

## Contrastive knowledge

DRAFT: final version is in Martijn Blaauw, ed. *Contrastivism in Philosophy* Routledge 2012 ( <http://www.amazon.ca/Contrastivism-Philosophy-MartijnBlaauw/dp/0415878608/> )

The concept of knowledge is a very ordinary one, in spite of its philosophical glory. Like the concepts of a thing or a person or an animal, or the concepts of cause, or of action, it is one that we use every day and would be lost without. We use it when we explain people's actions ("she visited you because she knew you would never visit her"), and when we say what testimony should be trusted ("she knows where he hid the loot: I'd pay attention to what she says and does"), and when we justify our own actions ("I looked in the fridge because I didn't know where else it might be.") This ordinariness makes philosophical skepticism threatening in an immediate way. But it combines uncomfortably with the high intellectual demands of some philosophical accounts of knowledge. There are reasons for these high demands arising from the function the concept has historically played in philosophy, associating it with the ideal of a rational intellectual agent. To reconcile the pulls from everyday life and the pulls from philosophy we need to understand the reasons why we have the concept: what are the core functions that it serves in our ordinary thinking?

The claim of this paper is that the everyday functions of knowledge make most sense if we see it as contrastive. That is, we can best understand how the concept does what it does by thinking in terms of a relation "*a* knows that *p* rather than *q*". There is always a contrast with an alternative. Contrastive interpretations of knowledge, and objections to them, have become fairly [common] in recent philosophy. The version being defended here is fairly mild in that there is no suggestion that we cannot think in

terms of a simpler not explicitly relation "*a* knows that *p*". Some, e.g. Schaffer(2005) and Karjalainen and Morton (2003) have hinted that this stronger possibility may be right. But all that I am arguing now is that facts that are easily expressed in contrastive terms are vital to understanding why we need the concept of knowledge. In a piece that is in some ways a companion to this one, (Morton forthcoming a), I give a general survey of theories of contrastive knowledge and the differences between them.

**beneath propositional attitudes: tracking** Knowledge is a factive relation: it holds between people and actual facts. You cannot know something that is not so. Facts are problematic things, somewhere between situations and propositions. Most relations between people and things, as between things and things, are just that, holding between individuals without involving anything sentence-like. And they hold just between the individuals they hold between, and not between nearby or alternative individuals. So when using them to explain or predict, some explicit or implicit reference to laws of nature, causation, or counterfactuals - something in the realm of the nomic - is needed. Many basic epistemic relations connect individual people to individuals in their environment: *a* sees *o*, *a* perceives *o*, *a* recognizes *o*, *a* remembers *o*. Sometimes we use an embedded sentence to describe what is really a relation between individuals: *a* knows that *o* is at location *l* (*a* locates *o* at *l*), *a* knows which species *o* belongs to (*a* classifies *o*). I shall assume that one important and central function of noting and stating such relations is to help anticipate and explain actions of individual people directed at individuals. Why did Alfred duck? Because he saw the stone whizzing towards his head. Why did Agatha return to the crossroads? Because she located her cellphone there. It takes some care to formulate these without using propositional attitude terms, and I take this to be due to the way that propositional attitude language dominates our descriptions of sentient life.

The point, though, is that explanations couched in this language can succeed because of the multitude of relations by means of which people can direct their actions at objects. (It would be easier to describe this if I could assume that propositional attitudes are a linguistic veneer over an underlying pattern of relational thought. But although I think something of this sort may be true, and have begun to explore the idea elsewhere - see Morton forthcoming b- this is not the occasion to defend it.)

Note the way that the epistemic element in these explanations "saw", "located", serves both as a fact and as a law. In saying that Alfred saw the stone (and its trajectory) we are saying that Alfred's information state is related in a particular way to the stone and what it was doing. There is a flow of information from the stone to Alfred, a flow being a causal process that relates objects in one kind of situation to a characteristic kind of result. Combined with tacit assumptions about his tendency to avoid injury, it is as if there is a causal flow from the motion of the stone to Alfred's action. It is a flow that we the observers or describers can ride along, as when we see him duck and duck ourselves in order to avoid the projectile we take him to be avoiding. The 'width' of the flow is left vague in the explanation, so we are not told exactly what other things Alfred did or would have been able to see. It may be necessary to say more about this, as we do when explaining why Alfred avoided the stone thrown by Martha but not the tomato thrown by Nelly. "There was a branch in the way, so that when he turned his head the tomato was behind it. If Martha had thrown the stone up rather than straight he wouldn't have seen it either." (I am appealing to what I take to be common to most accounts of explanation, though my expression may seem Hempelian. The idea of a flow of information is Dretske's. See Dretske 1981.)

We direct our actions at objects by keeping track of them. We note their

locations and attributes and how they change. Central to this are tracking relations: our representations of things are causally linked to the locations, colours, and other attributes of objects, so that if these change so do our representations. A cat chasing a mouse leaps where the mouse is. These require a certain sensitivity to the state of the object, which is supplied by different mechanisms for different aspects of the state. They all feed into a relatively uniform way of anticipating the actions of human and non-human agents, though: predict that some aspect of a relation between the agent and an object will remain invariant under changes of surrounding conditions. The cat's direction vector will remain pointed at the mouse; the bird-watcher's thumb will approach the page of the book for that kind of bird. The central point for present purposes is that in mere mortals these mechanisms of sensitivity are of limited accuracy and scope. We can keep track of where a prey animal is as long as it has not taken certain evasive actions (and a predator can keep track of us as long as we have not taken certain measures). We can keep track of roughly where it is, close enough for most of our purposes but not for all conceivable purposes. Even when normal conditions of detection are satisfied, we may be unable to distinguish between usual and rare trajectories or attributes: if the mouse is between the cat and the sun the cat may be misled by its reflection in a stream. As a result, tracking is inevitably contrastive. The cat can locate the mouse as being in front of it rather than 10 degrees to the left, but not as being in front of it rather than 1/2 degree to the left. And not as being in front of it rather than between it and the setting sun, above the reflective stream.

Contrastivity is inherent in tracking, and tracking is basic to the purposes for which we also use attributions of knowledge. I have just suggested how tracking connects with explanation and prediction, though obviously there is a lot more to say. Tracking also has a natural connection with testimony, via behaviour that indicates an agent's link to an object. Suppose that we are

using a dog to track some prey. After sniffing around the dog sets off in a definite direction. We take the dog to be tracking the prey, literally, and follow her. We follow her because we take there to be a counterfactual link between her behaviour and the path taken by the prey. If the prey had gone a different way the dog would have set off in a different direction, so her behaviour "tells" us which way to follow it. Here too there are limits. Dogs are notoriously prone to following a trail in the wrong direction. So we are told that the prey has gone along this route, but are less sure that it is towards the end rather than towards the beginning of it.

Tracking, with its limits, connects to the aspirations of inquiry, too. An owner training a hunting dog, wants it actually to track. On any particular occasion the dog's setting off in the direction of the prey does not count as success unless the dog would have gone off in another direction if the prey had taken another direction. Within limits: one doesn't expect a well-trained dog to fly, if the prey has taken a ride on a helicopter. A person wanting to become a good bird-watcher aspires to saying "finch" when it is a finch, and wants that had it been a grosbeak instead she would have said "grosbeak". She does not aspire to telling two year old finches from twenty five month old finches. An apprentice astronomer who guesses which planet is near the horizon is rebuked by his mentor even if he has guessed right: if it had not been that planet he would have made the same guess.

I have avoided the verb "know" throughout this section. (And I have only said "knowledge" once.) That is the point. The focus is on the purposes for which we use the concept of knowledge, and how they inevitably bring in considerations of the aspects of an object that an agent is and is not in contact with, themselves naturally expressed in a "rather than" idiom. When the topic is human belief we express many of these ideas in terms of knowledge, and then to express the contrastivity we say "knows that p

rather than q". But the roots lie deeper, and even if one balks at ascribing knowledge to, say, tigers, one will need to be able to say that e.g. the tiger has traced one to one's hiding place, that this means that if one had hidden in the next bush it would have found one there, and that it does not mean that if one had taken a ride out of the park it would have been waiting at the hotel. The tiger has located one at this bush rather than that one, but not at this bush rather than the hotel.

**contrastive counterfactuals** In an important article on contrastive causation, Jonathan Schaffer, who has played a central role in making philosophers take the idea of contrastive knowledge seriously, introduced the idea of a contrastive counterfactual (Schaffer 2005). If c rather than c' had occurred, then e rather than e' would have occurred. In the way of understanding this that I think is most relevant to our current concerns here, c' and e' are actual events. It is possible that there is a workable definition of this conditional in terms of the Lewis-Stalnaker counterfactual, or in terms of more recent refinements of it. (It is an important question how it relates to the non-contrastive conditional, that deserves serious attention.) It is also possible that it is a more fundamental idea, and the order of explanation ought to go in the other direction. This possibility becomes more plausible when we make the connection with contrastive knowledge, in particular with tracking.

We are fairly comfortable ascribing knowledge to individuals in many cases when they perceive their environment, even when the perception could easily have been fallible. In such cases tracking analyses can give the intuitively wrong answer. For example a person is watching a fly cross in front of a window. The window is partially reflective and a moment later a different fly will pass behind the person so that its image will appear on the glass, making it seem as if a fly is taking the path of the first fly. The

person's visual systems are working well, and she comes to think that there is a fly directly in front of her, as there is. The flight of both flies is very erratic, though, so that it could easily have been that the second fly was the one she saw, which would have been behind her. If the details are spelled out suitably, it is natural to say that she knows that fly one is directly in front of her. (Though we might conclude that it is natural but wrong, if we had good enough reasons.) One consideration we might use to back up the ascription is that she was tracking the fly: if it had been a little higher or a little to the left she would have taken it to be higher or to the left. But it is not clearly true that if the fly had not been directly in front of her she would not have taken it to be directly in front of her. Which is the more likely ('nearer', 'more accessible') possibility: the possibility in which fly one is not there because it is further away and she sees it as further away, or the possibility in which fly two is behind her and is taken to be fly one in front of her? The balance between these may be very delicate, and the English conditional is surely imprecise enough that there is sometimes no answer. But the more focused contrastive conditional "if fly one had been one degree to the left of centre she would have seen it one degree to the left of centre" may be unproblematically true.

The essential point is that a simple counterfactual "if p then q" may lack a truth value, or have one that is extremely hard to determine, while a contrastive counterfactual "if p rather than p' then q rather than q'" is straightforward. And it is very plausible that very often when we take a conditional as true we implicitly supply a pair of contrasts, which contrasts depending on context. Our evaluation of "if she had said that to me I would have been insulted" may be different depending on whether we take it as "if she had said that rather than ...", "if she had said that to me rather than to you ...", "... I would have been insulted rather than amused", "... it would have been me rather than you who was insulted", and various combinations

of these. So too apparently simple attributions of knowledge change their plausibility when we highlight different alternatives. She knew that fly one was *there* then; she knew that fly one was there *then*; she knew that fly *one* was there then; she knew that *fly* one was there then. And very often these contrasts will correlate with switches from one contrastive conditional to another, indicating one tracking relation rather than another. To say this is not to present an analysis of knowledge, contrastive or otherwise, in terms of tracking. The suggestion is just that considerations about tracking can influence our judgments about what a person knows, and that tracking, being a counterfactual concept, is sensitive to the contrastive considerations that tune our judgments about counterfactuals.

These considerations connect with observations made by both critics and defenders of tracking analyses of knowledge, to the effect that when defending the idea that to know one must track, we are choosy about what we are to count as a near alternative to the actual situation. A forceful exposition is found in Sherrilyn Roush(2005). (I am thinking especially of chapter 4; I do not mean to endorse the details of what Roush is describing, which may well be right but which I find very hard to follow.) For example if someone looks up and knows by seeing that it is a full moon, we do not consider the nearest situation in which at that very moment the moon is not full. That may be a situation in which the moon does not exist or the processes that form it give it a different orbit. In that situation that very person may not exist. Rather, we think in terms of “if she had been standing there on a day on which the moon was not full ...” or “if she had been somewhere from which the moon was not visible”, or “if Venus had been shining near that spot in the sky”. But these are different, and correspond to “knows that the moon is full rather than in some other phase”, “knows that it is the moon that is presenting that full appearance”, “knows that it is a full moon rather than a reflection in a contact lens”. (Of course we can also



attribute “knows that it is a full moon rather than an alien visitation”, and even “knows that the moon is full rather than her medications producing hallucinations”; see the section **full contrasts** below.) The conclusion to draw is that we describe the information-management of humans and other creatures in terms of how they keep track of facts around them, that these are very sensitive to the contrasts we read into associated conditionals, and that we take such informational states into account in attributing knowledge, which thus acquires an often hidden contrastivity.

**evidence** Processes of the kind that allow us to keep track of things, though fundamental, are just one source of knowledge. Another basic source is the force of evidence. We do not have a generally accepted understanding of the relation between evidence and theory in the philosophy of science, or indeed in statistics. But many cases are uncontroversial, and some general facts are fairly well established. One is the essential role of background assumptions, especially those that determine the probability that some observable consequence will be found if a hypothesis is true and if its main rivals are true, and which determine which hypotheses are the main contenders. Another is the role of those alternative hypotheses, leading to the statistical dialectic of null hypothesis, alternative hypothesis, and test of significance.

Both of these facts lead to contrastivity in the force of evidence. Suppose that we want to know whether a coin is fair. We assume that it has a constant bias to heads or tails, which will be zero if it is fair. We toss the coin twelve times, and observe that it falls HHTTTHTHTHTT, five Heads and seven Tails. Calculating, we find that this would be very unlikely if the coin had a strong bias to H (for example such that it will on average land heads 0.7 of the time), and fairly probable if the coin is fair. So the null hypothesis of Fair is favoured over the alternative hypothesis of strong bias to heads.

But our experiment has told us nothing to rule out a different hypothesis, that it has a much lesser bias to heads. And it does not provide evidence against the hypothesis that the probability of heads varies from one toss to another. (Perhaps as the coin ages the distribution of metal in it alters; after all the ratio in the first six tosses was 3:3, and in the next six was 2:4.) One might conclude from this in the vocabulary of Fisherian statistics, that we now know that the extreme-H hypothesis can be rejected, but not that the null hypothesis, is true. Yet this leaves out the possibility that from this or possibly more similar evidence we can become reasonably certain that the coin is fair. Bayesians will stress this point, though they too are uncomfortable talking of knowledge. What does seem clear is that one can have good evidence to decide between Fair and Very Biased, inadequate evidence to decide between Fair and Slightly Biased, and no evidence at all to decide between any of these and Variable Probability. Suppose that the coin is fair, and that the fact that it is fair rather than biased is a cause of its exhibiting the kind of pattern of which HHTTTHTHTTT is an instance. (This second fact could be seen as reducing to the truth of suitable contrastive counterfactuals.) Then one can be said to know that it is fair rather than very biased, and not to know either that it is fair rather than slightly biased or fair rather than having no constant bias.

These are not essentially different from contrastive knowledge based on limited powers of discrimination. Suppose you can tell dogs from cats but not from wolves, and you correctly identify the animal before you as a dog rather than a cat, but should not be treated as a good source on whether it is a wolf. You must be using some clues about what distinguishes dogs from cats. They may be obvious clues, but they may also be subtle and hard to access consciously, as it might be if you can tell small dogs from large cats in the moonlight. These serve the role of evidence: the characteristic dog walking gait is like the run of five heads and seven tails. Or to put it

differently, your ability to use either a string of random data or a typical dog feature, in order to evaluate the suggestion that you are dealing with a dog or a fair coin, is a limited discriminatory ability, which can indicate that some possibilities rather than others are actual.

It is worth stating explicitly here that the ability to discriminate two possibilities does not establish contrastive knowledge, if it is taken to mean just that if A or B is the case then it must be A. If your most informative report is "the coin is more likely to be fair than very biased, though it might be slightly biased", or "if it's either it's a dog", then what you have is not contrastive knowledge, at any rate not of the species or of the bias. In order to know that it is a dog rather than a cat, you must first believe that it is a dog, and then your belief must be linked in an appropriate manner to the fact that it is a dog. What counts as an appropriate manner is something that divides epistemologists, in particular internalists and externalists, for reasons that thinking contrastively is not going to dissipate. But many cases are uncontroversial, and it is clear that evidence is often essential, that a discriminatory skill is often essential, and that both typically separate one possibility from others, leaving further possibilities uneliminated. The requirement of evidence or a causal connection comes in here. In some cases the evidence that A is more likely than B has to be accompanied by reasons for ruling C out as a possibility. These reasons may consist not in direct evidence but in general considerations deriving from the structure of one's system of beliefs. In other cases the discriminatory capacity that tells As from Bs has to be accompanied by facts that make Cs rarely occur in the circumstances. In yet other cases these factors will be combined. One has evidence that supports the null hypothesis A in contrast to the alternative B, but does not eliminate alternative C, but in the circumstances of enquiry C is not to be found except when something really weird is going on. And in some cases C will be an unevidenced but not unreasonable assumption, as

described in the next section. It is hard to say which of these is the more fundamental element. Deep issues in epistemology arise here.

**assumptions and consequences** Seeking evidence that a coin is fair, you assume it has a constant bias. Trying to tell whether an animal is a dog or a cat, you assume that it belongs to one of your neighbours. In neither case do you have anything like direct evidence, but in both cases, let us suppose, your assumption is a sensible one. What makes it sensible is a controversial matter, as suggested just above. Having made the assumption, you use it in the formation of further beliefs, typically in eliminating alternatives to allow available evidence to get a grip on a situation or in allowing limited discriminatory capacities to operate effectively. You then treat some of these further beliefs as if they were definitely established. You take yourself to know them, in spite of the element of stipulation in their history.

Many epistemological theories will find this troubling. How can knowledge be based on mere assumption, even sane assumptions that in fact are true? From a contrastive point of view the situation is more manageable. You know that it is a dog rather than a cat, though you do not know that it is not a racoon that has wandered far from its usual habitat. You know that the coin is fair rather than strongly biased to heads, though you do not know that it is fair rather than of varying bias. This does not mean that beliefs downstream from any arbitrary assumption which happens to be true can count as knowledge. At the very least it has to be an assumption that is not undermined by other things you know and believe, and it has to be an assumption that you need to make in order for your enquiry to proceed. And at the very least the epistemic grounds for discriminating between the possibilities that, armed with the assumption, you can separate, have to be solid. Of course an illuminating account of when beliefs based on an

assumption are known would be extremely valuable.

But deciding what account of these matters is right cannot be a trivial matter. That can be seen by considering the possibility of Kantian contrasitivism. Human beings assume that they interact with a world of discrete objects located at points in three dimensional space and participating in events in a linear time. They assume that phenomena are explicable in terms of a stable set of knowable laws of nature. And they assume that people make decisions for identifiable reasons stemming from their desires. It is central to Kant's philosophy, particularly to the *Critique of Pure Reason* that we assume these things, and that the assumptions cannot be themselves grounded non circularly in any more basic evidence or experience. (I'm not doing Kant exegesis; he would have put it differently; that was 230 years ago.) One can argue that physics and psychology give us reasons for hesitating over all of these assumptions. One can do so while also arguing that evolutionary theory supports the idea that something like these assumptions are built into human thinking, and that for most humans in most circumstances thinking without taking these things for granted is not an option.

So consider a person who thinks as people normally do and concludes that the match lit because she struck it. Suppose that she is not familiar with any sophisticated reasons for doubting her natural assumptions. Does she have knowledge? Putting the question contrastively we are asking whether she knows that the match lit because she struck it rather than because it lights at random times, or because she willed it to light. We are accepting that she does not know that the match lit because it was struck rather than because its lighting was part of the computer program that gives an appearance of order to her experience, or because a pre-ordained destiny has laid out the

universe in advance, with the striking at time  $t$  and the lighting at  $t + \epsilon$ . (I think that granting that the person does not have knowledge of these contrasting cases is the right course for a contemporary Kantian. If you disagree, rename the position 'contrastive pseudo-Kantianism'. The issues remain as hard.) I think it is obvious that if you take two thoughtful intelligent epistemologists at random and ask whether our person has knowledge of why the match struck simply because of her immediate observations and her hard-wired Kantian equipment, there will be a 50% chance they will disagree. Therefore the issue is not trivial!

Now consider cases like those raised recently in the literature on knowledge and lotteries. You have an appointment with your dentist at 9 am tomorrow, and you are an effective planner who is compulsive about appointments. You plan to go to a movie after the appointment. When the dentist's secretary phones you to make sure you have remembered the appointment, you say "sure, I'll be there, barring nuclear war, hurricanes, or heart attacks". You are not saying to her "if there is no nuclear war. etc., I'll be there" but "I'll be there, and I'm assuming that there will be no war or hurricane and that I will not have a heart attack." You are also assuming that you will not have a traffic accident before 9 am, that your house will not burn down during the night, that no space debris will flatten you, and in fact that none of countless possible preventers will occur. When asked about any of them you will be happy to say that you are assuming they will not happen. And you will be very reluctant, of many of them, to say that you know that they will not occur. For good reason, since you have no evidence that they will not, and many of the factors that would make them occur if they do are random and essentially unknowable. Yet, based on these assumptions, you conclude that you will be in the dentist's waiting room before 9 am, and you do take yourself to know this. (see Hawthorne 2004, especially chapter 1;

the examples in Cohen 2004 have been especially influential.)

I am not going to make any suggestions about the core lottery problem, which is what distinguishes the beliefs that we refrain from calling knowledge even though they follow from things we do know, and why we do know these things even though they are based on unknown assumptions. In a contrastive context these amount to asking why although you know that you will be at the dentist rather than at the movie at 9 am, you do not know that you will be at the dentist rather than in an emergency morgue for victims of space debris. That is a hard question, but in asking it we are also seeing how the contrastive point of view takes the bite out of a skeptical paradox. Although it is not clear why we draw the line between knowledge and ignorance where we do, the fact that we fail to have knowledge of some familiar objects of confidence relative to some contrasts, is quite compatible with out having knowledge relative to other contrasts. You do know that you'll be at the dentist rather than at the movie.

**full contrasts** Ascriptions of non-contrastive knowledge make sense whatever proposition complements the verb. Shakespeare knew that London was in England, he did not know that the 2011 winter Olympics would be in Vancouver, he did not know that  $e^{\pi i} = -1$ , he did not know that  $2+2 = 6$ . We can even stick in a proposition that cannot be expressed in English, call it  $p$  and say that Shakespeare did not know *that*<sup>1</sup>. It is harder to do all this with contrastive knowledge. In particular it is hard to make sense of contrastive attributions with arbitrary contrasts. Did Shakespeare know that London

---

<sup>1</sup> choose a trajectory through space-time, and for each point along it choose the nearest possible world in which there is an electron in that point at that time; take the union of all these worlds. That defines a proposition. A god might be able to think it; no human can. Shakespeare did not know that proposition.

was in England rather than  $3 \times 21 = 64$ ? Did Shakespeare know that London was in England rather than on Alpha Centauri? The last paragraph of the previous section suggested that skeptical concerns are defanged if we distinguish between knowing that you will be at the dentist rather than at the movie, on the one hand, and knowing that you will be at the dentist rather than dead from the impact of random space debris. When you say "I know I'll be there" you mean the first. So assume that you do know that you will be at the dentist rather than at the movies, and do not know that you will be at the dentist rather than at the morgue after a space-debris attack. What follows from this? Is it ruled out that you know that you will be at the dentist rather than at the morgue from dropping in out of curiosity? Is it ruled in, just on general principles rather than as a result of details of your situation, that you know you will be at the dentist rather than at the north pole?

The problem is especially acute if we want to use contrastive knowledge to solve the problem of closure of knowledge under logical consequence. Dretske and Nozick pointed out that one can track  $p$  and not track  $q$ , even though "if  $p$  then  $q$ " is a logical truth. They defended the suggestion that the same is true of knowledge, and claimed that this resolves some issues about skepticism. Later philosophers have tended to disagree, though defence of closure - the principle that, when one knows propositions  $p_1, \dots, p_n$  and sees how to deduce  $q$  from  $\{ p_1, \dots, p_n \}$  one comes to know  $q$  - has proved to be difficult. (See Hawthorne 2004, especially chapter 3, also Luper) My own view is that appealing as the principle is, it is in the end indefensible in full generality, in part because of considerations like those about dentists and morgues. Suppose, though, that one wants to reconstruct closure in terms of contrastive knowledge (see Schaffer 2007.) This might be attractive because some putative counterexamples can be defused with well-placed contrasts. Notoriously, you know that you have two hands, and although



having two hands entails not being a brain in a vat, you do not know you are not a brain in a vat. But if we qualify the premise to “you know that you have two hands rather than two stumps”, the entailment to “you know that you are not a brain in a vat rather than having two stumps”, does not seem clearly false.

But it does not seem clearly false because it is so strange that we do not know how to evaluate it. Logical consequence can connect sentences which have so little intuitive connection with one another that they wreak havoc with sensible contrasts. Anyone trying to put together contrastivism and closure will need an attitude to assigning truth values to some very unfamiliar objects.

Even if we decide that clarifying epistemic vocabulary is not alone going to solve problems about closure, we still must face questions about truth values given unrelated contrast propositions. I do not think these questions are insoluble. I propose three principles for handling the issue.

First, the contrast proposition is always false and incompatible with the known proposition. If I know that  $p$  rather than  $q$ ,  $q$  is an alternative to  $p$ , and since  $p$  is true  $q$  is false. So we can rule out all of the following: Shakespeare knew that London was in England rather than Stratford was in England, Shakespeare knew that London was in England rather than  $2+2 = 4$ , Shakespeare knew that London was in England rather than cats chase mice. The list is easy to extend. In this connection it is worth pointing to an ambiguity. Sometimes in saying “Shakespeare knew that London was in England rather than Stratford was in England” we might mean something true. That would be when Shakespeare knew that London was in England, did not know that Stratford was in England, and someone had mistakenly asserted the latter. Then one might correct them by saying “No, it was

London that he knew was in England, not (rather than) Stratford." But in saying this one is not ascribing contrastive knowledge in the present sense.

Second, the person has to have discriminated the known proposition from the contrast proposition, by either a perceptual capacity, effective use of evidence, or reliable reasoning. So Shakespeare did not know that London was in England rather than London within the green belt is in England. Shakespeare did not know that the city from which Elizabeth ruled was in England rather than that the city which Ken Livingston would rule would be in England. Shakespeare did not know that uranium has two isotopes rather than six. He did not know that kings are to be obeyed because of divine command rather than for civic peace. And so on.

Third, we should look for systematic patterns of ascription and denial. Shakespeare knew that London was in England rather than in France, and also that London was in England rather than in Spain, Italy, or Illyria. He did not know that London was in England rather than being in the English sphere of influence or being a British crown dependency. Each of these lists can be continued because the initial contrastive ascription derives from the width of accuracy of Shakespeare's city-identification and city-location capacities, and these capacities make many knowledge contrasts hold and also many fail. When we cannot find such a systematic contrast space we should suspect that the ascription is false.

When these principles do not suggest that someone has contrastive knowledge then most likely they do not. So the vast majority of random ascriptions of contrastive knowledge are false, just as the vast majority of non-contrastive knowledge claims are. There are many ascriptions that are not settled. But that is as it should be: they have to be settled by data

about the particular cases and by an informative epistemology.<sup>2</sup>

**summing up: how fundamental?** The reason we attribute knowledge is very straightforward. We have reasons to be curious about what aspects of the world creatures have sufficiently accurate information about to guide their actions. Once we have determined this, we can use their actions, including their utterances if they are verbal creatures, as a guide to ours. But what is sufficient to guide one action may not be sufficient to guide another. Tracking-based information is particularly versatile in the variety of actions it can support, but it has its limits: we track some aspects of objects through some possible histories and not through others. As a result, we need a way of relating individuals to facts that does not pretend that the informational stream is wider than it is. So we do two things together. We relate individual agents to the objects they act on in a way that describes the stream of information - the set of action-guiding counterfactuals - that is relevant to actions towards those objects. And we describe the 'width' of those streams, the range of similar situations to which an explanation appealing to the same information would apply.

We identify the stream of information by saying what agents know. And we describe its limits by saying what contrasts their epistemic states will support. But the knowledge attribution alone, with no contrast specified, can go some way to describing the limits. This is clearest in cases where there is a close connection with tracking. Then the fact that *a* knows that *p* entails that *a* would believe *p* under a variety of similar conditions. But only in similar conditions, nearby possible worlds: there is no suggestion that if things had been more than a little different *a* would have got a true belief. So the width is vaguely specified by the range of situations *S* in which the

---

<sup>2</sup> Joel Buenting has made me appreciate the importance of the issues in this section.

fact that *a* actually knows that *p* entails that *a* has an accurate belief in *S*. Sometimes an explanation uses an ascription of knowledge in a proposition *p* to specify a flow of information, and in a rough way its limits, which is then used to explain an action which could be described independently of *p*. *p* serves only to pick out the information flow. For example we can explain why the police managed to arrest a wanted fugitive by the fact that there was a tracking device in a stolen car which he happened to be driving. They knew where the car was, and so they could arrest him. They didn't know where he was.

In this connection it is worth pointing out that explanations by knowledge do not just appeal to knowledge of single propositions. "How could she find you so quickly? Because she knows that forest very well." "Why did Ossie get lost although he knew where the road was? Because he knew that it was north rather than south, but not that it was ten miles north rather than five." "Why did Petra succeed in getting the plan approved? Because she knew who to bribe." The crucial phrases here are "knows the forest", "knows where ..", "knows who ...". These are all kinds of knowledge that do not centre on a particular fact. Instead they centre on a general kind of information possession, a set of counterfactuals, describing the flow of information without singling out one particular origin for it. These less focused knowledge attributions clearly need specifications of their width. The person who knows a forest very well does not know the location of every mushroom under every tree. So when trying to give the explanation more carefully we say "she knows locations in the forests very well". And in fact we will say "she knows the geography rather than the ecology of the forest very well."

The upshot is that the explanatory work is done by describing systematic information-links between individuals and their environments. We specify

these by describing central cases - usually with a "knows" locution - and by describing the limits of the links, usually with contrastive locutions. Sometimes, as I have pointed out, the central cases are not the information on which the action being explained is based. Especially then, we need to determine whether the action on which it is based lies within the limits of the link. Contrastivity is one device for doing this. Often we do not do it contrastively, but leave it to context. A good example of this is "knowing who". Notoriously, one can know who someone is, given the demands of one context, but not given the demands of another. The police may know that the motorist they have pulled over is Jane Doe, because that is what her license truly says, without knowing that she is the notorious graffiti artist wanted on three continents. In one sense they know who they have stopped, and in another sense they do not. We can rescue the attribution from context in many ways. One is contrastive: they know that she is Jane Doe rather than Mary Moos (or ...), but do not know that she is the wanted artist rather than the harmless commuter. They know that they have stopped the transgressor of a minor traffic offence now rather than a major traffic offence now, but do not know that they have stopped a traffic offender rather than someone wanted by interpol. We can also use terms that are not explicitly contrastive: they know which citizen they have stopped, but not which criminal. (The awareness of the ambiguity of "knows who" dates to Kaplan 1968; see also Boer and Lycan 1985.)

Contrastive knowledge ascriptions give us information about information links that is essential to using them for explanatory purposes. There are other ways of giving us the information, other explicit linguistic devices and dependence on general contextual inference. There are always other ways of saying things. (Most of the time we can avoid using "knows" if we really want to.) But that does not prevent the information being essential to epistemically based explanations.

Adam Morton  
 University of Alberta  
 adam.morton@ualberta.ca

## References

- Boer, Steven, and William G. Lycan (1985) *Knowing Who* (Bradford Books: MIT Press)
- Cohen, Stewart (2004) Knowledge, Assertion, and Practical Reasoning, *Philosophical Issues* Volume 14 Issue 1, Pages 482 - 491
- Dretske, Fred (1981) *Knowledge and the flow of information* Cambridge, Mass. : MIT Press.
- Hawthorne, John (2004) *Knowledge and lotteries* Oxford : Clarendon ; New York : Oxford University Press
- Kaplan, David (1968) Quantifying in *Synthese* 19, 178-214
- Karjalainen, Antti, and Morton, Adam (2003) Contrastive knowledge *Philosophical Explorations*, **6 (2)**, 74-89
- Luper, Steven The Epistemic Closure Principle  
<http://plato.stanford.edu/entries/closure-epistemic/>
- Morton, Adam (forthcoming a) Contrastive knowledge, in Duncan Pritchard, ed. *The Routledge Companion to Epistemology*.
- Morton, Adam (forthcoming b) From tracking relations to propositional attitudes. *European Journal of Philosophy*
- Roush, Sherrilyn (2005) *Tracking truth : knowledge, evidence, and science* Oxford : Oxford University Press, 2005.
- Schaffer, Jonathan (2005) "Contrastive Knowledge," in *Oxford Studies in Epistemology 1*, eds. Gendler and Hawthorne, 235-71, Oxford: Oxford

University Press

Schaffer, Jonathan (2007) Closure, Contrast, and Answer *Philosophical Studies* 133.2, 233-55

Schaffer, Jonathan (2005) Contrastive Causation *The Philosophical Review*, **114**, No. 3 297-328