



Collective ignorance: an information theoretic account

Christopher Ranalli¹ · René van Woudenberg¹

Received: 20 February 2019 / Accepted: 13 August 2019 / Published online: 13 September 2019
© The Author(s) 2019

Abstract

We are ignorant knowers. This paper proposes an information theoretic explanation of that fact. The explanation is a conjunction of three claims. First, that even in those dimensions where we are capable of picking up information, there is information that we don't pick up. Second, that there can be dimensions of information for which we lack the capacity to pick up any information whatsoever. Third, that we don't know whether the faculties and cognitive capacities we are endowed with process all the information that they pick up.

Keywords Collective ignorance · Information · Dretske · Epistemic finitude · Intellectual humility · Information insensitivity

There are lots of things we don't know but little by way of a unified explanation of why we don't know these things. Much of the literature on ignorance¹ focuses on the extent,² the nature³ and varieties⁴ of ignorance, and the value⁵ of ignorance. Much of it

¹ See Gross and McGoey (2015).

² Rescher (2009).

³ Van Woudenberg (2009) and Le Morvan and Peels (2016) argue against its being absence of knowledge and in favor of its being absence of true belief. See El Kassir (2018) for an extensive overview and critique of the different conceptions of individual ignorance. She favors a proposal on which ignorance is a disposition that manifests itself in false beliefs, no beliefs, as well as in certain sorts of epistemic attitudes or intellectual character traits.

⁴ Nottelman (2016) distinguishes factual, objectual and erotetic ignorance; he also argues that some kinds of ignorance allow of degrees, whereas others don't. There are also varieties of ignorance, such as white ignorance, which makes links between ignorance and injustice. See Mills (2007).

⁵ Pritchard (2016).

✉ Christopher Ranalli
c.b.ranalli@vu.nl

René van Woudenberg
r.van.woudenberg@vu.nl

¹ Department of Philosophy, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

also focuses on whether ignorance can excuse one for some of one's moral⁶ and legal⁷ missteps, and on when ignorance is culpable.⁸ And almost all of it concerns ignorance of individuals. This paper, by contrast, is about a kind of *collective ignorance*, what we call *collective human ignorance*, and its aim is to both describe and explain it in a unified way in information theoretic terms.⁹

By collective human ignorance we have in mind what remains after totalizing all that humans collectively know. And 'Humans collectively know that p ' is true if and only if at least one human being, now or in the past, knows or has known that p . This definition entails that even if only Jones, and no one else, knew what he did on New Year's Day 1641, we collectively know what Jones did that day. If no one knows (no one has a true belief about) how many breaths Mary drew yesterday, not even Mary herself, we are collectively ignorant of the number of Mary's breaths yesterday. This example of collective ignorance could well *not* have occurred—not, for example, when Mary kept the count of her breaths. This example of collective ignorance is therefore, in one sense of that word, accidental: *as a matter of fact* we collectively don't know the number of breaths Mary drew, but *as a matter of principle* we could have known it. The truth about that number isn't known, but it sure is knowable. This paper is about collective human ignorance of what isn't, in fact, known and what may not even be knowable to us at all. This paper aims to describe and explain this particularly deep form of collective ignorance—collective *propositional*, as opposed to acquaintance and ability-based ignorance.

The description and explanation are couched in information theoretic terms, which will be explained in the next section. We say up front that the explanation offered is partial in two respects. First, it only explains collective *empirical* ignorance, as opposed to ignorance of necessary truths. Second, it does so only in terms of human finiteness, whereas a full explanation should also take human fallibility into account (more on this in Sect. 1). The explanation itself has three parts: we are ignorant because we are insensitive to information even in dimensions of information we are familiar with (Sect. 2); because there may be dimensions of information that are inaccessible to us (Sect. 3); and because it cannot be established that our faculties process the information they do pick up adequately (Sect. 4).

We now turn to preliminaries.

⁶ Zimmerman (2016).

⁷ Alexander (2016).

⁸ Smith (2016).

⁹ Explanations of individual ignorance, so of a person's ignorance about X , will often consist in pointing out that (1) an individual person was at the wrong place at the wrong time so as to miss X , and/or (2) that there was no testimonial chain that led from X to her, and/or (3) that she was not free from bias. Some forms of collective ignorance, albeit not the one we are focusing on, can be explained by analogue of (1) and (3), namely when all humans were at the wrong place at the wrong time, or when all are in the thralls of a bias. But it cannot be explained by an analogue of (2), as true testimony must start with someone knowing something not on the basis of testimony.

1 Preliminaries: fallibility, finitude and information

Human beings are ‘ignorant knowers’ in the sense that while there are lots of propositions we know and lots of propositions we don’t but can know, there are also lots of propositions we won’t and cannot know. Collective human ignorance should be explained by reference to two distinct facts: *human fallibility* and *human finiteness*. In what follows, we’ll explain what we mean by ‘human fallibility’ and ‘human finiteness’.

First, to be a fallible knower is to have beliefs that are possibly false. Fallibility here is not understood as a property of knowing but a property of believing and its relation to our capacity to have justification or reasons for our beliefs: you can know that p —and in this way not possibly be wrong about p —even if the reasons for which you believe that p don’t guarantee that your belief is true.¹⁰ Of course, there can be justifiers which guarantee that your corresponding beliefs are true, as is the case with factive states or events such as seeing that p or proving that p . Fallibility here is thus a property of our intellectual capacities and their relation to our beliefs. For example, perception as well as, say, mathematical reasoning can yield false beliefs, even though, when all goes well, they can also yield justification which excludes the possibility of false belief.

Second, to be a fallible knower is not the same as a *finite knower*. Being a finite knower is to have *limited* knowledge of things—where knowledge is limited provided one knows some but not all truths about an object, event, or state of affairs.¹¹ For example, you have limited knowledge of the number 2 if you know that it’s even but not that it is prime. You have limited knowledge about your computer if you know that it’s a Mac but not whether it’s MacBook Air or a MacBook Pro. Human cognitive finitude has to do with the sort of beings we are, and so with the sort of faculties we have: roughly, how they work, and what they are capable of. Finite knowers could, in principle, be infallible ones: when all the beliefs they have are true, or when every reason they have for every belief they have is factive, while at the same time there are truths they don’t know or perhaps cannot know.

We said that a finite knower is someone who has ‘limited knowledge of things’. There are seemingly trivial cases of this, such as when the number of truths expressible in a system (say, set theory) exceeds the number of truths knowable by any one person or even humanity as a whole. A purely *quantitative account* of epistemic finiteness says that the *number* of truths there are exceeds the *number* of truths we can come to know. Plausibly, this will be the case whenever there is an infinite amount of truths in the domain. To be epistemically finite, on this account, just *is* to be in the condition whereby the number of truths there are exceeds the number of truths we can come to know.

One might worry that this quantitative account of epistemic finiteness is false because we cannot really make sense of how propositions, and therefore true propo-

¹⁰ Fallibility finds its expression in biases as well as in cognitive malfunction. Gilovich (1991): part 1 and 2 discusses a number of ‘determinants’ of questionable beliefs, and Kahneman (2011): part 2 is a discussion of a number of ‘heuristics and biases’. Sacks (1985) is a classic presentation of a various form of malfunction.

¹¹ This leads the way to an account of what *knowing more* is. See Treaner (2013).

sitions, can be counted. To see this, consider the following example from Treanor (2013):

(1) grain of sand x is so-and-so-many millimeters from grain of sand y

This certainly looks like an atomic sentence, a sentence which contains no logical connectives. Intuitively, we'd say that (1) expresses a single truth: a truth about how many millimeters x is from y . Contrast it with the following sentence:

(2) grain of sand x is so-and-so-many millimeters from grain of sand y and x is heavier than heavier than y .

This certainly looks like a compound sentence, a sentence which contains a logical connective. Intuitively, we'd say that (2) expresses more than one truth (and perhaps only two truths): a truth about how many millimeters x is from y *and* a truth about how x is heavier than y . As Treanor argues, however, matters aren't so clear with (1):

On the most natural way to read what that sentence expresses, the same proposition could have been expressed by the distinct sentence ' x is a grain of sand and y is a grain of sand and x is so-and-so-many millimetres from y '. But that makes it three truths, not one (Treanor 2013, p. 5).

In that case, (1) is really a compound of three truths. Or is it? Treanor argues that matters still aren't so clear:

For instance, doesn't ' x is a grain of sand and y is a grain of sand and x is so-and-so-many millimetres from y ' really mean ' x exists and y exists and x is a grain of sand and y is a grain of sand and x is so-and-so-many millimetres from y '? But that would be five truths, not three (ibid).

What Treanor's examples highlight is that we cannot count the number of truths we have by counting the number of sentences used to express those truths. Nevertheless, we don't think Treanor's arguments problematize the quantitative account of epistemic finiteness so understood. For it is consistent with that view that "you cannot measure how much truth you have by counting the number of sentences you use to express it" as we did with examples (1) and (2) (Treanor 2013, p. 7). Relatedly, we cannot measure how *finite* we are in some domain by measuring how much truth we have—compared to how much truth there is to have—by counting the number of sentences we use to express the truths we do have. But this doesn't support the idea that we cannot measure how finite we are in some domain by measuring how much truth we have against how much truth there is to have. What it tells us is that one way of making this measure—by *counting* the sentences which express those truths and mapping them one-to-one on to the true propositions we know or believe—will yield unreliable results. That's a worry for how we go about finding out just *how finite we are* and not worry for the quantitative account of epistemic finiteness.

We will explain *how* we are ignorant knowers in information theoretic terms. Following Fred Dretske we think of 'information', or semantic content, as something that can be 'carried' by signals. Consider one of Dretske's examples:

there are eight mischievous boys and a missing cookie. Who took it? An inspection reveals crumbs on Junior’s lips. ... Any one of the eight boys could have taken the cookie, and any one was as likely to take it as any other one. Hence, Junior’s eating the cookie represents a reduction of eight possibilities to one. ... With some routine assumptions we may suppose that the crumbs on Junior’s lips carry information about who ate the cookie.¹²

The crumbs on the lips are a signal that carries the information about Junior that he took the cookie. Another example is the red spots that Agnes developed on her cheeks which are signals that carry the information about Agnes that she is having measles. A final example: in certain conditions two knocks on the door is a signal—and the information it carries is that the coast is clear.¹³

“Information”, says Dretske, “is a commodity that, given the right recipient, is capable of yielding knowledge”,¹⁴ or, a bit more elaborated: “What information a signal carries, is what it is capable of ‘telling’ us *truly*, about another state of affairs. Roughly speaking, information is that commodity capable of yielding knowledge, and what information a signal carries is what we can learn from it”.¹⁵ That information is “capable” of yielding knowledge means that the information that is contained in a signal does not depend for its *being information* for a recipient of the signal that the recipient *actually* learns something from the signal; it is possible that a recipient of the signal is unable to decode or interpret the sign.

Generally speaking, one and the same signal may carry a number of different informational contents. Dretske illustrates the point as follows: “The acoustic signal that tells us someone is at our door carries not only the information that someone is at our door but also the information that the button is depressed, that electricity is flowing through the doorbell circuit, that the clapper in the doorbell is vibrating, and much else besides. Indeed, all the information that is nested (analytically or nomically) in these states of affairs is also part of the acoustic signal’s informational content.”¹⁶

So Dretske says that *signals* are the carriers of information, and that, given the right recipient, they are capable of yielding knowledge.¹⁷ From this, it doesn’t immediately follow that a deprivation of information yields ignorance unless it is only by having access to information at some time that are you in a position to gain knowledge at that time. So, we’ll be presupposing in this paper that that premise is true. From this it follows that we are ignorant when one of the following scenarios obtains:

[BAD INTERPRETATION] There are information carrying signals that we pick up but interpret wrongly;

[INABILITY] There are information carrying signals that we are unable to pick up;

¹² Dretske (1981, pp. 21–22).

¹³ Dretske (1981, p. 40).

¹⁴ Dretske (1981, p. 47).

¹⁵ Dretske (1981, p. 44).

¹⁶ Dretske (1981, p. 72).

¹⁷ As we shall say: the world *contains* or *embodies* information that signals only *carry*.

[**SIGNAL FAILURE**] There is information but no signals to carry it and hence no signals capable of being “picked up” by us.

What these scenarios have in common is that in each of them there *is* information—but it is information that isn’t picked up by us, thereby leaving us ignorant.

2 Finite because of Information Insensitivity even in familiar dimensions

Empirical knowledge is knowledge of contingent features of the world. Such knowledge results, ultimately,¹⁸ from the operation of our sense organs, i.e. from what are traditionally called the five senses, as well as from proprioception, thermoception, nociception and a number of other senses. Generalizing, we can say that our senses pick up information that subsequently occasions belief in us—belief that often qualifies as knowledge. If this characterization of the acquisition of empirical knowledge is adequate, it enables us to note three ways in which we must be deemed finite knowers. These ways, in turn, suggest three parts of the explanation of our epistemic finitude and ignorance. This section discusses the first part.

Without a doubt the world contains more information than our unaided senses have collectively picked up so far. But more interesting and more relevant for our concern here is that it is likely that we are ignorant in a principled way too: it is likely that there is information that we, collectively, *cannot* pick up by our unaided senses. The argument for this draws on a comparison with animal cognition.

Consider first the well-known fact that dogs and whales are able to hear sounds that we cannot hear; that elephants are able to sense seismic vibrations that we cannot sense; that bears can smell smells that we cannot smell; that eagles can see things that we cannot see. The world contains information that their senses can directly pick up but ours cannot.¹⁹ This entails that the world contains more information than our unaided senses are directly sensitive to. Call this *sensory information insensitivity*.

What follows from sensory information insensitivity? Notice that it doesn’t yet establish that our ignorance in certain dimensions is a matter of principle. For although we sometimes are not directly sensitive to certain information, we may be *indirectly* sensitive to it. For instance, humans cannot normally hear sounds outside the 20 Hz–20 kHz range—yet we know that there are sounds outside that range. For we know that some animals hear sounds outside that range. And we know that because sound waves can be visually represented on a screen; *and* we notice that some animals respond behaviorally to waves outside that range. From which we infer that they *hear sounds*. This, then, is a case in which we have *indirect* access to acoustic information. The access is indirect because it is through vision of the amplitudes on a screen that we have access to the acoustic realm. We have *indirect* access to information that is

¹⁸ Ultimately, for although testimony may transfer knowledge between persons, it is not an original source of empirical knowledge.

¹⁹ For these and many other examples of ways in which some animals cognitively out-compete most human beings, and humans out-compete some animals, see Lund (2002), and Wynne and Udell (2013). The latter explicitly describe and discuss how we can investigate the perceptual worlds of animals. (pp. 17–37).

carried by acoustic signals. The acoustic signals that humans cannot notice are “translated” into visual signals that they can directly notice (on a sound meter representing visual graphs). There are many cases of signal-translation. Vibrations of the earth’s surface that we cannot directly feel can be translated into signals that we can directly see (on a seismograph). Fluctuations in temperature that we cannot directly sense can be translated into signals that we can directly see (on a thermometer). The general structure of this kind of “translation” is this: signal S to which Sense-1 of an organism *is not* directly sensitive is translated into signal S^* to which Sense-2 of the organism *is* directly sensitive. In that case, we say, the organism is *indirectly sensitive* to signal S (and to the information it contains).

So far we are supposing that the information we are discussing comes to us through signals of a kind (or in a dimension) that we know one of our senses is sensitive to. Examples of kinds or dimensions of signals include: perceptual signals, auditory signals, and signals in all the sense modalities. The question we should now like to consider is whether we know, or have solid reasons to think that the following translation principle is true:

Translation Principle 1

If there is a kind of signals (or signals in a dimension) to which one of our senses is directly sensitive, then the following is true: if that dimension also contains signals that fall outside the sensitivity range of that sense, then there is always another sense within the sensitivity range of which the translated signal falls (and through which we thus have *indirect* access to the information).

One reason to think this is false is that even when we avail ourselves of technical instruments by means of which we can “expand the scope of perception”, the outputs of the instruments will always have to be items that, in the end, our unaided senses must be sensitive to. If certain sounds, that cannot be heard by us, can be translated in signals that lay in the dimension of the visible, but outside the range in which we are sensitive to them, then this translation is of no avail to us. And if certain colors in the infrared spectrum, that cannot be seen by us without aid, are made visible to us through some technical device, then the output of the device must be something to which our senses are sensitive. Even if there are translations that can give us indirect access to signals, this doesn’t mean that every signal that falls outside the range of one of our unaided senses can be trusted to be translatable into signals that fall inside the range of what one of the other senses is sensitive to. The point, then, is that we just don’t know whether or not Translation Principle 1 is true, and that there doesn’t seem to be any clear reason why we should expect it to be true. But if that is correct, it means that we should not assume that for all signals that cannot be directly sensed, there are translations so that the signals can be indirectly sensed. Hence we are ignorant in a very specific sense: we don’t know whether the signals in the dimensions in which our senses are sensible to them, either directly or indirectly (via translations), are *all* the signals in that dimension. We don’t really know whether all we know in realms in which we know we can know, is all there is to know in that realm. We may be missing out on important truths.

The basic argument, then, is this: there are types of information for which we are directly insensitive, and we lack reasons to believe that for every type of information

for which we are directly insensitive, we could be made sensitive to it indirectly. Of course, *if* we had reasons to believe that the translation principle is true, then we *would* we have reasons to believe that we could be made sensitive to the aforementioned type of information indirectly. But so far we don't have much if any reason to believe that such a translation principle is true. Hence, we lack reason to believe that we are sensitive to all of the available types of information.

This, then, is a first way in which are ignorant: we have no reason to suppose that our senses pick up (directly or indirectly) *all* the information in the dimensions in which they are capable of picking up information. This describes and explains why we are finite, and hence ignorant.

3 Finite because of possibly inaccessible information dimensions

Moving on to the second part of our description and explanation of the sort of collective ignorance that we have our eyes on, we note that there may very well be information that is embodied in signals to which our senses are entirely insensitive. So it is possible that there are information carrying signals for which we have no faculties to pick them up. Hence there may be information to which we have no access at all, neither direct nor indirect access. We offer two arguments to make this claim plausible.

First, there is an inductive argument for the conclusion that such inaccessible dimensions exist. It starts out by noting that *we know* that some animals can pick up more signals in certain dimensions (in which humans can also pick up some signals) than humans can. Next it notes that *we also know* many animals can pick up less signals in certain dimensions (in which we can also pick up some signals) than we can. And it notes furthermore that *it is very plausible to think* that there are organisms that pick up *no* signals in dimensions in which we, and many higher animals, *are* able to pick up signals. If we try, imaginatively, to “place ourselves in the shoes” of these organisms (ants, flies, amoeba), then we must say that *to them* certain dimensions of information are entirely closed off: they cannot pick up *any* information in those dimensions. From this we construct the following inductive argument:

(P1) Organisms O1 pick up more information carrying signals in dimension D1 than we do.

(P2) Organisms O2 pick up less signals in dimension D2 than we do.

(P3) Organisms O3 pick up more signals in dimension D3 than organisms O4.

(P4) Organisms O5 pick up less signals in dimension D4 than organisms O6.

(P5) Organisms O7 pick up no signals in dimension D5, while we do.

.

.

(C) There are information carrying signals to which we are insensitive.²⁰

²⁰ It may well be that a class of organisms referred to in one of the premises, is the same class as a class that is referred to in one of the other premises—even if they bear different name tags; it may also be that a dimensions of information referred to in one of the premises is the same as one that is referred to in one of the other premises—even if they bear different name tags.

We claim three things for this argument. First, that its premises are known to be true; the examples of animal cognition in the previous section should be taken as gestures toward a much fuller set of data that support the premises of this argument. Second we claim that the conclusion follows inductively from the premises: it is made plausible by them. Third, it is impossible to show that the conclusion is false.²¹ One might doubt the correctness of the third claim by referring to the history of science. After all, one might think, information that once was directly inaccessible to us, became indirectly accessible to us via technological innovative instruments. This has happened with cells, with magnetism, with various sorts of radiation, with germs etc. We note, however, that this does not undermine our claim. For whatever becomes indirectly accessible to us, must in the end always become accessible to us via the senses and faculties that we are stuck with. However, as we argued in the previous section, there is no reason to think the Translation Principle is correct: we have no reason to think that all the information there is, is embodied in signals that we can pick up either directly or indirectly.

Finally, it is important to highlight that our inductive argument is not a skeptical argument, or at least, not a *radical skeptical* argument: it is not an argument for the conclusion that we don't know anything about the world at all, or that none of our beliefs about the world is justified. At best, it motivates a kind of intellectual humility, specifically, the recognition that we don't know how far our ignorance stretches.²² Since we have good reason to believe that there are information carrying signals to which we are insensitive, we don't know how much information we are constitutionally deprived of, and thus how ignorant we are. And rather than supporting radical skepticism, you might think this supports a kind of *epistemic anxiety* or *angst*. After all, if you have a desire to know something—say, to know the solutions to all the major unsolved mathematical problems—but you learn that you and everyone else is constitutionally unable to know all of these solutions, you might become anxious and develop angst in response to this (putative) epistemic limitation.²³ We are not claiming here that humanity ought to be anxious and respond to our epistemic finitude with angst. Rather, we are only flagging that recognition of our epistemic finitude in this regard can be at least as much of a reason to develop intellectual humility as it is a reason to develop epistemic angst.

Our second argument is a conceivability argument. Over the centuries many people have claimed that at least *some* people have exceptional faculties that enable them to pick up information from dimensions that are closed off to virtually all others. Perhaps those who are so closed are not *essentially* closed off from them. At least, sometimes it is claimed that a process of mental growth, or education, or a regime of training

²¹ That there are dimensions of information to which we are entirely insensitive cannot be shown to be false for the same sorts of reasons that the claim that there is design in the natural world cannot be shown to be false. A discussion of some of those reasons is Van Woudenberg and De Ridder (2017).

²² This has affinities with what Church and Barrett 2016 call the 'limitations-owning' account of intellectual humility. In this case, we recognize and thereby own the fact that our intellectual capacities are unable to directly or indirectly pick up all of the information carrying signals, thereby leaving us ignorant with respect to objects, properties, or states-of-affairs to which that information is a part.

²³ Cf. Pritchard (2015) on epistemic angst. Notice that for Pritchard, the source of epistemic angst is also a fundamental type of epistemic limitation, but it has to do with recognizing the rational 'groundlessness' of our most fundamental commitments, such as that there is a physical world, or that our senses are reliable. Translating his position into our vocabulary of epistemic finitude, we are in part epistemically finite because we are constitutionally deprived of rationally assessing our fundamental commitments.

and discipline might eventually “awaken” or “reap” the faculty at hand. Below we will provide some examples. Before doing that, however, we need to be clear about the following. Some claims in this area are claims to the effect that some people have non-standard faculties that enable them to pick up information in dimensions in which most other people are *also* able to pick up information, but by standard faculties. Other claims in this area, however, are claims to the effect that some people have faculties that enable them to pick up information in dimensions to which standard faculties have *no* access. Let us call claims of the first type “checkable claims”, and the faculties that are claimed to exist in some persons “checkable faculties”. Let us call claims of the second type “non-checkable claims”, and the faculties claimed to exist in some persons “non-checkable” faculties.²⁴

One claim to an allegedly non-standard but checkable faculty, is the ability to know what some other person thinks or feels, without the other person using any form of language or any other standard mode of communication. Such claims to telepathy are checkable, as there are standard ways to know what someone else thinks or feels, viz. linguistic communication by that person. Another claim to a non-standard checkable faculty is the ability to see persons and events distant in space and/or time. Such claims to clairvoyance are, in principle, checkable, as there are standard ways to know the whereabouts of persons (their own testimonies, historical records), and there are also standard ways to know about past events (historical records, personal recollections), and even future events (wait and see). *In practice*, claims to clairvoyance will often be un-checkable, as there will be no personal testimonies, no historical records available of the events clairvoyed. But *in principle* a check is possible.

One claim to an alleged non-standard non-checkable faculty is a claim to ‘spiritual eyes’ that, if appropriately cultivated, give access to presumed higher worlds to which our standard senses don’t give us access.²⁵ Another example is the claim to an ability to communicate with the dead. The faculty that supposedly enables their possessors to know what the dead say or think, is non-checkable too, as no standard faculty is able to penetrate into the realm of the dead—if there is such a realm.

We have no interest whatsoever in even mildly suggesting that these alleged un-checkable faculties exist in some persons and do indeed pick up information in a dimension in which our standard faculties don’t. The point of bringing up these alleged faculties is, rather, that we lack faculties that we could *conceivably* have. That is to say, the point is that the following seems true: there are possible worlds in which we exist and in which we possess, in addition to the standard faculties that we possess in the actual world, also non-standard faculties like the ones mentioned. If non-standard faculties are not inherently impossible, it is possible to have them. That is, if they are possible, there is a possible world in which we *do* have them. This very possibility, we aver, should make us think that it is possible that *in the actual world* there are

²⁴ Claims in both categories have been made, among others, by members of the famous, or infamous, Society for Psychical Research, founded in 1882, which counted among its members the philosophers Sidgwick (1838–1900), James (1842–1910), Price (1871–1947) and Broad (1887–1971), all of whom have served as presidents of the Society. See Gauld (1968) and Haynes (1982).

²⁵ Such claims have been issued by, for example, Emanuel Swedenborg, *Arcana Coelestia* (1749–1756), and Rudolph Steiner, *Wie Erlangt man Erkenntnisse der höheren Welten* (1904). An English translation of the latter work is Steiner (1947).

dimensions of information from which we are unable to pick up any information, due to our lacking the relevant faculty.

These, then, are our two arguments for thinking that there can be dimensions from which we simply cannot pick up any information whatsoever: an inductive argument, and a conceivability argument. This is the second part of our description and explanation of collective ignorance.

These conclusions also undercut the idea that our best, future scientific theories—based on all the available information we could ever gather—will be fundamentally correct. After all, if there are dimensions of information that are fully inaccessible to us, then there is no reason to be assured that a mature scientific description of the world—one that, to speak with Peirce, human inquirers “ultimately agree on”²⁶—would be fundamentally correct,²⁷ for it may be *fundamentally incomplete*. To see this, consider the following principle:

Incompleteness Principle 1:

If a description D of the world is fundamentally incomplete, then we cannot be assured that it is fundamentally correct.

It is important not to confuse this with

Incompleteness Principle 2:

If a description D of the world is fundamentally incomplete, then it is fundamentally incorrect.

By way of illustration, suppose we have an incomplete description of what is in fact a car’s engine: we have a description of its battery, its spark plugs, and its fan. From this it doesn’t follow that *what is* described, viz. battery, spark plugs and fan, is fundamentally incorrectly described. However, since the description at hand is not a description *of a car’s engine*, it *does* follow that since the description is fundamentally incomplete *we cannot be assured* that it is fundamentally correct. A fundamentally correct description would be a description that makes it clear that what is described is *a car’s engine*. In general: if there is a THING composed of parts t, h, i, n and g, then if we describe t, h, and g, then we have an incomplete description of the THING. And since it is incomplete, we cannot be assured that it is fundamentally correct.

Again, this is by no means radical skepticism. It is a form of higher-order skepticism: we don’t know whether what we know concerns all that there is to be known. Our knowledge may very well be fundamentally incomplete. But if it is, then, we cannot be assured that it is fundamentally correct.

4 From contingency to epistemic finiteness

In this section, we’ll argue that the contingent relations between our informational inputs and outputs is contingent and that this contingency provides *prima facie* reason to think that we are epistemically finite.

²⁶ Peirce (1992 [1878], pp. 139).

²⁷ Rosenberg (2011, p. 6) claims that what current science tells us is fundamentally correct. A discussion of Rosenberg’s scientism is Van Woudenberg (2018).

Our argument begins with the following observation: We pick up lots of information through our perceptual faculties and our cognitive capacities. Characterized at a very high level of abstraction, we can say that these mechanisms (a) “work on” (b) an “input”, and occasion (c) an “output”. It has been maintained that the relation between the “inputs” and the “outputs” of our cognitive mechanisms is *contingent*. Thomas Reid, for example, has it that

“No man can give a reason, why the vibration of a body might not have given the sensation of smelling, and the effluvia of bodies affected our hearing, if it had so pleased our Maker. In like manner, no man can give a reason, why the sensation of smell, taste, or sound, might not have indicated hardness, as well as that sensation, which, by our constitution, does indicate it. Indeed no man can conceive any sensation to resemble any known quality of bodies. Nor can any man show, by any good argument, that all our sensations might not have been as they are, though no body, nor quality of body, had ever existed.”²⁸

Alvin Plantinga strikes a similar chord when he says that “the connections between experience and belief are multiply contingent”²⁹ as all of the following scenario’s are possible: (i) one may have a sensuous experience, and form no perceptual belief at all—as may happen when one is just enjoying the play of one’s sensuous experience; (ii) one may have a sensuous experience, and *not* form belief B that one normally forms, but belief B* instead—as is the case when one looks at one’s backyard and *not* form the belief that one normally forms, viz. that the tiger lilies are now blooming, but instead the belief that there are tigers in one’s backyard, or even the belief that Goldbach’s conjecture is true; (iii) one may have something like a perception, without one’s being appeared to in any sensuous way—as is testified by the phenomenon of “blindsight”.³⁰

We now suggest that the contingency of the relation between the inputs and outputs of our faculties betokens epistemic finitude and collective ignorance, as follows. We know that our cognitive mechanisms “do” something with their “inputs”, “work” on them, “process” them in ways for which there are alternatives. Take vision, for instance. The information that reaches the retina is much more abundant than can be transmitted to the brain. A lot of information is lost along the way.³¹ Losing information is one thing that the faculty of vision ‘does’ with its inputs. Similar things happen to our other faculties. This betokens our epistemic finitude.

We note in passing that the fact that the relationship between the inputs and outputs of our faculties is contingent is *another* fact than the fact, assuming it is one, that there are information carrying signals that we pick up but interpret wrongly (what

²⁸ Reid (1764 [1997], p. 57). For extensive discussion of passages such as these, see Nichols (2007).

²⁹ Plantinga (1993, p. 92).

³⁰ A book length study of blind sight is Weiskrantz (1986).

³¹ See Hardin (1993). Note that this is similar to what Johnston (1996) says about the epistemology of color experience: “Barring a pre-established harmony no such causal process will preserve and transmit information so as to secure a nature-revealing match between how some feature of the cause, say the greyness of my dog’s coat, is and the way I am caused to represent that feature as being. To see involves having the natures of visible properties revealed by a causal process, but this is just what no causal process actually does” (Johnston 1996, p. 191).

we have called BAD INTERPRETATION). But the facts are related in the following way: signals are the inputs of our faculties and the faculties produce beliefs in us. But the fact that the relation between inputs and outputs is a contingent one doesn't mean or entail that every interpretation we give of a signal therefore is a bad one. Nothing in the contingency of the input–output relation suggests that every interpretation of signals is wrong, or that no interpretation can be correct. By way of illustration: *a fox's footprints in the snow* carries the information *that a fox was here*. Suppose now the footprints are in your visual field, you notice them. Suppose also that you form the output belief *that a raccoon was here*. Then you have misinterpreted the signal and formed a false belief. But next suppose that you form the belief *that a fox was here*. Then you have formed a true belief and given a correct interpretation of the signal. In both cases the relation between input (the signal, in this case *a fox's footprints in the snow*) and the outputs (the belief that a raccoon was here, and the belief that a fox was here, respectively) is contingent. But this doesn't mean or entail that the interpretations are both bad ones. For the second one isn't a bad one at all! The standard for goodness and badness here is the world—the world as it is independent of interpreters, the world “as it is anyway”, as Bernard Williams put it.

The relationship between the inputs and outputs, to repeat, is a *contingent* one. This is similar to what Hume thought in the *Enquiry* about the relationship between events. As Hume observed:

Suppose that a highly intelligent and thoughtful person were suddenly brought into this world; he would immediately observe one event following another, but that is all he could discover. He wouldn't be able by any reasoning to reach the idea of cause and effect, because (firstly) the particular powers by which all natural operations are performed are never perceived through the senses, and (secondly) there is no reason to conclude that one event causes another merely because it precedes it.³²

According to Hume, our sensory experiences provide us with no reason to believe that there is a necessary connection between one event and another event. We don't *perceive* a necessary connection between the events. This seems intuitively plausible. For it's hard to see how we could perceive *necessity* in the way that we perceive *color* or *shape*. Put another way, when we reflect on our experiences, it doesn't seem to us as if, say, the brittle glass's *having to shatter* after being struck by a rapidly moving ball is a feature of our experience of the event.

Moreover, Hume thought that it's hard to see how we could, by reasoning, get justification to believe that there is a necessary connection between the events. For it *could* be that one event occurs (the rock striking the glass) and then a second event follows it without a causal connection between them (the glass's shattering). “Their occurring together may be arbitrary and casual, with no causal connection between them”, he says. Since we can conceive of one event happening (the rock striking the glass) and then another event just so happening independently (the glass's shattering), this provides *prima facie* reason to think that there is no necessary connection between the two events. There is no evident contradiction in conceiving of this.

³² Hume (2008 [1748], pp. 19–20).

We don't want to suggest that some events don't cause other events, or that we don't have any reason to believe that some events cause other events. We are not suggesting here that Hume is right to think that "custom or habit" is the only reason why (or the only reason we have) to make the causal judgments that we do. What we think is that Hume's view about our deprivation of information about necessary connections between external events is true of the relationship between *external inputs* and our *cognitive outputs*. It could be that square buildings—external inputs—produce triangular visual experiences in creatures with our capacities—cognitive outputs—or that what is really fluid produces tactile experiences of solidity in creatures like us, and so on. There is no contradiction in conceiving of this.

If that's possible, it makes it hard to see how we could know that our representations, whether cognitive or sensory, *must* accurately represent their distal external causes; that they *have* to represent their causes accurately. It could be that they systematically misrepresent their causes. What we mean here is that it's hard to see *how* we could come to know that:

Higher-order veridicality: For any of our types of representations (cognitive or sensory outputs) E, necessarily, E is generally veridical.

This is not to be confused with two other ideas, namely:

First-order inaccuracy: For any of our types of representations (cognitive or sensory outputs) E, necessarily, E is generally unveridical.

First-order ignorance: For any of our types of representations (cognitive or sensory outputs) E, we cannot know, of some particular instance E_1 of the type E, that E_1 is veridical.

We think we are *finite* with respect to our knowledge of the accuracy of our representations. That is, we endorse:

Higher-order finiteness: We are not in a position to know that, for any of our types of representations (cognitive or sensory outputs) E, that it is generally veridical. We are epistemically finite with respect to their external inputs.

Our argument for higher-order finiteness deploys two premises. The first premise is that the only way we currently have of getting information about external inputs is by having certain representations (outputs) which accurately represents those inputs, such as by having sensory experiences, or by making inferences from those experiences. Second, it's possible that the way the external inputs are—what properties they have, or what they are fundamentally like—is very different from the way we represent them as being. Again, to take one example, it could be that objects which are essentially square typically produce sensory representations in human beings as of something which is triangular.³³

³³ This seems to be at odds with Davidson's view about the general veridicality of belief, one type of representation. He says that: "If anything is systematically causing certain experiences (or verbal responses), that is what the thoughts and utterances are about. This rules out systematic error. If nothing is systematically causing the experiences, there is no content to be mistaken about" (Davidson 1991, p. 195). We take it that we can generally identify the content of our representations successfully even if we are not in a position to know that the relevant type of representation, as a class, is generally veridical.

Now imagine that we are seeing what is in fact a square object O and we have a corresponding visual experience, E , which represents O in a particular way, V . Perhaps it represents O as being a V (e.g. it represents the square object as being square), and thus it represents O accurately. In this way, we are in a position to know that O is V . First-order knowledge is within our reach. But can we know that E accurately represents O when E represents O as V ? That is, can we acquire higher-order knowledge of the accuracy of our representations? It seems as if this is possible, as per our first premise, only by getting information about this representation. So, imagine that you have some further experience, E^* , which indicates to you that E accurately represents O as V . E^* would then take the experience E as its object O^* and represent O^* in a particular way, namely, as being accurate with respect to O . Call this property V^* . This raises the question of whether E^* , by representing O^* as V^* , is accurate. Are you in a position to know that E^* , by representing O^* as V^* , is accurate? And so on. In short, it seems as if to get knowledge about whether our representations are generally veridical, we'd need to have some *non-representational* way of knowing about them. Otherwise there will always be a further question of whether the particular higher-order representation we have, which takes the relevant first-order representations or a class of them as its object, is veridical.

Here's a worry for our view. For all that we have said thus far, you can still get *circular* epistemic reasons for believing that our representations generally accurately represent external reality and that such reasons can be *good reasons*.³⁴ We doubt this. To see this, consider the following case:

Cave. Joey has lived his whole life in a cave with its exit blocked. He has no outside information about his life in the cave (or that it's even a cave as such: for Joey the walls of the cave are simply the boundary of his world). Nevertheless within the cave are screens with images on them which in fact represent their distal causes accurately. Joey looks at the screen and 'image A', which is an accurate representation of a tree, has an outside tree as its cause; 'image B', which is an accurate representation of a river, has the river outside as its cause, and so on, for nearly all the 'images' he names. Joey has never left this cave and can never leave this cave. He believes that image A (of the tree) and image B (of the river) accurately represent what the images seem to present.

We are bracketing the question of whether Joey can know that there are trees and rivers on the basis of appealing to the images on his screen. What we are interested in here is whether he is in a position to know the following kinds of claims: 'that image A accurately represents x ', or 'that image B accurately represents y ', where x is the tree and y is the river. It seems as if he would need some evidence to believe those images (outputs) aren't likely to be inaccurate representations of their causes (inputs) in order to know that they are accurate representations. After all, a closure principle supports this much:

³⁴ See Alston (1993). See Sosa (2009) (chapters 8–9) for a defense of the view that epistemic circularity doesn't prevent higher-order knowledge.

Closure: If S knows that P, and competently deduces Q from P, whilst retaining their belief that P throughout, S is thereby in a position to know that Q.³⁵

Suppose that Joey can know that if what he identifies as image A accurately represents x (he calls it ‘the tree’), then that output is not caused by some other input, say y (what he calls ‘the river’). So, if Joey knows that image A (the target output) accurately represents x (the alleged input), then he knows that image A isn’t caused by some other entity. But surely he has no reason to believe image A isn’t caused by some other entity. For all he can do is form his belief about the relationship between the alleged inputs and the manifest outputs based on what he sees depicted on the screens—the outputs—he has within the cave. For all he knows, the images on the cave have systemically *different* inputs, disconnected from what the outputs depict. And, as we argued, such a situation is possible. This echoes Hume on the causes of our sensory perceptions:

By what argument can it be proved that the perceptions of the mind must be caused by external objects that are perfectly distinct from them and yet similar to them (if that were possible), rather than arising from the energy of the mind itself, or from the activities of some invisible and unknown spirit, *or from some other cause still more unknown to us?* (Our emphasis).³⁶

All of this belongs to the third part of the explanation *why* we are ignorant knowers: we cannot know that the “outputs” that our faculties generate “match up” with the way the world is; and we cannot know this essentially because the relation between “inputs” and “outputs” is contingent, and we cannot avoid using the very faculties whose reliability we want to establish. This is part of our epistemic finiteness, and explains our ignorance.³⁷

5 Two thought experiments that support collective ignorance

In this section we present what we think of as an interesting augmentation of what we have been arguing so far, especially in Sect. 2. To get started, consider the following two thought experiments.

One. Suppose you slowly walk up to what is in fact William Holman Hunt’s painting *The Light of the World*. From the distance where you start walking you cannot see that it is that particular Hunt painting you are approaching, you cannot even see it is a painting. All you see is something far in front of you. As you progress, however, you see clearer and clearer what it is that you are walking up to: something on a wall—something *hanging* on a wall—something *darkly colored* hanging on a wall—a darkly colored

³⁵ Cf. Pritchard (2015) and Williamson (2000).

³⁶ Hume (2008 [1748], p. 79).

³⁷ None of this entails that we don’t know things. But it is a powerful reason to doubt what has been called the KK-principle—the principle that says that if one knows that p , one knows that one knows that p . There are many things that we know through the operation of our faculties. But this doesn’t entail that we know this; it doesn’t entail that we know that the faculties we are endowed with adequately process the information that they pick up. There is disagreement about the KK-principle; see Rescher (2005, p. 22) and Hintikka (1962, pp. 103–106).

painting hanging on a wall—etc. and within a certain range of distances, you have *The Light of the World* in full view, you see the whole of it. Progressing just a bit further you see more clearly a number of details that escaped your notice when you looked at the whole of the painting. Not stopping there, you get nearer to the surface of the painting bit by bit. The effect is that you no longer see the painting as a whole, and when your eyeballs almost touch the canvas, you no longer see that it is a painting. You only see dark colors, and when your eyes almost touch the canvas you don't even see that anymore.

Two.³⁸ Think of the captivating melody of Sibelius' *Finlandia*. Its performance in a music hall usually takes some nine minutes. Now think of the same piece as being performed at a slower tempo—so that its performance now takes 18 min. In fact think of it as being performed at ever slower tempi—as taking a full hour, or even a full day, a full week, a full month, a full year, a full decade, a full century. Now ask yourself: what would you hear when you hear the ever slower performances of Sibelius' masterpiece? In the initial stages of the progression you would still hear, and recognize what you hear as a performance of, *Finlandia*. But as the progression continues, it is likely that in some interval you no longer recognize this, although you may still recognize it as music, and later, perhaps, as “organized sound”. But at some further point you even no longer do that—then it is just noise. Similar things “in the other direction” (so to speak) happen when the tempi are sped up.

Hunt's painting and Sibelius' composition (or the various performances thereof) are, as common sense metaphysicians think of it, “wholes” or “units”, as opposed to jumbles of paint, or jumbles of sound. They have “individuality”, their parts compose something, as opposed to being a jumble of things that compose *nothing*—as a tennis ball, a Boeing 747, and the British Prime Minister jointly compose *nothing*, not something that is a unity.³⁹

Common sense strongly suggests that wholes, things with unity, exist. The examples it can refer to abound: molecules, cells, organisms, artifacts like books, computers, cars, Boeing 747 s, and, of course, paintings and (performances of) musical compositions. These wholes embody “information” of their unity in the sense in which I have been using that notion all along. Now, intuitively speaking, a whole embodies more information than the sum of its constituting parts. If a, b and c compose a whole or a unity (to be represented as $a + b + c$) then $a + b + c$ carries more information than the sum of a, b, and c contain separately.

What the two experiments suggest is that in order for us to be able to pick up all the information that is needed to perceive wholes or “unities” we need to be at a right distance range to them. In the first experiment the relevant distance is *spatial* distance and orientation. Standing too far too near to Hunt's painting as well as upside down or at the wrong angle prevents one from seeing *The Light of the World*. In these cases seeing the painting doesn't enable one to pick up the amount of information

³⁸ This thought experiment is inspired by the performance of John Cage's composition ORGAN2/ASLSP (“As slow as possible”) in Halberstadt, Germany. The performance started on September 5, 2001, and is scheduled to end in the year 2640.

³⁹ We are assuming here a common sense metaphysics (as opposed to a revisionary one like Van Inwagen 1990) that provides a natural home for ordinary objects like paintings and musical compositions. A lucid defense of common sense metaphysics is Thomasson (2007).

that is needed for seeing that what one is looking at is a particular whole, a unity. One will pick up *some* information from the area one is looking at, of course, but not all of the information that is needed to see the unitary object that is Hunt's painting. And this suggests the following general idea: we may be surrounded by wholes, things with unity, without us being able to pick up the information that is needed for us to *see* wholes, and to *believe* they are, because we don't have the right spatial distance towards them. We may even go further and say that *it is likely* that we are so surrounded.

The second experiment is a bit more complicated, but suggests the same point. Sibelius' *Finlandia*, or any of its regular performances in a music hall, I assume, is a unity, it is *one* thing (and not many, even though it is composed out of many different sounds), it is a whole. When the tempo of the performance is slowed down, the piece doesn't thereby stop being a unity, or a whole. I take this to be true even when the tempo of its performance is slowed down so as to take a full day, or a full week. However, somewhere down the line we stop hearing *Finlandia*, we even stop hearing music. Somewhere down the line we become incapable of picking up the information required to hear the unitary whole that the *Finlandia* is, even if we do pick up information coming from what is in fact a performance of Sibelius' most famous composition. The problem here is not, as it was in the first experiment, our (spatial) "distance" to the information (at least not primarily), but [i] the temporal "distance" of the parts of the information, and [ii] our temporal "distance" to the parts of the information. Those parts of the information that would have to be "picked up" by us if we were to experience the performance as a whole, are not, and, if the performance takes a century or more, *cannot* be "picked up" by us. For us the parts are temporally too far removed *from each other*, but also from temporally too far removed *from us*—too far for us to be able to pick up the information relevant to experiencing what we hear as a unitary whole.

That information can *temporally* be too far removed from us so as to prevent us from experiencing (what is in fact) a unitary whole, has a *spatial* analogue. There may be spatial unitary wholes, whose parts are so widely spread out over vast stretches of space that we are unable to see them as wholes—the information relevant to experiencing such whole, is not and perhaps *cannot* be "picked up" by us.

All of this, then, suggests that the following is possible, and presumably even likely: we are surrounded by unitary wholes, from whose parts we are able to pick up some information, but not enough so as to experience, or know, them as wholes. And the explanation of this possibility can be given in terms of "distance": the temporal distances between the parts of the whole, the spatial distances between the parts of the whole, the temporal distance of the whole and its parts to us human beings, and the spatial distance of the whole and its parts to us human beings—these distances are too great or too small and thus they prevent us from noticing unitary wholes. These distances prevent us from picking up the information relevant to experiencing wholes and hence from knowing *that* they are wholes.

6 Conclusion

The explanation as to why we are ignorant knowers, then, consists in the conjunction of the following theses: (1) Even in those dimensions where we are capable of picking up information, there is information that we don't pick up. (2) There may very well be dimensions of information for which we lack any capacity to pick up information from. (3) We cannot establish that the faculties we are endowed with process the information that they do pick up reliably. This explanation is not a *full* explanation, as it focusses only on our finiteness and leaves out the other aspect that a full explanation should include, viz. our fallibility.

Thesis (1) was augmented by pointing out the problem of temporal and spatial distance for knowing unified wholes: in order to be able to know unified wholes, neither the distances *between the parts* of the wholes must be too large, nor must *our distance* to the parts. Since we cannot be sure these conditions are satisfied, we cannot be sure that there are no wholes we don't know.⁴⁰

Funding Funding was provided by Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NL) (Grant No. 276-20-024).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Alexander, L. (2016). Ignorance as a legal excuse. In R. Peels (Ed.), *Perspectives on ignorance from moral and social philosophy* (pp. 205–216). London: Routledge.
- Alston, W. P. (1993). *The reliability of sense perception*. Ithaca: Cornell University Press.
- Davidson, D. (1991). Epistemology externalized. *Dialectica*, 45, 191–202.
- Dretske, F. (1981). *Knowledge and the flow of information*. Oxford: Blackwell.
- El Kassar, N. (2018). What ignorance really is examining the foundations of epistemology of ignorance. *Social Epistemology*, 32(5), 300–310.
- Gauld, A. (1968). *The founders of psychical research*. London: Routledge & Kegan Paul.
- Gilovich, T. (1991). *How we know what isn't so. The fallibility of human reason in every day life*. New York: The Free Press.
- Gross, M., & McGoe, L. (Eds.). (2015). *Routledge international handbook of ignorance studies*. London: Routledge.
- Hardin, C. Larry. (1993). *Color for philosophers. Unweaving the Rainbow*. Indianapolis: Hackett.
- Haynes, R. (1982). *The society for psychical research 1882–1982*. London: MacDonald.
- Hintikka, J. (1962). *Knowledge and belief. An introduction to the logic of the two notions*. Ithaca: Cornell University Press.
- Hume, D. (2008) [1748]. *An enquiry concerning human understanding*. In J. Bennett (Ed.). <http://www.earlymoderntexts.com/assets/pdfs/hume1748.pdf>. Accessed Apr 2019.
- Johnston, M. (1996). Is the external world invisible? *Philosophical Issues*, 7, 185–198.
- Kahneman, D. (2011). *Thinking fast and slow*. London: Allan Lane.
- Le Morvan, P., & Peels, R. (2016). The nature of ignorance: Two views. In R. Peels & M. Blaauw (Eds.), *The epistemic dimensions of ignorance* (pp. 12–32). Cambridge: Cambridge University Press.

⁴⁰ For comments and discussion on earlier drafts of this paper we should like to thank Jeroen de Ridder, Rik Peels, Wout Bisschop, Naomi Kloosterboer, Lieke Asma, and Hans van Eyghen. We are also indebted to two anonymous referees of *Synthese* for their very helpful and insightful remarks.

- Lund, N. (2002). *Animal cognition*. New York: Taylor & Francis.
- Mills, C. (2007). White ignorance. In S. Sullivan & N. Tuana (Eds.), *Race and epistemologies of ignorance* (pp. 11–38). New York: State University of New York Press.
- Nichols, R. (2007). *Thomas Reid's theory of perception*. Oxford: Oxford University Press.
- Nottelman, N. (2016). The varieties of ignorance. In R. Peels & M. Blaauw (Eds.), *The epistemic dimensions of ignorance* (pp. 33–56). Cambridge: Cambridge University Press.
- Peirce, C. S. (1992) [1878]. How to make our ideas clear. In N. Houser & C. Kloesel (Eds.), *The essential peirce. Selected philosophical writings* (Vol. I, pp. 124–141). Bloomington: Indiana University Press.
- Plantinga, A. (1993). *Warrant and proper function*. New York: Oxford University Press.
- Pritchard, D. (2015). *Epistemic angst: radical skepticism and the groundlessness of our believing*. Princeton: Princeton University Press.
- Pritchard, D. (2016). Ignorance and epistemic value. In R. Peels & M. Blaauw (Eds.), *The epistemic dimensions of ignorance* (pp. 132–143). Cambridge: Cambridge University Press.
- Reid, T. (1764) [1997]. *An inquiry into the human mind on the principles of common sense*. Edited by D.R. Brookes. Edinburgh: Edinburgh University Press.
- Rescher, N. (2005). *Epistemic logic. A survey of the logic of knowledge*. Pittsburgh: Pittsburgh University Press.
- Rescher, N. (2009). *Ignorance. On the wider implications of defective knowledge*. Pittsburgh: University of Pittsburgh Press.
- Rosenberg, A. (2011). *The atheist guide to reality*. New York: Norton.
- Ross, D., Ladyman, J., & Spurrett, D. (2010). In defence of scientism. In J. Ladyman & D. Ross (Eds.), *Every thing must go. Metaphysics naturalized* (pp. 1–65). Oxford: Oxford University Press.
- Sacks, O. (1985). *The man who mistook his wife for a hat*. London: Duckworth.
- Smith, H. (2016). Tracing cases of culpable ignorance. In R. Peels (Ed.), *Perspectives on ignorance from moral and social philosophy* (pp. 95–119). London: Routledge.
- Sosa, E. (2009). *A virtue epistemology. Apt belief and reflective knowledge* (Vol. II). Oxford: Oxford University Press.
- Steiner, R. (1947). *Knowledge of the higher worlds and its attainment*. Spring Valley, NY: Anthroposophic Press.
- Thomasson, A. L. (2007). *Ordinary objects*. Oxford: Oxford University Press.
- Treanor, N. (2013). The measure of knowledge. *Nous*, 47, 577–601.
- Van Inwagen, P. (1990). *Material beings*. Ithaca: Cornell University Press.
- Van Woudenberg, R. (2009). Ignorance and force. Two excusing conditions for false beliefs. *American Philosophical Quarterly*, 46, 373–386.
- Van Woudenberg, R. (2018). An epistemological critique of scientism. In J. de Ridder, R. Peels, & R. van Woudenberg (Eds.), *Scientism: Prospects and problems* (pp. 167–189). Oxford: Oxford University Press.
- Van Woudenberg, R., & de Ridder, J. (2017). Design hypotheses behave like skeptical hypotheses (or: Why we can't know the falsity of design hypotheses). *International Journal for the Study of Skepticism*, 7, 69–90.
- Weiskrantz, L. (1986). *Blindsight: A case study and implications*. Oxford: Oxford University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Wynne, C. D. L., & Udell, M. (2013). *Animal cognition. Evolution, behavior & cognition*. London: Palgrave Macmillan.
- Zimmerman, M. J. (2016). Ignorance as a moral excuse. In R. Peels (Ed.), *Perspectives on ignorance from moral and social philosophy* (pp. 77–94). London: Routledge.