this, imagine that in yet another bleak development to the story, the father suffers a stroke and dies as soon as he says, ‘her name is ‘Juliana’, expressing his belief. The nurse mishears him saying, ‘her name is ‘Juana’’, and as a result, the wrong name is placed in the birth certificate of the child. In that case, the father’s belief would be false. Or again, suppose that the person exercises authority, but it is disregarded (as in the case of the baptism of John the Baptist that I discuss in Section 5). Once again, forming a belief on the basis of exercising authority could yield a faulty belief.²⁸

²⁸ Another way in which one might form a false belief is if one mistakenly uses one’s authority. For instance, the person performing the proof might use a variable that she is not permitted to use at that step of the proof (e.g. because it already has been given a referent).

Thus, I do not see an obvious reason to stop the generalization of the argument from stipulative knowledge to all cases of authority. I therefore conclude that authority in general is a basic knowledge source.

4. Keeping Track by Perception, Keeping on Track by Authority

I have argued that authoritative knowledge is direct. It is not warranted in virtue of bearing an inferential relation to other beliefs. In virtue of what, then, is authoritative knowledge warranted? In this section I present a positive account of authoritative knowledge that attempts to answer this question, and that helps us address some outstanding worries. I shall start by presenting an outline account of my favoured view of perceptual knowledge, to use as the basis for an account of authoritative knowledge.²⁹

Intuitively, we are in a position to have perceptual knowledge by exercising our perceptual systems, because such an exercise puts us in special contact with the world. The question is: what is the nature of this contact, and how is it achieved through perception, such as to yield direct knowledge? My favoured response has two central components. First, when exercising our perceptual systems, we are the subjects of experiences that put us in a position to exploit a world-dependent informational link to have thoughts about the world.³⁰ The perceptual link is informational or contentful, because it has correctness conditions (it is correct just in case the experience represents the world faithfully). It is world-dependent insofar as the link consists in the subject's bearing a relation to the world that is sensitive to changes in it. To illustrate, when I

²⁹ The two central sources of inspiration for the view are Ernest Sosa's virtue-theoretic account of knowledge in terms of competently formed judgment, developed over a series of writings (2015, 2010, 2007), and Gareth Evan's account of the discriminative knowledge which, following Russell, he took to be a requirement for singular thought (1982).

³⁰ My understanding of this first condition is where I am most indebted to Evans (1982, especially ch.5), though his influence is also evident in the discriminating component of the second condition.

experience a timer going off, there is an informational link between me and the timer. The link is informational because I am in a state that is correct just in case a timer is going off,³¹ and it is world-dependent because the link is maintained in virtue of the perceptual relation that I bear to the timer, which is sensitive to changes in the world (my experience will change as e.g. the volume of the timer changes, or when it stops). I can thus exploit my experience to make the judgment that a timer is going off.

Yet, such a link is insufficient for knowledge, because the information provided by the link, or the beliefs formed on its basis may be only accidentally correct. For instance, if I perceptually experience the timer going off, but I cannot discriminate this sound from other sounds likely to occur in my environment (e.g. a whistling teapot), then I might be the subject of an accurate experience of a timer going off, but if a true belief is formed on the basis of this experience, its truth will be accidental. Hence, it won't amount to knowledge. This suggests a second condition: the perceptual exercise must manifest a set of skills (such as discriminating skills) in virtue of which the beliefs formed on the basis of the system are reliably true.

Since my topic of investigation is not perceptual knowledge, I shall not explore this account further. However, by way of highlighting its attractions, I should emphasize the stringency of its epistemic demands. First, the appeal to perceptual experience renders this account stronger than so-called conservative views according to which having an appearance that *p* is sometimes sufficient for warrant, since having an experience that *p* entails having such an appearance.³² Second, the appeal to the notion of an information link renders this account stronger than causal accounts according to which a certain type of causal relation generates warrant.³³ Finally, the

³¹ We could give graded conditions in terms of accuracy, and there are some advantages to this view, but for our purposes the non-graded account will suffice.

³² See Huemer (2013) for a helpful summary of work on conservatism.

³³ See e.g. Goldman (1967).

appeal to reliable skills makes this account stronger than many broadly reliabilist views, according to which belief-formation through a reliable process is sufficient for warrant (the appeal to skills makes this a version of reliabilist virtue theory).³⁴ Thus, if we can show that an agent who exercises authority meets parallel conditions to the ones outlined, we will thereby have shown that the claim that authoritative knowledge is direct can be vindicated from a broad range of epistemic paradigms.

With this in mind, what I seek to show is (i) that exercises of authority forge an informational link of the sort described between the bearer of authority and the world that can be exploited by her so as to have a world-tracking thought, and (ii) that when she does this, she exercises some kind of skill that ensures the reliability of these thoughts. The aim, as with the case of perception, is not to show that (i) and (ii) obtain whenever authority is exercised, but rather to show that they are often enough met in tandem, yielding direct knowledge in a broad range of situations.

When one exercises authority, one is thereby in a contentful state. This is so regardless of what we take the locus of such an exercise to be, whether a mental state like a decision, or a speech act like a command. For both decisions and speech acts have determinate correctness conditions. Of course, the direction of fit is different in the case of authority, because authority is a practical rather than a theoretical power. Thus, as noted while explaining why authoritative knowledge is practical, it is the representation that sets the standards of correctness for the world: what the sailors do is judged by the standards set by the order.

Arguably, exercises of perception, unlike exercises of imagination, are object-dependent. One cannot be in a perceptual state with respect to an object unless the object exists and is the source of information for the perceptual state, whereas one can imagine an object that doesn't

³⁴ As noted in n28.

exist. In this way, being in a perceptual state ensures that what one represents in perception is actually the case. At first sight, it is tempting to group exercises of authority with exercises of the imagination in this respect, since one can exercise authority to the effect that p , and yet p may not obtain. The captain's orders can be disobeyed or ignored, and the father's baptismal act may be ineffectual. Yet, there is an important disanalogy grounded in the essence of authority. The cases where an exercise of authority is ineffectual are abnormal, calling for further explanation. By contrast, it is normal to imagine what is not the case, and that doesn't call for further explanation. This means that a match between mind and world in the case of authority will be non-accidental in the way in which it would be in the case of imagination. Since the aim is to validate the view that authority *sometimes* yields knowledge directly, I take this non-accidental match to be sufficient for warrant.

In the good case where the exercise of authority is effective, there will thus be an informational link of the needed sort between mind and world. Sometimes the informational link will be largely characterized in causal terms, as with perception. This is so in the case of the builder and of the captain, whose orders are the first stage of a causal chain that proceeds from them unto their subordinates and then unto the world through their respective skills. In other cases, the informational link will be characterized mainly in conventional terms. The transfer of information is made possible by the existence of certain conventions in virtue of which certain acts settle some matter. This applies to the artist and the father. Artistic and linguistic conventions ensure that the nature of the drawing and the name of the child are settled by the determinations of the artist and the father, correspondingly. Both sorts of links, however, enable a world-tracking thought, since the content of the representation is world-sensitive in the way outlined. Again, though, we find that due to the practical nature of authority, the manner in which the tracking takes place is

reversed. Whereas experiences adjust in response to the world, it is the world that adjusts in response to the dictates of authority. Put otherwise, theoretical systems keep track of the world, whereas practical systems keep the world on track.³⁵ If the artist had decided that the picture represents a falling tent rather than a boa, or if the sailors had encountered a problem with their initial strategy for cleaning the cabin's floors, the world would have adjusted to match the representations. In this way, an exercise of authority can give rise to a world-sensitive informational link that can be exploited by the subject to have thoughts about the matter, as per (i).

To explain how exercises of authority meet condition (ii), we need to take account of its social nature. For, clearly, the skills that can be manifested by the individual exercising authority won't account for the reliability necessary for knowledge. In this respect, authority is analogous to testimony rather than perception. The recipient of testimony may exercise certain skills in determining whether they should trust the testifier, but they must ultimately rely on the testifier's own epistemic competence. Indeed, this is part of the point of testimony, that it enables us to expand our knowledge well beyond the limit of our own epistemic powers, by relying on others's.³⁶ The existence of this reliance system enables us to take credit for epistemic work that we may not even be capable of doing on our own. This is what happens when, for instance, a patient trusts the doctor's verdict as to what an MRI image shows even when she herself is utterly incapable of discerning what the doctor sees in the MRI. The patient can learn that (e.g.) the MRI shows a tumour even though she cannot discern it, by relying on the discriminatory skills of the doctor.

³⁵ The contrast shouldn't be overemphasized, though. It would be hard for creatures like us either to keep the world on track without the help of our theoretical systems (e.g. checking that matters are going as they should), or keep track of the world without the help of our practical systems (moving around an object to better determine its nature as revealed by sight).

³⁶ I am here indebted to Goldberg (2006) and Baker and Clark (2017).

In a similar way, systems of authority are fundamentally reliance systems. Those in authority are relied upon to make good determinations, while subjects of authority are relied upon to act on those determinations.³⁷ The skills of those subject to authority ensure the reliability of the information available through the informational link created by authority. Conversely, where these skills are lacking, the conditions for direct knowledge will be lacking too. Thus, if the sailors are incompetent at cleaning, or the workers incompetent at building, the captain and the director will not gain knowledge on the basis of exercising authority even if they form a true belief on this basis.³⁸ However, when such competence is exercised in response to the dictates of authority, the authoritative exercise possesses the needed reliability, and the reliability is established through the exercise of skills that the bearer of authority can take credit for, given that a system of reliance is in place. In occasions when this happens, the bearer of authority will also meet condition (ii), and will thus be in a position to gain direct knowledge on the basis of exercising authority.

5. Objections and Implications

With my positive account on the table, I shall now turn to clarifications and objections, beginning with a question about the scope of the thesis: Why do I restrict my claims to authority? Why not say that there is a more general kind of knowledge—‘powerful knowledge’, if you like—that someone possesses whenever they exercise any power whatsoever? Addressing this question will provide an occasion to further clarify the positive account.

³⁷ This is a central theme in Raz’s writings on authority (2010, 1985a, 1985b, 1981).

³⁸ In the conventional cases, such as those of the artist, the father, and stipulative knowledge, the skills of the person in authority are the most fundamental, though even here uptake skills may be relevant, in light of the kind of failures we find in cases such as the baptism of John the Baptist, discussed below. We could easily imagine a similar case where a baptismal act fails on account of failure of competence on the part of an audience that fails to recognize the baptismal character of an act.

First of all, not every exercise of a power will meet condition (i) of the account. For not every exercise of a power is contentful in the required way. For instance, my power to crush a beetle may be manifested without any awareness, as when I step on it on my way to turn off the timer in my kitchen. These powers thus fail to create an informational link between the world and the subject.

Still, there seem to be exercises of non-authoritative powers that are conscious, and would thus seem to provide for the necessary link. Consider the bully in the playground who tells another child, ‘You will give me your lunch right now’, and suppose the child gives him his lunch. The bully thereby exercises power without possessing authority, and does so in a fully conscious way, one that seems to give rise to an informational link of the right sort, by the lights of my account. Finally, it is not hard to imagine that the bully’s power is extremely reliable, so that he generally gets what he wants. So, is he in a position to acquire direct knowledge? Not so, in my view. For it is part of my account that the reliability must arise from an exercise of skill for which the person can take epistemic credit. And though the bully may get some credit for the way in which he exercises his own power (the fierceness in his voice, a menacing glance, etc.), he does not get credit for the skills exercised by the bullied child in ensuring that his demands are met. What unifies the action of truster and trustee, and commander and subordinate, is the existence of established systems of reliance that enable different parties to take credit for another’s actions. Since such a system is not in place in the case of the bully and the bullied (this is what makes the bully’s demand an exercise of mere power, rather than authority), I am of the mind that there is also no direct knowledge in this case (though, of course, he can gain inferential knowledge).³⁹

³⁹ Perhaps it is better to treat this as a case where the bully has or gains de facto authority, even though it is illegitimate (cf. n5). *If* this is the right analysis of the case, then I hold that the bully can acquire direct knowledge on the basis of exercising authority (supposing the other features of the account are met). I should nonetheless register

Another important worry stems from reflection on certain beliefs that seem to be required to gain knowledge on the basis of authority, beliefs that have not been given a role in the direct account. For instance, focusing on the case of the captain, it is tempting to think that to acquire authoritative knowledge he must rely on beliefs such as (a)-(d):

- (a) The belief that she is exercising authority.
- (b) The belief that if she gives the order, it will be accomplished.
- (c) The belief that her commands are reliably obeyed.
- (d) The belief that the sailors are competent at the given task.

After all, we may be unwilling to ascribe knowledge to the captain if we found that she lacked these beliefs. Could she really know that the cabin would be cleaned if she exercised her authority unwittingly, or if she was completely ignorant about how obedient or competent her sailors are at accomplishing this task? Certainly, a captain who did not take such matters into account in giving orders would not be very effective leader at seas!⁴⁰

As I have emphasised, though, the direct account is neutral with respect to whether the captain would need to *have* any of these beliefs to acquire knowledge. The sole commitment is to the claim that they need not be used as part of an inference to gain authoritative knowledge. We saw this in connection with knowledge that is required for concept-possession. Therefore, a defender of the direct account can agree that one must have (a)-(d) to gain authoritative knowledge.

Moreover, recall that my aim is to show that authority is as good a candidate for being a basic knowledge source as other well-recognized sources such as perception or testimony. However, a parallel objection seems equally powerful against a direct account of perceptual

that my views on this matter are hesitant, and would be open to being convinced that the thesis should be broader in scope, encompassing forms of consciously exercised powers other than authority.

⁴⁰ I thank an anonymous reviewer for forcefully pressing me on these points.

knowledge. For it seems equally plausible to hold that my knowledge that a timer is going off is dependent on:

- (e) The belief that I am perceiving a timer going off.
- (f) The belief that if I am perceiving a timer going off, then there is a timer going off.
- (g) The belief that my perception is reliable.
- (h) The belief that my hearing system is competent at detecting such sounds.

This suggests that the present objection is of a more general sort, which must be answered by a general account of direct knowledge, which is beyond the scope of this paper.⁴¹

On the other hand, it would be more satisfying if we could offer an outline of an error theory that would explain why we are inclined to think that such beliefs are important without ascribing to them an inferential role. It is here where it helps to appeal to the positive account.

Consider, first, (a) and (e), beliefs that you are exercising authority or perceiving. The grain of truth here is that on any account you must be exercising the relevant systems, and you will normally know when you are. We normally know that we perceive when we perceive, and that we are exercising authority when we are. When we are forced to consider cases where we do not have the relevant knowledge then we must consider either (i) cases where we are not in the relevant condition (perceiving, or exercising authority), or (ii) abnormal cases. Either way the requirements for direct knowledge are lacking, given my positive account. This is easy to see for (i), since in these cases there isn't an informational link of the required sort. With respect to (ii), consider the kinds of cases where one might perceive or exercise authority without knowing that one is: they would have to be cases where, for instance, one lacks the relevant concepts, or the perceptual or

⁴¹ For such a general reply, see Alston's (1983) response to what he calls the "Level Ascent" argument (pp.80-91).

authority system is unreliable or malfunctioning. However, lack of concepts and unreliability are again predicted by the direct account to hamper knowledge acquisition.⁴²

To account for the other beliefs on the list, we can appeal to the second part of the account, in terms of reliable skills. There are two (compatible) response strategies. The first strategy is to claim that the skills required to acquire direct knowledge on the basis of authority are partly constituted by capacities that bear a close relation to these beliefs. So, for instance, one might claim that corresponding to (b), (c), and (d), the captain must possess the skills of:

(b*) Ensuring that if she gives the order, it will be obeyed (e.g. by only giving orders that are likely to be efficacious).

(c*) Ensuring that her orders are reliably obeyed (e.g. by enforcing norms that ensure obedience).

(d*) Ensuring that the sailors are competent at the given task (e.g. by correctly selecting which subordinate to assign to which task).

A defender of the direct account can then claim that beliefs (b)-(d) seem required only because the skills at tasks (b*)-(d*) are required.

A second, more concessive strategy is to grant that the relevant beliefs are required, but only because someone who lacked them would lack the required reliability. For, imagine a captain that (e.g.) was altogether doxastically ambivalent about the competence of her sailors, even when they are fully competent. It would be natural to think that such a captain is not correctly monitoring her epistemic situation, and is therefore unreliable with respect to matters concerning her authority.

⁴² Or, at least, the agent will have good reason to think the systems are malfunctioning. There are three alternative ways to treat these cases: (1) treat them as cases of defeat, and, as many accounts of direct knowledge do, add an anti-defeat condition to the account; (2) treat them as cases of defeat, but argue that they amount to knowledge despite our intuitions to the contrary (as in Lasonen-Aarnio (2010)); or (3) argue, along the lines argued below for beliefs (b)-(d), that someone who forms beliefs in situations where they have good reason to think the systems are malfunctioning fails to form beliefs reliably.

Hence, she cannot gain knowledge on the basis of authority. Again, though, this story makes belief-*possession* relevant only insofar as it affects reliability. It is silent on whether the person must *use* these beliefs to gain knowledge. Hence, it is consistent with the direct account.

Throughout much of the paper I have emphasized the similarities between authoritative knowledge and other forms of basic knowledge, such as perceptual and testimonial knowledge. However, I also emphasized from the start that authoritative knowledge is practical knowledge, “the cause of what it understands”. This reverse-dependence is reflected at various points in the positive account, especially in the direction of information transmission in the links forged by authority, and in the fact that someone in authority gains knowledge largely by using skills that keep the world on track, rather than by keeping track of it. However, the very notion of practical knowledge sometimes occasions epistemic suspicion because it is viewed as a form of wishful thinking, giving rise to beliefs on the basis of how we want the world to be rather than how it is (e.g. the Captain wants the cabin to be cleaned, so she forms the belief that it will be so). But wishful thinking is widely thought to be incompatible with knowledge.

This objection is often motivated by reflection on cases such as the one famously presented by William James, with a climber considering whether to take a leap (Langton 2004). Here is James’s advice to the climber:

Have faith that you can successfully make it, and your feet are nerved to its accomplishment. But mistrust yourself, and think of all the sweet things you have heard the scientists say of maybes, and you will hesitate so long that, at last, all unstrung and trembling, and launching yourself in a moment of despair, you roll into the abyss. In such a case (and it belongs to an enormous class), the part of wisdom as well as of courage is to

believe what is in the line of your needs, for only by such belief is the need fulfilled. (James 1979: 53-4)

James's contention appears to be that since the climber's belief is self-fulfilling and he cannot accomplish his plan without it, he should form it.⁴³ Now, many of us agree that it is correct *all things considered* for him to form these beliefs. Nonetheless, they are epistemically faulty, in a way incompatible with knowledge. But wherein lies their fault? It is widely thought to lie in the practical nature of their origin, the fact that they are based on a practical attitude.⁴⁴ If correct, this explanation would raise problems for the notion of authoritative knowledge, knowledge gained just on the basis of exercising authority. Against this, I wish to argue now that a better explanation is available just in terms of the purely epistemic features of the case.

On a natural reading of James's story, the climber doubts whether he would succeed at the leap: the climber seems to regard his failing as an epistemic possibility up until the time when he has acted.⁴⁵ That in itself provides us with a sufficient explanation of the faultiness of the belief: a belief that *p* cannot amount to knowledge if the subject regards it as a live possibility that *p* is false (a lesson from lottery cases). A belief formed in such a case would therefore be necessarily faulty, falling short of knowledge.

Suppose we sidestep this worry by considering a climber that is fully confident in his success. Now the question is whether such confidence is warranted. If the positive account I have

⁴³ This alerts us to a different problem: knowledge cannot be warranted by practical reasons (reasons, for instance, to do with desirability or worthwhileness). I agree that forming beliefs on the basis of such reasons is epistemically suspect (though see, e.g. Marušić and Schwenkler (2018) and Rinard (2017, 2018) for the contrary view). However, the question is besides the point here. On my account, the epistemic grounds for authoritative beliefs are exercises of authority rather than the reasons for the exercising authority, which may of course be practical. To use a distinction familiar from ethical contexts: the former are backward-looking grounds, the latter forward-looking.

⁴⁴ The most sophisticated treatment of this worry is by Langton (2004), targeting Velleman's (2001) cognitivist account of intention. Note, however, that on Langton's considered view (2009), an externalist account of knowledge may suffice to allay her initial worries. The present considerations are meant to support this view.

⁴⁵ This is suggested by the sentences surrounding the quoted passage, where James speaks of "may"s and "maybe"s that are most naturally taken as epistemic (p.53-4).

offered is on the right track, this is a question about whether the climber has an appropriate informational link to the matter, *and* the link is reliable on account of his skills. As the story of the climber is usually told, it is tempting to suppose that it is not appropriately reliable: the fact that simply musing about what one has “heard the scientists say of maybes” could lead to failure arguably makes the belief too flimsy to constitute knowledge. Such considerations should either not be in the horizon of a skilled agent, or else they should not make a difference to his success. Once more, then, we have a perfectly good explanation of the belief’s faultiness that doesn’t appeal to the belief’s practical origin.

Let us therefore focus on the case of a climber who has no doubts about his ability to make the jump, and whose decision is part of a manifestation of skill, leading to a true and reliable belief. Then I have a clear intuition that the climber *does* know that he will make the jump. He knows this, as well as I know that I will make coffee when I take a break from writing. If this is right, then the problem with the climber’s belief is not found in its practical origin. Rather, if we apply the same epistemic standards to practical knowledge as we would anyway apply to paradigm examples of theoretical knowledge (as my positive account does), we see that the factors that mar his belief and prevent him from gaining knowledge are factors that would equally prevent him from gaining perceptual or testimonial knowledge.

Recognizing the existence of a form of practical knowledge that extends beyond the agent’s own actions may give us insight into the nature of other phenomena such as semantic knowledge. It has been thought that one is always in a position to know what one means: meaning is luminous. Part of the purchase for this view comes from the fact that what we mean in a context of conversation is quite generally up to us. Just as the author of the picture with which we began can decide what the picture represents, we can generally let our words mean

what we intend them to mean. As such, the knowledge that one has of what one means is (at least partially) authoritative. However, the qualifications (“generally” and “partially”) are important: language is a communal phenomenon, and no single user has sole authority over what is meant by certain signs. She shares that authority with other speakers of the language, so that meaning is dependent on many facts beyond her control. As such, one cannot always know what one means, contrary to the thesis that meaning is luminous.⁴⁶ Recognizing semantic knowledge as a species of authoritative knowledge allows us to account for what is special about semantic knowledge without holding the implausible thesis that meaning is luminous.

The case of linguistic meaning highlights only one of several complications that arise when we turn to cases of shared authority. Another complication arises from the fact that to be effective, authority must be recognized, which it often is not. This can jeopardize even authoritative knowledge that seems quite secure, like artistic or baptismal knowledge. The story from *The Little Prince* with which we began may be an example of the first failure, as the child’s drawing is initially believed by adults to be of a hat (this is what forces the child to draw the second picture). A similar failure is illustrated by the story of the baptism of John the Baptist, as told in Luke’s Gospel:

And it came to pass, that on the eighth day they came to circumcise the child; and they called him Zacharias, after the name of his father. And his mother answered and said, Not so; but he shall be called John. And they said unto her, There is none of thy kindred that is called by this name. And they made signs to his father, how he would have him

⁴⁶ That semantic knowledge is not luminous is of course a central lesson drawn by Williamson in the context of his anti-luminosity argument (2000). In Piñeros Glasscock (forthcoming-b) I show that a powerful version of the objection applies to practical knowledge as well.

called. And he asked for a writing table, and wrote, saying, His name is John. And they marvelled all. (Lk. I:59-63)

In the story, Elizabeth attempts to name her child ‘John’. However, her authority to name the child is rejected by her audience, who turn instead to Zacharias, whom they recognize as an authoritative source. Only once *he* exercises authority is John’s name recognized as such. An unjust distribution of practical resources thus results in an unjust distribution of epistemic resources: Zacharias’s authority is a source of direct knowledge, whereas Elizabeth’s is not, because only the former’s is appropriately recognized.⁴⁷

I view such cases as deviations from the standard one where authority is recognized or obeyed, and the person can believe accordingly. This is to treat them as the epistemological analogues of illusions and hallucinations, understood as deviations from the case where the person perceives correctly, and can believe on this basis. This is not to deny that the questions raised by these deviant cases are easy. They are obviously not, but I believe they require a more general treatment. Since this is also a good place to stop, I close with one last illustration of the phenomenon I have aimed to elucidate: I can tell you—because I know so—that the paper ends here.⁴⁸

References

⁴⁷ I believe this falls short of being an epistemic injustice in Fricker’s sense (2007). This is because although Elizabeth incurs an epistemic harm, the primary way in which it is inflicted is not epistemic.

⁴⁸ Earlier versions of this paper were presented at the Epistemology Colloquium at Queen’s University in Kingston, Ontario, and at Yale University. I thank the participants for discussion and comments, especially David Bakhurst, Jeremy Butler, and Chirstine Synowich. I also thank Jewelle L. Bickell, Maxime Lepoutre, and Christopher M. Cloos for comments, and Allison Glasscock, Mike Deigan, Jo Demaree-Cotton, Eric Guindon, Mark Maxwell, Daniel Moerner, Jessie Munton, Paul Schilling, Tim Williamson, and Yuan Yuan for discussion. The paper was substantially improved by the comments of two anonymous referees. Finally, thanks especially to David Charles, Steve Darwall, Reier Helle, Jason Stanley, Gideon Yaffe, and Zoltán Gendler Szabó for comments and discussion on earlier drafts. Research on this project was supported by the University of Toronto-Mississauga Postdoctoral Fellowship Award Program, sponsored by their Philosophy Department and the UTM Office of the Vice-Principal, Research.

- Alston, William P. 1983. 'What's wrong with immediate knowledge?', *Synthese*, 55: 73-96.
- Anscombe, G. E. M. [1957] 1963. *Intention* (Basil Blackwell: Oxford).
- Aristotle. 1995. *Complete works of Aristotle* (Princeton, NJ).
- Austin, John L. 1975. *How to do things with words* (Harvard University Press: Cambridge).
- Baker, Judith, and Philip Clark. 2017. 'Epistemic buck-passing and the interpersonal view of testimony', *Canadian Journal of Philosophy*: 1-22.
- Bennett, Jonathan. 1988. *Events and their Names* (Oxford University Press: Oxford).
- Bird, Alexander. 2016. 'Evidence and Inference', *Philosophy and Phenomenological Research*: 299-317.
- Breckenridge, Wylie, and Ofra Magidor. 2012. 'Arbitrary reference', *Philosophical Studies*, 158: 377-400.
- Burge, Tyler. 1993. 'Content preservation', *The Philosophical Review*, 102: 457-88.
- . 2003. 'Perceptual entitlement', *Philosophy and Phenomenological Research*, 67: 503-48.
- Chisholm, Roderick M. 1987. 'The truths of reason'.
- de Saint-Exupéry, Antoine. 2001 [1943]. *Le Petit Prince* (Harcourt, Inc.).
- DeRose, Keith. 2002. 'Assertion, knowledge, and context', *The Philosophical Review*, 111: 167-203.
- Douven, Igor. 2006. 'Assertion, knowledge, and rational credibility', *The Philosophical Review*, 115: 449-85.
- Evans, Gareth. 1982. *The varieties of reference* (Oxford University Press: Oxford).
- Fine, Kit. 1983. 'A Defence of Arbitrary Objects', *Proceedings of the Aristotelian Society*, 57: 55-77.
- . 1985. *Reasoning with Arbitrary Objects* (B. Blackwell).
- Fricker, Miranda. 2007. *Epistemic injustice: Power and the ethics of knowing* (Oxford University Press).
- Goldberg, Sanford. 2006. 'Reductionism and the distinctiveness of testimonial knowledge.' in Jennifer Lackey and Ernest Sosa (eds.), *The Epistemology of Testimony* (Oxford University Press).
- Goldman, Alvin I. 1967. 'A causal theory of knowing', *The Journal of Philosophy*, 64: 357-72.
- Hawthorne, John. 2004. *Knowledge and lotteries* (Oxford University Press).
- Hill, Christopher S., and Joshua Schechter. 2007. 'Hawthorne's lottery puzzle and the nature of belief', *Philosophical Issues*, 17: 120-22.
- Huemer, Michael. 2013. 'Phenomenal Conservatism.' in *Internet Encyclopedia of Philosophy*.
- James, William. 1979. *The Will to Believe and Other Essays in Popular Philosophy* (Harvard University Press: Cambridge, MA).
- King, Jeffrey C. 1991. 'Instantial terms, anaphora and arbitrary objects', *Philosophical Studies*, 61: 239-65.
- Lackey, Jennifer. 2007. 'Norms of assertion', *Noûs*, 41: 594-626.
- Langton, Rae. 2004. 'Intention as faith', *Royal Institute of Philosophy Supplement*, 55: 243-58.
- . 2009. 'Speaker's Freedom and Maker's Knowledge.' in *Sexual solipsism: Philosophical essays on pornography and objectification* (Oxford University Press).
- . 2015. 'How to Get a Norm from a Speech Act', *The Amherst Lecture in Philosophy*, 10: 1-33.
- Lasonen-Aarnio, Maria. 2010. 'Unreasonable knowledge', *Philosophical Perspectives*, 24: 1-21.
- Marušić, Berislav, and John Schwenkler. 2018. 'Intending Is Believing: A Defense of Strong Cognitivism', *Analytic Philosophy*, 59: 309-40.

- McGrath, Matthew. 2017. 'Knowing what things look like', *Philosophical Review*, 126: 1-41.
- McKinnon, Rachel. 2013. 'The Supportive Reasons Norm of Assertion', *American Philosophical Quarterly*, 50: 121-35.
- . 2015. *Norms of Assertion: Truth, Lies, and Warrant* (Palgrave-Macmillan).
- Moore, Michael S. 2010. *Act and crime: The philosophy of action and its implications for criminal law* (Oxford University Press).
- O'Brien, Lucy. 2007. *Self-knowing agents* (Oxford University Press).
- Paul, Sarah K. 2009. 'How We Know What We're Doing', *Philosophers' Imprint*, 9: 1-24.
- . 2012. 'How we know what we intend', *Philosophical Studies*, 161: 327-46.
- Peacocke, Christopher. 2003. *The Realm of Reason* (Oxford University Press).
- Piñeros Glasscock, Juan S. 2019. 'Reason in Action in Aristotle: A Reading of EE 5.12/ NE 6.12', *Journal of the History of Philosophy*, 57: 291-417.
- . forthcoming-a. 'Alientation or Regress: On the Non-Inferential Character of Agential Knowledge', *Philosophical Studies*.
- . forthcoming-b. 'Practical Knowledge and Luminosity', *Mind*.
- Raz, Joseph. 1981. 'The Authority of Law', *Mind*, 90: 441-43.
- . 1985a. 'Authority and justification', *Philosophy & Public Affairs*: 3-29.
- . 1985b. 'Authority, Law and Morality', *The Monist*, 68: 295-324.
- . 2010. *Between Authority and Interpretation: On the Theory of Law and Practical Reason* (Oxford University Press).
- Reynolds, Steven L. 2002. 'Testimony, knowledge, and epistemic goals', *Philosophical Studies*, 110: 139-61.
- Rinard, Susanna. 2017. 'No Exception for Belief', *Philosophy and Phenomenological Research*, 94: 121-43.
- . 2018. 'Believing for Practical Reasons', *Noûs*.
- Schechter, Joshua. 2017. 'No Need for Excuses: Against Knowledge-First Epistemology and the Knowledge Norm of Assertion.' in J. Adam Carter, Emma C. Gordon and Benjamin W. Jarvis (eds.), *Knowledge First: Approaches in Epistemology and Mind* (Oxford University Press: Oxford).
- Schwenkler, John. 2015. 'Understanding 'Practical Knowledge'', *Philosophers' Imprint*, 15.
- . 2019. *Anscombe's Intention: A Guide* (Oxford University Press).
- Setiya, Kieran. 2007. *Reasons without rationalism* (Princeton University Press).
- . 2008. 'Practical knowledge', *Ethics*, 118: 388-409.
- . 2009. 'Practical Knowledge Revisited', *Ethics*, 120: 128-37.
- . 2016. 'Anscombe on Practical Knowledge.' in, *Practical Knowledge: Selected Essays* (Oxford University Press).
- Slote, Michael. 1979. 'Assertion and belief', *Papers on language and logic*: 177-90.
- Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I* (Oxford University Press).
- . 2010. 'How competence matters in epistemology', *Philosophical Perspectives*, 24: 465-75.
- . 2015. *Judgment and agency* (Oxford University Press, USA).
- Thomason, Richmond H. 1990. 'Accommodation, Meaning, and Implicature: Interdisciplinary Foundations for Pragmatics.' in Jerry Morgan Philip Cohen, and Martha Pollack (ed.), *Intentions in communication* (MIT Press: Cambridge, MA).
- Unger, Peter. 1975. *Ignorance: A Case for Scepticism* (Oxford University Press USA).

- Velleman, J David. 2001. "Practical Reflection " In. Stanford, CA: David Hume Series: CSLI Publications.
- Weiner, Matthew. 2005. 'Must we know what we say?', *The Philosophical Review*, 114: 227-51.
- Williamson, Timothy. 2000. *Knowledge and its Limits* (Oxford University Press).
- Witek, Maciej. 2013. 'How to Establish Authority with Words: Imperative Utterances and Presupposition Accommodation.' in Anna Brożek (ed.), *Logic, Methodology and Philosophy of Science at Warsaw University, Warszawa 2013*.
- Zagzebski, Linda Trinkaus. 2015. *Epistemic authority: A theory of trust, authority, and autonomy in belief* (Oxford University Press).