

# Chapter 7

## Perspectivalism About Knowledge and Error



Nick Treanor

**Abstract** Knowledge and error have a quantitative dimension – we can know more and less, and we can be wrong to a greater or lesser extent. This fact underpins prominent approaches to epistemic normativity, which we can loosely call truth-consequentialist. These approaches face a significant challenge, however, stemming from the observation that some truths seem more epistemically valuable than others. In this paper I trace out this perspectivalist challenge, showing that although it arises from a mistaken picture of the quantitative dimension of knowledge and error, when we reconceive how that quantitative dimension should be understood we find the perspectivalist challenge has survived unscathed.

**Keywords** Veritism · Epistemic normativity · Truth · Perspectivalism · Similarity

### 7.1 Introduction

Consider an omniscient being and a blank slate. One has perfect, complete, immaculate knowledge, the other none at all. Each of us is somewhere in between. Our knowledge grows and increases, both overall and with regard to specific subject matters or domains. How much each of us knows can decrease as well, both overall and about subject matters. Moreover, interpersonal comparisons are possible: you know more now than I knew as a child, and there are topics about which you now know more than me and others about which I now know more than you. All this points to the idea that there is a quantitative dimension to knowledge. Indeed, it seems built into the very concept of knowledge that it is a quantitative notion in the sense that one can have more and less of it. It is not just that one knows or fails to know, or has knowledge or doesn't; it is that one can know more or less, one can have more or less knowledge.

We also differ from the omniscient being and the blank slate in that there is, in all of us, some admixture of error. It is not just that how much we know can increase,

---

N. Treanor (✉)  
University of Edinburgh, Edinburgh, UK  
e-mail: [nick.treanor@ed.ac.uk](mailto:nick.treanor@ed.ac.uk)

diminish, and in principle be compared to how much some other person knows. It is that we can go wrong to a greater or lesser extent. That is again a quantitative notion, more-ness or magnitude in the domain of error.

Finally, it is common to think not only that knowledge and error each have a quantitative dimension, but that there is an intelligible sense to the idea that there is some degree to which a person is, at a time, getting the world right. It's not just that I here, now, know a certain amount and am wrong a certain amount. It is that I have a picture of the world that can in principle be assessed overall with regard to the faithfulness with which I represent the world. How faithfully I represent the world, overall, is some function of how much I get right and how much I get wrong, with increases of knowledge or true belief increasing overall faithfulness, and increases in error decreasing overall faithfulness.

The question I want to explore in this paper is whether there is an interesting sense in which quantities of knowledge and error are perspectival. I will not do the topic justice. I will also not do justice to the interesting work on perspectivalism in science explored by Michela Massimi (2016, 2018a, b) and others and how it bears on the overall question I am here interested in. Nonetheless, I hope there are some points of contact.

This will be a story in four parts, where we visit the issue of perspectivism briefly near the beginning to foreshadow a more sustained discussion near the end. First, I spend some time talking about a picture of epistemic normativity that is, I believe, rampant in contemporary philosophy. I say 'rampant' rather than 'prevailing' or 'standard' to adumbrate the fact this picture is, in my view, a sort of disease, one that has spread quietly and unnoticed. The second part of the story is the explanation or argument concerning why I think this picture of epistemic normativity is mistaken, why the prevalence of this picture is a problem rather than testimony to its soundness. The third part of the story concerns my effort to offer something better, something that preserves the best part of the picture I want to replace while avoiding its mistakes. The fourth and final part of the story focuses on the perspectival challenge, which we shall see survives unscathed through a radical reconception of a prevailing approach to epistemic normativity. What I mean will become clear in due course, but the basic idea is that the picture of epistemic normativity that I think is wrong threatens a kind of perspectivalism about epistemic normativity, *but so too* does the picture that I argue should replace it. I will elaborate and defend both of these points. The aim will not be to conclude with perspectivalism, but to throw down a challenge to my own view and to show how the spectre of perspectivalism haunts truth consequentialist approaches to epistemic normativity. This is the case, we shall see, whether one endorses the traditional interpretation of that approach or the interpretation that I propose has to replace it.

## 7.2 Truth-Conduciveness and Epistemic Normativity

The picture of epistemic normativity I focus on in this paper is one that is truth-consequentialist. There are two aspects to it. The first is a claim about the structure of epistemic normativity, that it is consequentialist. Roughly, the idea is that epistemic processes (or whatever the locus of evaluation is) should be assessed by appeal to how conducive they are to the achievement or existence of some epistemic good. The second adds to this structural story the claim that the good in question is the good of truth, or more perspicaciously the good of more truth and less falsehood. To see the basic idea, think of a consequentialist view in ethics that we can call Simple Hedonism, which holds that if action A leads to more net hedons than action B, A is morally better than B. Simple Hedonism has a consequentialist structure – actions should be evaluated by appeal to the goodness or badness of the states they bring about – and it makes a particular claim about what the good and bad states are – that they are pleasure and pain. This is the same sort of two-part structure that you see within truth-consequentialist approaches to epistemic normativity. As an illustration of the approach within epistemology, consider the following remarks: “A very plausible set of [cognitive] goals are the oft-cited aims of believing the truth—as much truth as possible—and avoiding error” (Goldman 1980, p. 32); or “An intellectual virtue is a quality bound to help maximize one’s surplus of truth over error” (Sosa 1985, p. 227). The Goldman quotation articulates the claim that believing the truth is good, or that what is good is believing more truth and less falsehood. The Sosa claim adds to this the idea that it is by appeal to this good that intellectual virtue should be understood. Goldman, of course, also uses this good to develop reliabilism and a veritistic approach to social epistemology. This sort of approach is very common in epistemology and I cite Goldman and Sosa largely because their views are well-known and the lines quoted compact expressions of it.<sup>1</sup>

What I want to focus on here in Sect. 7.2 is not reliabilism, or virtue reliabilism, or even truth-consequentialist pictures of epistemic normativity more generally. Rather, I want to focus on something that is implicit in this general approach as it has been articulated or developed throughout philosophy, a dimension of the picture that is not normative but descriptive. This is a certain conception of what truth-conduciveness is, or of what it is to believe as much truth as possible and as little error as possible, to adopt Goldman’s line, or of what it is to maximise one’s surplus of truth over error, to use Sosa’s. On this conception, there are some number of truths and some number of falsehoods and maximising truth is a matter of increasing the number of truths believed and minimising error is a matter of decreasing (or holding at zero) the number of falsehoods believed. Note that there are two separate claims here. The first is a claim about the in-principle countability of what is true

---

<sup>1</sup>There are interesting differences between truth-consequentialist and truth-teleological approaches to epistemic normativity that I here elide as I don’t think they are relevant to the points made in this paper.

and of what is false. The second is a claim about such countability exhausting the quantitative structure of what is true and of what is false. Note that this second claim is not otiose or redundant on the first. Some things can be counted (e.g., the pains one has suffered, the wrongs that one has committed), without it being the case that counting exhausts their measure (how much pain one has suffered, or how much wrong one has committed). The second claim therefore is that not only does what is true and what is false have a countable structure, cardinality gives the correct measure on each.

These two claims are seldom explicit, but they underpin discussions of truth consequentialism throughout epistemology.<sup>2</sup> In Sect. 7.3, I will criticise both aspects of this conception of the quantitative dimension of truth and error. For now, let me bracket these issues to discuss an influential line of criticism that has been directed at truth-consequentialist approaches to epistemic normativity, for doing so will foreshadow more extensive discussion later of the threat of perspectivalism. The basic objection is that the good of more truth and less falsehood can't be the ground of epistemic normativity because not all truths and falsehoods are equal—some are better or worse, epistemically better or worse, than others, in the sense that some are such that believing them makes a big difference to one's epistemic awesomeness whereas others make only a little difference, or perhaps no difference at all.

Here is how Dennis Whitcomb sums up what is known as the trivial truths objection, which has been almost universally accepted:

It is better epistemically to know deep theoretical truths about, for example, metaphysics or physics, than it is to know trivial truths such as truths about the number of grains of sand on the nearest beach. Indeed, even if one were to know a lot of trivial truths, and thereby fulfil the epistemic value of having more rather than less knowledge, one's epistemic states would still be deficient owing to their triviality (2015, p. 313).

Or as Paul Moser put it almost 35 years ago:

It is somewhat implausible to hold that an epistemic agent should aim just to obtain as much truth and avoid as much error as possible. For it is plausible to suppose that some truths are epistemically more important than others (1985, p. 5).

I will not discuss this objection in detail,<sup>3</sup> but want to note two things. First, that the objection threatens not only a truth-consequentialist account of epistemic normativity, but what we can loosely call the objectivity and interest-independence of epistemic normativity. That is, it leaves us with the claim that some truths are more epistemically *important*, *significant*, or *interesting* than others. These adjectives are stand-in labels, of course, not accounts, but they at least suggest a worthiness that is contingent on our particular interests and cares. To be sure, one could provide an objective, interest-independent analysis of any of these, given that they are mere labels. But they are connotative labels and suggest (and are often intended to suggest) a movement from an objective, interest-independent account of epistemic normativity to a kind of perspectivalism wherein the epistemically normative

<sup>2</sup>For an elaboration and defence of this, see Treanor (2018), especially pages 1064–1067.

<sup>3</sup>See Treanor (2013) and (2014) for more discussion.

depends on our situation, or history, or particular and peculiar desires and interests. The second thing I wish to note about the trivial truths objection is that it is apt only if the truth that, for example, my phone number was once 416–684-0019 is the same amount of truth as the truth that, for example, all ordinary matter in the universe is composed of species of atoms with the same number of protons in their atomic nuclei. Those who give the trivial truths objection, or who are moved by it, have simply assumed that every truth is the same amount of truth, something like one unit of truth. We can see why they think this—every truth is *a* truth. How could it not be? But the fact every truth is a truth doesn't entail that every truth is the same amount of truth. We will discuss both points in more detail below. For now, note just that as standardly construed truth-consequentialism faces a perspectivalist threat, but the challenge is defused when we notice that it arises from an objection that conflates the numerosity and measure of *true sentences* with the numerosity and measure of *what true sentences express* (to wit, the truth, or true content), and *mutatis mutandis* for false sentences and what false sentences express.

### 7.3 A Problem for How Truth-Conduciveness Is Understood

This is the part of the story where I argue that the background, prevailing picture of *more truth* and *less falsehood* is mistaken. I describe this picture as background because it has not been explicitly elaborated and defended but rather implicitly assumed. I describe it as prevailing because it is so deeply-seated and pervasive that philosophers treat it as almost analytic, as capturing the meaning of the quantitative vocabulary concerning truth rather than as expressing a substantive philosophical theory of the quantitative dimension of truth. I don't want to spend a great deal of time on this part of the story, as I have argued it at length in several papers. It is important, however, to do more than merely gesture at arguments given elsewhere. The claim I want to dislodge is so engrained and implicit in our understanding that the criticism I want to make of it would probably not be intelligible if I were to merely cite arguments given elsewhere without rehearsing them in outline.

In this section, therefore, I will first outline in more detail the picture that I think is false, and then trace out two main lines of criticism. The aim will be to convey the rough shape and intelligibility of a set of criticisms that have been given in more detail elsewhere. The point of doing this will be to set up a positive story of how the measure of knowledge and error should be understood, one that, like what it supplants, threatens to be deeply perspectival.

What is the picture I think is mistaken? There are two ways to put it, which are importantly different but for the purposes of this paper I will sometimes run together, or at least not take pains to disentangle. One way of thinking of the picture is as the claim that more truth is a matter of more truths or of more true propositions, while more falsehood is a matter of more falsehoods or of more false propositions, where the 'more' in the analyses is the more-ness of cardinality. A second way of thinking of the picture is as the claim that more true belief is a matter of more true beliefs,

while more false belief is a matter of more false beliefs, where again the ‘more’ in the analyses is the more-ness of cardinality. How exactly these views are different turns on questions about the nature of belief that are orthogonal to the issues I focus on in this paper, so I will principally speak of more and less truth, but occasionally of more or less true belief (or of more and less knowledge and error) when it seems stylistically felicitous.

As I mentioned in Sect. 7.2, there are two quite different aspects to this picture. First, there is a claim about structure, about what is true and what is false having a countable structure or a cardinality. Second, is a claim about this countable or cardinality structure exhausting the quantitative structure of what is true and what is false. The second claim could be true only if the first is, but the converse doesn’t hold. This point can be easily overlooked because it is natural to think that if there is such a thing as the number of Xs, then who has more Xs is settled by counting. That observation is itself accurate, but the mistake is to assume that who has more X (or more of X) is always given by who has more Xs. Sometimes that is true – which city has more people in it is determined by whether the number of people in one city is greater than the number of people in another city. That is because ‘more people’ just means a greater number of people. But in other cases it is not true. Who has committed more harm, for instance, isn’t entailed by who has committed more harms. The count of harms done and the quantity of harm done come apart. If we assume for a moment that truths are in principle countable, then it may be that more truth is just the same thing as more truths (plural). But that is a substantial philosophical view, not something that follows just from the fact that, for example, five truths are a greater number of truths than three truths. In this section, I will argue against both of these claims. There is compelling reason to doubt that what is true is countable or has a cardinality (and likewise for what is false). And there is compelling reason to doubt that, even if what is true and false is countable or has a cardinality, counting or cardinality gives its measure.<sup>4</sup>

Why should we think that what is true and what is false isn’t denumerable? The main reason is because we can find nothing good to count. Our central grasp of what is true and of what is false is via sentences of natural language, and although the set of true sentences of a natural language is in principle denumerable, let us grant,<sup>5</sup> its cardinality is arbitrary. This is because, to put it simply, it’s possible to say the same thing in different ways. For this reason, although we can (let us grant) count true sentences, they are not the right thing to count, since different arrangements of words, or arrangements of different words, would yield a different number of sentences while saying the very same thing. To zoom in on one example, the sentence

---

<sup>4</sup>So far in this paper I haven’t carefully distinguished countability from cardinality, although they are of course different: the real numbers are not denumerable yet have a cardinality, the cardinality of the continuum. For the most part this doesn’t matter in what follows and where it does I will be more precise.

<sup>5</sup>The concession is needed for two reasons. First, because it is hardly clear that even sentences have the identity conditions that are required. Second, because, for example, for every real number there could be a sentence that says that number is a real number; but then there would be uncountably many sentences.

‘John is a bachelor’ and the sentence ‘Richard is male’ are each one sentence, exactly one sentence, and so in a matter of speaking are one truth each. The sentence about John says more than the sentence about Richard, however, and if each is true then the sentence about John expresses more truth than does the sentence about Richard. To put this another way, the truth that John is a bachelor is more truth than the truth that Richard is male (more truth, not more true), even though each is one truth. (For simplicity I assume the proper names in each sentence have no descriptive content, or that if they do the scale of that content is identical.)

With this particular example it is tempting to think there is a straightforward containment relation that preserves the intelligibility of more truth being a matter of more truths. A truth to the effect that S is a bachelor decomposes, one might think, into a set of truths that includes the truth that S is male. We can therefore make intelligible that “John is a bachelor” is more truth than “Richard is male” by conceiving of what is said of Richard as a proper part of what is said of John. The promise of this as a general account, however, is illusory, since it is implausible that every sentence of natural language decomposes into a concatenation of sentences that are genuinely atomic in the sense of saying exactly one thing, no more and no less. If sentences did decompose into atomic sentences so understood, then we could in principle arrive at a set of atomic sentences that has a cardinality. But they don’t so we can’t.

The second aspect of the prevailing picture of the quantitative dimension of knowledge and error goes beyond the claim that what is true and false is denumerable to claim that, in addition, counting gives the proper measure of it. To put this another way, the second aspect of the prevailing picture is the claim that not only does the world divide into facts, to use Wittgenstein’s expression, the facts into which the world divides are all the same size. There is compelling reason to doubt this, however. There are intuitive cases where one person knows much less than another even though they know the same number of truths, however we wish to understand that. The point is most easily made if we consider not knowledge simpliciter but knowledge of a restricted domain. Consider Edinburgh, the city, and two people Charlie and Zachary who wish to know it. Suppose Charlie knows some large number of truths,  $N$ , akin to these:

Edinburgh is the capital city of Scotland.

In 2019 approximately 500,000 people live in it.

It was a medieval city formed around a castle, with a defensive wall.

In the medieval period, it was crammed with people living in high rises and deep underground so that they could live within its protective wall.

As conflict between Scotland and England abated, the city expanded northward, outside the walls, with the creation of New Town, an extensive and splendid example of Georgian town planning.

Today Edinburgh New Town is almost perfectly intact and part of a UNESCO world heritage site, along with Old Town, the medieval core.

Edinburgh was the seat of the Scottish Enlightenment, a remarkable flowering of science, philosophy, literature, and culture generally.

Meanwhile, suppose Zachary knows exactly the same number of truths about Edinburgh, but the truths are akin to these:

Genghis Khan never visited Edinburgh.

Socrates didn't visit Edinburgh.

Edinburgh would not fit on the head of a pin.

Edinburgh is not identical to the number 5.

Edinburgh is not identical to the number 6.

Edinburgh is named 'Edinburgh'.

Edinburgh is not named 'Audrey Hepburn'.

If we compare who knows more about Edinburgh, the right thing to say is that Charlie does. This is because Zachary, as we imagine him, has no clue what Edinburgh is; he just knows that, whatever it is, Genghis Khan never visited it, it wouldn't fit on the head of a pin, it is not identical to the number 5, etc. What he knows about Edinburgh is not enough to distinguish it from things that are wildly unlike Edinburgh, such as the Andromeda Galaxy or Cheese Whiz. In contrast, what Charlie knows about Edinburgh homes in on Edinburgh, it is such that although there are many things he could not distinguish from Edinburgh, anything that he could not distinguish from Edinburgh would have to be a great deal like it. The takeaway is that Charlie and Zachary know as many truths about Edinburgh as each other, but Charlie knows more about Edinburgh.

The argument just given moves from amounts of knowledge to amounts of truth. This should be unobjectionable, since what would more knowledge be if not knowledge of more truth? If one is worried about this, however, here is a similar argument that talks just of truth and not of knowledge. Think of the whole truth about Edinburgh; this includes all the truths that Charlie knows about Edinburgh, all the truths that Zachary knows, and countless more that neither does. Now take just those truths that Charlie knows about Edinburgh and compare with just those truths that Zachary knows. Which of those is more truth about Edinburgh? One might here be tempted to insist that they are the same amount of truth – some number  $N$  truths's worth of truth, as it were. But we are asking this question in an effort to discern whether the number of truths gives the measure of truth, and this response simply assumes that it does. What we need to do is step back a bit and consider each body of truth, asking which seems, intuitively, to be more truth about Edinburgh. My claim is that the body of truth that Charlie knows is more. Anything that had just the properties the Charlie-truths ascribe to Edinburgh would have to be a lot like Edinburgh; but something could have the properties the Zachary-truths ascribe to Edinburgh and still be wildly unlike Edinburgh, e.g., Cheez Whiz and the Andromeda Galaxy.

I don't think this quick sketch of each line of argument is full enough to be convincing, but I hope it gives some sense to the idea that the measure of truth is not given by counting or cardinality. It is important to note that one should be on board with this conclusion even if one only finds the first line of argument persuasive. The second line of argument assumes that truths are denumerable or has a cardinality and argues that, even then, more truth is not a matter of more truths. But first line of argument, which says there is nothing good to count, is all that's needed to undermine the standard approach to understanding truth-conduciveness within epistemology.

## 7.4 What, Then, Is More Truth and Less Error?

In Sect. 7.2, I introduced truth-consequentialist approaches to epistemic normativity and showed how, as truth-conduciveness is standardly understood, they threaten a kind of perspectivalism about epistemic normativity. In Sect. 7.3, I showed how this standard understanding of truth-conduciveness is mistaken. In this part of the story, I want to propose something better. It will be an approach to the quantitative dimension of knowledge and error that avoids the critical problems I argued the standard understanding faces. We will see in the next Section that although this approach radically reconceives how truth consequentialism should be understood, the perspectivalist threat survives. If anything, in fact, it grows even stronger and more pressing. But we get to that in the fullness of time; for now, let us just see how the quantitative dimension of knowledge and error should be understood.

The standard conceptualisation of truth consequentialism rests on an unstated picture of the quantitative dimension of knowledge and error that takes the object of belief to be countable and cardinality to exhaust its quantitative dimension. Both aspects of that picture are wrong, as I have argued. I now want to show how we can conceive of the quantitative dimension of knowledge and error in a way that doesn't rest on or require these problematic claims. The proposal will be schematic and unsatisfying, in that it could hardly be thought to be a fully elaborated, detailed picture. But it will be well-motivated, I believe, since it will connect the issue of the quantitative dimension of knowledge and error to related foundational issues in metaphysics and philosophy of language in an illuminating way. Moreover, to the degree that mysteries survive—and they do!—they are mysteries we face generally, mysteries in metaphysics and philosophy of language that we are already stuck with. The point of this section, therefore, will not be to offer a full or final analysis of the quantitative dimension of knowledge and error, but to show how the problem collapses into a more familiar problem that pervades philosophy. This move is what will let us avoid the problematic aspects of the standard picture, but it is also, as I will show in the next section, the very thing that threatens to deepen the perspectivalist threat.

I'll introduce and outline the positive proposal with an analogy.<sup>6</sup> Suppose two artificial-intelligence devices are each charged with the task of building a physical duplicate of a target object, say some particular apple. Device A goes about the task, scanning the target region to discern what is there and what it is like, and then setting out to acquire appropriate materials and assemble them in an appropriate way. Device B does the same. After some period of time, both devices come to a rest, their task complete, or as well done as they are capable of doing. Each device has produced a physical object. The object that Device A produced is a spitting image of the target apple; it is visually indistinguishable to the human eye, but moreover has the very same mass, density and density distribution, a remarkably similar chemical structure, and so on. Let us suppose, in fact, that although it is not a perfect

---

<sup>6</sup>For more extensive discussion, see Treanor (2019), pp. 35–38.

molecule-for-molecule duplicate, it would take a team of scientists some effort to tell the apples apart, or to speak more precisely, to discern which is the apple and which is the artefact. The object that Device B produced, in contrast, is much less similar to the target apple. It is roughly the same size and shape, but is a bit taller, not quite as wide, has a pebbly rather than smooth skin, has an interior structure that is white, moist and sweet like the apple's but with the density of foam, and so on. It could be mistaken for the target apple after a superficial inspection, but generally the similarities between it and the apple are not nearly as deep and pervasive as with the object produced by Device A.

If we think about this situation with an eye to judging which device, A or B, did a better job of duplicating the apple, there is no question that Device A did. The object that it produced was not a perfect duplicate, but it was very close. The object that Device B produced hardly competes. To be sure, that object is still a remarkable achievement. This is because if one imagines the universe of possibilities, all the things that Device B could have produced, the object that B made is much more similar to the target apple than most of that; Device B could have gotten much less right and much more wrong. It might have made an electron, or a supernova, or a doppelgänger of Joan Rivers, or any of countless other things that are wildly dissimilar to the apple. The object that it made is still, however, substantially dissimilar to the target apple compared to the object that A produced. That much is clear.

Let us consider, though, whether we can say *why* this is so. By this I don't mean what it is about the two devices, how they were designed, that led one to do a better job than the other. The question concerns why it is that the object that Device A produced is better, qua duplicate of the target apple, than the object that Device B produced. That is a question not about the devices but about the relation between the two artefacts the devices produced and the target apple.

A natural answer to this question is that Device A, in making the object it did, *got more right*, or more fully, *got more right and less wrong*. It is almost as natural to make the further claim that Device A, in making the object it did, *got more things right and fewer things wrong*. This further claim, in effect, says that the object that A produced shares a greater number of properties with the target apple than does the object that B produced, while having fewer properties not in common. That is what it would be for Device A, in making the object it did, to get more things right and fewer things wrong. This further move is appealing but it is essential to recognise that it is a mistake. As Goodman pointed out decades ago, "any two things have exactly as many properties in common as any other two" (1972, 443).<sup>7</sup> We should agree, therefore, that Device A, in making the object it did, got more right and less wrong, but not take the further step of claiming that it did this by getting more things right and fewer things wrong, that is, by giving the object it produced a greater number of properties the apple has and fewer properties the apple doesn't have.

---

<sup>7</sup>Goodman, of course, took the point to be that we should be sceptical of similarity. But the momentum in philosophy has been in the other direction, to accept similarity and affirm a non-egalitarianism about properties, with some being more natural than others. See Lewis (1986), pp. 59–69.

The answer to the question of why the object that A produced is better, qua duplicate of the apple, than the object that B produced lies already in the situation as we described it. The object that Device A produced was more similar, overall, to the target apple than was the object that Device B produced. This is not a matter of the number of properties had in common. We don't know what it is a matter of (or even whether it is a matter of anything as opposed to basic). But it is the right thing to say, for it is better to say something unclear but true than something less unclear but false. Moreover, there is something else in this analogy that should be drawn out, as it will be relevant later. It's not just that similarity doesn't consist in the number of properties. It is that some properties make more of a contribution to similarity than others. This is a consequence of there being greater and lesser similarity but not a greater and lesser number of properties in common. But we can also see the force of it intuitively. The two apples on my kitchen table are substantially similar by virtue of sharing the property of being the fruit of a McIntosh tree, but not substantially similar (or at least less substantially similar) by virtue of sharing the property of having a volume smaller than a billion light years cubed.

This analogy concerns the similarity relation that holds between objects. But it is easy to see that the same issues are at stake when we think about quantities of knowledge and error. Instead of thinking of the objects that A and B produced, think of mental representations that correspond exactly, save being mental rather than real, to those two objects. One mental representation ascribes to the apple all and only the properties that the object produced by Device A has, and a second mental representation ascribes to the apple all and only the properties that the object produced by Device B has.<sup>8</sup> Each of those mental representations of the target apple seem to get the apple right to some degree and wrong to some degree, or to be, as I put it early in the paper, an admixture of knowledge and error. The question that concerns us is how to assess the quantitative aspect of that. We already know, from discussion earlier in the paper, that we cannot count how many truths and how many falsehoods each mental representation consists in, both because it doesn't consist in some number of truths and falsehoods at all, and because even if it did, cardinality would not exhaust the quantitative structure. We are now in a position to see that this should not have been a surprising conclusion, despite how prevalent the counting conception of the quantitative dimension of knowledge and error is. If it is widely accepted that one object is not more similar to another by virtue of sharing a greater number of properties while differing on fewer properties, why should we have ever thought that a representation of an object gets more right and less wrong by correctly representing a greater number of the object's properties while incorrectly representing a smaller number of the object's properties? To illustrate this with an example, think again of the apples on my kitchen table. We grant that the one on the left is more similar to the one on the right by being, like it, the fruit of the McIntosh tree than it is by being, like it, such that its volume is less than a billion light years

---

<sup>8</sup>For simplicity I am here focusing on intrinsic properties understood as those that do not differ across duplicates. The proper understanding of 'intrinsic' is vexed, but as I appeal to it only to simplify the issue for the purposes of presentation I think it can be bracketed.

cubed. Why, therefore, should we have ever thought that knowing that it is the fruit of a McIntosh tree is to know exactly as much (exactly as much truth or true content) about it as knowing that it is smaller than a billion light years cubed? This is, I think, a remarkable instance of how parallel lines in philosophy have drifted apart.<sup>9</sup>

## 7.5 Perspectivalism Renewed

Where are we? We have seen what is wrong with the counting or cardinality approach to the quantitative dimension of knowledge and error. Just as you can't make sense of how similar an artefact is to an apple by counting the properties they share and differ on, you can't make sense of how much a mental representation gets right and how much it gets wrong by counting the number of truths and falsehoods believed. It either can't be done, at all, or it can be done but is unreliable, generating the wrong answer at least much of the time. We have also seen how the question of how to understand the quantitative dimension of knowledge and error is intimately related to questions in metaphysics about the nature of similarity. What remains to be shown is a certain vulnerability to an objection, one that is similar to one that I considered at the end of Sect. 7.2 of our story. I will close the paper by drawing out and discussing this vulnerability.

Recall the objection, considered earlier in the paper, that insists that since some truths are epistemically more valuable than others, epistemic normativity cannot be grounded in the good of more truth and less falsehood. The idea is supposed to be that a so-called trivial truth and a so-called significant truth are each one truth, so believing each contributes equally to how much truth a person believes. But if that is the case, and it is epistemically better to believe the significant truth than the trivial truth, something else (beyond a mere increase in the amount of truth believed) has to explain the difference in epistemic normativity. This is a poor argument, despite its widespread acceptance, because it says nothing to establish that the epistemically more valuable truths are not *more truth* than the epistemically less valuable truths; that is a hidden assumption that arises from conflating the numerosity and quantitative dimension of the vehicle (declarative sentences of natural language) with the numerosity and quantitative dimension of the content that such vehicles express. Recall, however, that the objection, were it sound, would threaten the objectivity and interest-independence of epistemic normativity. This is because it wouldn't be that what is epistemically valuable or good is more truth, or more truth and less error, where facts about what is more and less truth or error are objective,

---

<sup>9</sup>Goodman's insistence that similarity is not a matter of shared properties was given within a framework of thinking of properties as extensional. But notice that his scepticism becomes even more vehement when you turn toward the semantic analogy by thinking of intensional properties: "The inevitable suggestion that we must consider intensional properties seems to me especially fruitless here, for identifying and distinguishing intensional properties is a notoriously slippery matter, and the idea of measuring similarity or anything else in terms of number of intensional properties need hardly be taken seriously." (1972, 444)

impersonal and interest-independent. Instead, the objection has it that what is epistemically good is more truth that is significant or important, where the correct analysis of that seems (most plausibly) to depend on our interests and cares, or at least on contingent properties about us as knowers and cognizers. The arguments in Sects. 7.3 and 7.4 of this paper show how this objection is mistaken. We *can* make sense of why so-called significant truths are more truth than so called trivial truths. It is because they are truths about properties that make a greater contribution to overall similarity. We don't need to say that each is one truth, so therefore the same amount of truth, so therefore that the epistemic difference between them needs to be explained by something other than the contribution they make to truth-conduciveness. In other words, the arguments we have considered so far in this paper salvage the idea that truth-conduciveness alone, rather than truth-conduciveness plus significance, can ground epistemic normativity.

The problem, however, is that the positive story I have offered closes this door on perspectivalism with a draught that opens another. To articulate the worry, let us turn again to Goodman. After pointing out that any two objects share the same number of properties, he points out that one could hope in response to count not just any properties, but properties that are privileged in some way:

More to the point would be counting not all shared properties but rather only *important* properties—or better, considering not the count but the overall importance of the shared properties. Then *a* and *b* are more alike than *c* and *d* if the cumulative importance of the properties shared by *a* and *b* is greater than that of the properties shared by *c* and *d*. But importance is a highly volatile matter, varying with every shift of context and interest, and quite incapable of supporting the fixed distinctions that philosophers so often seek to rest upon it (1972, 444, italics in original).

Here the spectre of perspectivalism enters very forcefully, even more forcefully than it entered on the standard understanding of truth-conduciveness wherein more knowledge and less error is just a matter of counting. If the right way to understand more knowledge and less error is itself by appeal to similarity, as I have proposed, then a kind of perspectivalism or interest-relativity will infect not only epistemic normativity, but even the more basic question of whether one body of knowledge is or isn't *more* knowledge than another. That is a much more thorough and threatening perspectivalism.

The quotation from Goodman illustrates that what we can call 'importance' has to enter *somewhere*. It could enter after quantities-of-truth, as it does with the objection that I have been concerned to refute. That is, one might say that every truth is one truth and that amounts of truth are just numbers of truths; but then one has to explain epistemically normative differences as a function of amount-of-truth and importance-of-truth. Or one could take 'importance' to play an ineliminable role in how amounts-of-truth should be understood, as I do. In this case, nothing is needed to ground epistemic normativity other than amounts-of-truth, but 'importance', and its attendant threat of perspectivalism, is still in the picture, just lurking in a different shadow. Moreover, this threat of perspectivalism is more haunting. In the original position, one could take epistemic normativity to ultimately depend in some way

on our interests and cares, but at least, one could say with relief, how much one knows, fails to know, and is wrong about are all objective facts of the matter. On the positive story I have offered, in contrast, *even that* may be deeply perspectival. That is a much worse place to be in for those who long for the lonely indifference of interest-independent objectivity, especially concerning matters as deep and central as truth and falsehood.

My aim in this paper is principally to trace out the threat of perspectivalism rather than defeat it. This is in the spirit of knowing thy enemy, as my own sympathies are anti-perspectivalist, for better or worse. That said, I want to close by indicating why I think the positive story I offer has more potential to deliver an analysis of ‘importance’ that isn’t interest-relative or perspectival than does the original position. I will unfortunately be able to do little more than gesture at an argument I hope to develop more fully elsewhere.

Here is the thought in outline: A truth-conduciveness approach to epistemic normativity has to let ‘importance’ or ‘significance’ into the picture somewhere. It could be that amounts of truth are independent of importance/significance and they enter later, as a link between amount of truth and epistemic value. That is the standard position, so to speak. Or it could be that amounts of truth themselves are the amounts they are because of the importance/significance of the truths involved, as I prefer. If the standard position is correct, it is difficult to see how there could be an analysis of ‘importance’ or ‘significance’ that doesn’t depend on our interests and cares, since importance and significance are at the level of value, and need to be understood as modifiers of value. Now, to be sure, value might be as objective as anything can be, but it is at least very plausible that if two truths are the same amount of truth but one is more epistemically valuable than another, the value has to be explained on the mind side of the mind/world relation rather than on the world side, and this brings to the fore perspectivalist explanations. The story that I prefer is very different. ‘Importance’ and ‘significance’ refer to an inegalitarianism on the world side of the relation, specifically to a hierarchy in the domain of properties. Some properties are more natural than others, in the sense in which that term has become familiar in metaphysics, and epistemic importance or significance rests not on our private cares and position but on the impersonal structure of reality.<sup>10</sup> To be sure, that might ultimately be unintelligible as independent of our interests and cares (as Goodman thought). But there are two related reasons to be hopeful.

First, it seems plausible or intelligible that a world bereft of persons is bereft of value, but much harder to see how a world bereft of persons is bereft of an inegalitarianism in the domain of properties. Electrons would still exist, one wants to think, and they would be similar to each other, and properly a class, by virtue of being negatively charged. But the class of things that are negatively charged or citrus fruits, or the class of things that are smaller than the moon, would be much less natural. To be sure, this is difficult to explain or understand. But it is also much harder to reject. If this is right, the inegalitarianism of amounts-of-truth is rooted in

---

<sup>10</sup> See Almotahari ([forthcoming](#)) for an illuminating and sustained discussion of how the quantitative dimension of truth relates to realism about the normativity of metaphysical structure.

the incurious and perspectiveless world. Second, if ‘importance’ or ‘significance’ are located in the domain of value, as the link between amounts of truth and epistemic goodness, it seems *sui generis*, in the sense of standing alone and requiring an explanation of its own. It is at least plausible that there is little to appeal to other than persons and their situations. In contrast, if importance/significance is located in the domain of properties, the epistemic inequality is assimilated to an inequality that is widely thought to be required to make sense of all sorts of things (e.g., but hardly uniquely, Lewis 1983).

## Bibliography

- Almotahari, M. (forthcoming). *Rescuing realism*.
- Goldman, A. (1980). The internalist conception of justification. *Midwest Studies in Philosophy*, 5, 27–51.
- Goodman, N. (1972). Seven strictures on similarity. In *Problems and projects* (pp. 437–446). New York: Bobbs-Merrill Company.
- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61, 343–377.
- Lewis, D. (1986). *On the plurality of worlds*. Oxford: Blackwell.
- Massimi, M. (2016). Three tales of scientific success. *Philosophy of Science*, 83, 757–767.
- Massimi, M. (2018a). Perspectival modelling. *Philosophy of Science*, 85, 335–359.
- Massimi, M. (2018b). Four kinds of perspectival truth. *Philosophy and Phenomenological Research*, 96, 342–359.
- Moser, P. K. (1985). *Empirical justification*. Dordrecht: D. Reidel Publishing Company.
- Sosa, E. (1985). Knowledge and intellectual virtue. *The Monist*, 68(2), 226–245.
- Treanor, N. (2013). The measure of knowledge. *Noûs*, 47, 577–601.
- Treanor, N. (2014). Trivial truths and the aim of inquiry. *Philosophy and Phenomenological Research*, 89, 552–559.
- Treanor, N. (2018). Truth and epistemic value. *European Journal of Philosophy*, 26, 1057–1068.
- Treanor, N. (2019). The proper work of the intellect. *Journal of the American Philosophical Association*, 5, 22–40.
- Whitcomb, D. (2015). Epistemic value. In A. Cullison (Ed.), *The Bloomsbury companion to epistemology*. London: Bloomsbury Academic.

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as 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.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

