

Contrastive knowledge

Philosophical Explorations 6,2, 2003, pp. 74-89.

Antti Karjalainen and Adam Morton

Abstract

We describe the three place relation of contrastive knowledge, which holds between a person, a target proposition, and a contrasting proposition. The person knows that *p* rather than that *q*. We argue for three claims about this relation. (a) Many common sense and philosophical ascriptions of knowledge can be understood in terms of it. (b) Its application is subject to fewer complications than non-contrastive knowledge is. (c) It applies over a wide range of human and non-human cases.

The aim of this paper is to argue for the intelligibility and usefulness of a concept. That concept is *contrastive knowledge*: person *S* knows that *p* rather than that *q*. Our central claim for contrastive knowledge is that it combines both robustness and stringency.

That is, it both applies unproblematically in a wide range of cases, and it involves the kinds of standards that make it qualify as knowledge.

We do not intend this paper to make a water-tight or conclusive case for any particular claims about contrastive knowledge, as much as to make an attractive presentation of the concept: here it is, please pay it some attention.

Since we are claiming robustness, we set the stage with a non-human example.

Consider Reilu, a sniffer dog at Helsinki airport. She is very good at her job. She can detect the slightest trace of a number of illicit substances in luggage and clothing, and she can tell them from tobacco, Vodka, Aquavit, Lakka, and a vast range of other similar or distracting smells. Of course she has her limitations. She cannot tell real pot from synthesized marijuana smell: she'll bark and thrash her tail given a cotton swab that smells right even if it has no other characteristic of the forbidden stuff. Reilu is malleable as well as discriminating: if her handlers prime her suitably she can on a particular occasion seek out cocaine or seek out explosives, or any of a number of other substances, ignoring ones that she is sensitive to but has not been set to find.

Reilu knows that a suitcase contains cocaine rather than powdered plastic explosive. (It will take the lab hours to confirm this. But if set to find cocaine she will single out the case, and if set to find explosives she will ignore it.) She does not know that it contains cocaine rather than a synthetic alkaloid that smells the same. She does not know that there is cocaine in the suitcase rather than minute traces of cocaine being carried by air currents in such a way that they move

towards her nose in just the way traces from a cocaine-containing suitcase would. But what she has is a kind of knowledge: an important kind, since her expertise is irreplaceable. One of her handlers asks 'might that box with no labels on it contain explosives?' And another replies 'Send Reilu over; she knows explosives from fishing tackle.'

Now a human case. A person looks out of a window and sees a robin on the ledge. He thinks 'there's a robin.' His eyes are good, and the light is fine, but his ornithology is limited. He couldn't tell a robin from an oriole or a towhee. Still, he knows where the robin is. He knows *where* the robin is, though he does not know *what* it is. Take his knowing where the robin is to be based on his knowing that it is on the ledge rather than on the tree or the lawn or the television aerial.

So he knows that the robin is in one location rather than many others though he does not know that it is a robin rather than an oriole. So he does not have non-contrastive knowledge that the robin is on the ledge; for to know that he would have to know that it is a robin. (Even more clearly, he does not know that there is a robin on the ledge, though he does know that there is a robin on the ledge rather than on the tree. He does know that there is a robin on the ledge rather than a crow, but he does not know that there is a robin rather than an oriole.)

Cases with a *de re* aspect often are best described in contrastive terms. What is reliable about the person's belief is *of* the robin, even though his labeling of it as such is right partly by accident. So there is something knowledge-like about his state, although it is less than full knowledge. In ordinary language we reflect this fact by saying that he knows *where* the robin is, just as he knows *when* it landed on the ledge, and does not know *what kind* of bird it is. We can also say that he knows where there is a robin, etc. Note, though, that if he had thought that the robin was an oriole but had tracked its location carefully, though he might have known of it that it was on the ledge and not on the tree he could not have been said to know that there was a robin on the ledge rather than on the tree. (Compare: if Reilu gives her 'marijuana!' reaction to explosives in a case, and in fact classifies the contents mentally the way she does marijuana, then she does not know that there are explosives in that case rather than another.) So contrastive knowledge is not just the *de re* version of regular non-contrastive knowledge. (All the combinations of contrastivity and *de re/de dicto* are possible: you can know that there is a robin rather than an oriole on the ledge, but not know that there is a robin on the ledge rather than levitating a micron above it; you can know that there is a robin on the ledge rather than on the tree, but not know that there is a robin on the ledge rather than an oriole; you can know of the robin that it is on the ledge rather than on the tree but not know of the robin that it is a robin rather than an oriole; and so on.) But the range of ways in which an epistemic process can give reliable information, which can then link to some aspect of a believed proposition, is wider than the range of *wh-* idioms that we happen to have. So if you accept that people often know some but not all of where, when, why, who, what, and so on, you are well on the way to understanding that a person's epistemic situation is often well described as knowing that *p* rather than that *q*.

We return to *wh*-knowledge below. The important point now, while we are convincing you that you can almost grasp the concept we are trying to evoke is the similarities between the bird-observer and Reilu. He is in a situation in which the processes he is using can reliably distinguish between several possibilities – the bird's being on the ledge and its being in various other places – and when they do fasten on one of a set of distinguishable possibilities the one that they select will be the case. Similarly Reilu knows that it is marijuana rather than tobacco in a suitcase because her innate capacities have been trained to make that discrimination, and the discrimination is robust in her actual circumstances, in spite of the fact that had synthesized marijuana-smell been present she would have responded to it the same way. She has this contrastive knowledge because she can discriminate a range of possibilities and she is in a situation where this capacity operates well and guarantees that the selected possibility is actual.

Do not take any of this as a definition! There may be ascriptions of contrastive knowledge that do not fall into any similar pattern. The aim is to make it plausible that there is an intelligible and useful idea here. (After all, no one would now begin a paper on knowledge by simply laying down necessary and sufficient conditions for 'S knows that p' and then going on with the assurance that what these conditions picked out was knowledge.) In the rest of the paper we discuss a series of other natural habitats for contrastive knowledge, cases where we readily say that someone knows that P rather than that Q or that a person knows whether P or Q, or can tell P from Q. We return only at the very end to the question of whether there is a single general characterization of it.

Before going on to these habitats it may be as well to mention some confusions to avoid. Keep these cordoned off in your mind and what we say below will be clear. First, our concerns are different from Fred Dretske's, though contrastivity is central to both. Dretske is concerned with full one-propositional knowledge and has brought to philosophers' attention cases in which the reasons that a person knows that p involve a contrast between the person's relation to p and to some other propositions.¹ The relevance of these contrasts varies from one context to another – the possibility that someone has been painting stripes on mules is important for some attributions of 'knows that there is a zebra in the paddock' and not for others – so that hard questions of the context relativity of knowledge ascriptions inevitably arise. But our concern is different. We are not concerned with the contrastive aspect of 'knows that p' but with a different, inherently contrastive and less context-dependent concept, 'knows that p rather than that q'.

There are many connections between the two: the person who does not know that the animal before her is a zebra rather than a painted mule may or may not know that it is a zebra, depending on factors that are currently a matter of intense controversy. If we get into this controversy the position we take on it will obscure our main point.

1 See especially Dretske (1981), and more recently, with Zebras, Dretske (2000).

To keep contrastive knowledge separate from non-contrastive knowledge, one has to avoid running together 'knows that p rather than that q' with 'knows that (p rather than q)', or 'knows that p rather than knows that q'. Both of these latter are complex non-contrastive knowledge. The distinction one has to make is similar to that made in the literature on contrastive explanation. Explaining why a ball reflected at angle θ rather than angle μ is the same neither as explaining why the ball reflected at θ rather than explaining why it reflected at μ , nor as explaining why the ball reflected at (θ rather than μ). If the second of these is taken literally, as a non-contrastive explanation, it would often be unavailable when the contrastive explanation was straightforward: for example when simple mechanics says why the angle of reflection equals the angle of incidence, but does not say why the ball reflects at all rather than breaking through the wall, or disintegrating. Similarly 'knows that (there is a robin on the ledge rather than on the tree)' will often fail, taken literally, of a person who knows where the robin is but does not know that it is a robin.²

Of course very often we do not take such ascriptions literally; we have more sense, given the ascriptions that a practical situation requires. When we say 'he knows that there is a robin on the ledge rather than the tree' out loud we stress the 'ledge' slightly, to make it clear that we are saying 'he knows that there is a robin on the ledge rather than on the tree, but some other contrast may fail'. If we want to be even more explicit we can say 'he knows that there is a robin on the ledge rather than on the tree, but he does not know that it is a *robin* on the ledge rather than an oriole.' Or 'he knows that there is a robin on the ledge rather than on the tree, but he does not know that it is on the ledge rather than on a transparent pane just beside it.' Rather than making a richly context-dependent ascription we make a weaker ascription; we aim at a different target.³ Keep that target in mind – the bird observer knowing that the robin is on the ledge rather than the tree, though not knowing that it is a robin rather than an oriole, Reilu knowing that it is cocaine rather than explosives, though not knowing that it is cocaine rather than synthesized cocaine odor – and the concept should stay fixed in your mind.

- 2 When ascribing contrastive knowledge in English one often truncates the second proposition and drops the that: she knows that her pen is in her bag rather than on the table. However our official formula is "knows that p rather than that q" in order to signal the distinctions made in this paragraph. Thanks to a referee on this point.
- 3 Someone might speculate that whenever one wants to say something with the force of 'a knows p rather than q' one can say just 'a knows p' with a stress which in context brings out the intended meaning. If this were so then a dyadic knowledge relation plus rules of conversation would be enough. Even were this so, in order to say independently of the conversational context what the content of the assertion is, one needs a way of making the contrastivity explicit. We have a host of devices for doing this. We can say 'She knows he was *there*', 'She knows *he* was there', 'She knows that it was him there', 'She knows that it was there that he was', and also, when we want to be clear that it is none of these 'She knows this: he was there.'

First habitat: questions and wh- knowledge

When the contrast concerns time, location, or agent, we often express contrastive knowledge with wh- idioms. We say that someone knows where and when something happened, or who made it happen. This time, place, or agent, rather than those other ones. Wh- knowledge is usually more robust than simple knowledge. Looking at a bird on a branch you can know where the robin is, even if it is unclear whether your ornithological expertise is adequate for your knowing that there is a robin on the branch. Language gives us a fixed list of wh- idioms. But we can express an enormous number of contrasts in terms of them, since we can say 'which P' for any category of predicates. We can say that someone knows which make of carburettor his car has, or even which type of Ford carburettor it is. We can say that someone knows which cat ate his budgie, or even which female Siamese cat did the deed, or even which of Gretchen and Angelica did it.

In one respect attributions of contrastive knowledge are simpler than those of wh- knowledge. Wh- knowledge nearly always involves parameters that are set conversationally in subtle ways. If you see that one masked man rather than any of twelve others fired the fatal shot then for the purposes of one conversation you may be said to know who fired the shot. But for the purposes of another conversation you may not be, since you cannot produce a name or an address. (Similarly with knowing where, when, and the rest.) This relativity does not apply to your knowledge that the second masked man from the left rather than any of the others fired the shot. He rather than any of them did it, but not he rather than the president of France or the first horn of the Berlin Philharmonic, who for all you know may be behind the masks. Though we ascribe contrastive knowledge when we make a wh- knowledge attribution, what contrastive knowledge we ascribe will depend on context.

Someone might argue that wh- knowledge claims have presuppositions that contrastive knowledge claims lack. To say that someone knows which cat ate his budgie might seem to presuppose that this person knows that some cat ate his budgie. On the other hand to say that someone knows that some cat rather than some other cat ate their budgie is not, to say that they know, non-contrastively, that some cat ate their budgie. (For they might not have adequate grounds for knowing that the animal is a cat.) But in fact it does not seem that the wh- idiom has this presupposition. If someone knows which cat ate their budgie then they *believe*, truly, that some cat ate their budgie, and they have the *contrastive* knowledge that cat A rather than cats B, C, D ... is the murderer. They may not have literally construed uncontrastive knowledge that (cat A and not cat B killed the budgie) and thus their belief that the murderer was a cat may not be knowledge.

Wh- knowledge connects with wh-questions. When someone knows where/who/which they can normally give a non-accidentally correct answer to a corresponding question 'Where/who/which...?'. The link transfers to contrastive knowledge. Looking at the branch you can answer "where is the robin?"; in fact you know the answer to it. But if your bird-recognition skills are very rudimentary any correct answer you give to 'what kind of bird is that on the window

ledge?’ is going to be to some extent right by accident. Or, to put the point more cautiously, it is quite possible to know the answer to ‘where is the bird?’ and not know the answer to ‘what kind of bird is it?’ even though one has true beliefs on both counts. This suggests some equivalences, which though imprecise help pin down contrastive knowledge within a larger circle of concepts. To state them we adapt the fact/foil terminology from the literature on contrastive explanation.⁴

When someone knows that *p* rather than that *q* then *p* is the fact and *q* is the foil. Wh- knowledge invokes fact and foil too. For example if someone knows where a thing is, the foil is the set of propositions ‘*a* is at *l*’ where ‘*a*’ and ‘*l*’ are suitable designators of the thing and of a location, and the fact is a true proposition ‘*a** is at *l**’ believed by the person.⁵ And so too with questions. The foil of a question as asked in a context is the set of relevant answers, and the fact is the right answer. These are purely stipulative definitions, but with them in place we can state some conjectures which would make them non arbitrary.

Claim one: When a person has wh-knowledge, she knows that *p* rather than that *q*, where *p* is the fact and *q* is a subset of the foil of the wh-knowledge.

Claim two: When a person has wh-knowledge, there is a question and a suitable context such that the fact and foil of the wh-knowledge are the fact and foil of the question in the context.⁶

Claim three: When a person knows that *p* rather than that *q*, *p* is the answer to some question whose foil is *q*.

Or, to put it very simply: *wh-knowledge constructions attribute contrastive knowledge, as do attributions of correct answers to wh-questions.*

Second habitat: *facticity*

Many propositional attitudes can only be directed at true propositions. And, as several writers have pointed out, many of these attitudes can only be directed at propositions that are not only true but known. If you see that the bird is on the window ledge you know that the bird is on the window ledge. If you remember that Helsinki was founded in 1550 then you know that Helsinki was founded in 1550.⁷ These entailments constrain the interpretation of the corresponding

4 See Lipton (1991) pp 38-42. Also Garfinkel (1980), chapter 5 of van Fraassen (1980), and Temple (1988). For an early version see Hanson (1958) p 54.

5 It is not at all trivial to specify what makes a suitable designator. In fact, most of the literature on knowing wh- grapples with this question rather than with directly epistemological issues. Christopher Hookway has pointed out to us that the designators are not always explicitly singular, even when it is knowing who or what. ‘Do you know who told them?’ ‘Yes, an accountant.’ One could try to reduce knowing wh to knowing that. In effect this is the project of Boër and Lycan (1986). See also Hintikka (1992). It is very controversial whether this can succeed; from our point of view this would be like trying to reduce contrastive explanation to non-contrastive, a project that has never succeeded, or to reducing the definite article to quantifiers and identity, a project that in succeeding raises more serious questions.

6 See Hookway and Grimaltos (1995)

7 See Gordon (1988), Williamson (1995).

propositional attitudes. But they also constrain the interpretation of 'know'. We have to understand 'know' in such a way that the entailments are preserved.

There is a danger here that 'know' will be pulled in different directions, for the various attitudes that entail knowledge differ in fundamental respects: can all the facts and purposes that shape their interpretation be satisfied with a single interpretation of 'know'? For example, the *see/know* link suggests that typical instances of seeing will also be instances of knowing. But suppose that the seeing part of the process is absolutely normal while other cognitive elements tend against knowledge? Suppose for example that you have been dreaming about robins all night and you are predisposed to see almost any bird as a robin. Then you do not know that the bird on the ledge is a robin. You do not even know that the robin is on the ledge, rather than levitating a millimetre above it. (So a skeptic who took this as grounds for denying that you know that the robin is on the ledge, and who accepts the facticity thesis, should have to conclude that you are blind!) But not only do you see the robin, you see that the robin is on the ledge.

Related examples turn on misidentification. You can see the robin while thinking it is an oriole. So the entailment from 'sees that the robin is on the ledge' to 'knows that the robin is on the ledge' has to be understood carefully. We would often say that you see that the robin is on the ledge even if you had failed to identify it as a robin. Similarly, you can remember that the keys to the Citroën are in your bag even though you think that it's a Peugeot. (In both kinds of case there is also a bias of the contrastive reading to a *de dicto* rather than *de re* reading. We return to this in the last section of the paper.)

Examples like these suggest that the entailments in question are not unrestrictedly true. Weaker entailments involving knowledge *wh*- seem more generally true. If you see a robin, or see that it is on the ledge, then you generally know where it is. If you remember an event's occurrence, or remember that it occurred, then you know when (in some sequence of events) it happened. But, as we argued above, *wh*-knowledge attributions are just a contextually sensitive way of attributing contrastive knowledge. If a person sees a robin then she generally knows that it is in one of a class of locations rather than another. A full statement of the contrast would have to build in the inherent imprecision of visual localization, as we pointed out above.⁸ Each knowledge-factive idiom generates a complex contrast space, peculiar to it, which speakers come to understand implicitly and then usually express with a handy inaccurate shorthand of 'know' and 'know *wh*'.⁹

8 In fact, a full statement of the contrast-space appropriate to seeing would be very subtle and complicated. Someone who sees their true love in a hall of mirrors may only in the most contextually delicate sense be said to know where their true love is. What they know is that the person is at a location got to in a direct line starting in a specific direction and changing direction appropriately as it meets each of an unknown number of mirrors, rather than some location that cannot be so characterised. Thanks to Christopher Hookway for the example.

Third habitat: *the force of evidence*

The better your evidence the stronger your claim to knowledge. But evidential support is a complex matter. It is misleading to describe it by analogy to a deductive argument, starting from evidential ‘premises’ and following a path of reasoning to justify more or less strongly a conclusion. At any rate it is misleading if we think of evidence in traditional statistical terms. Statisticians rarely collect a mass of data and then announce that some hypothesis that happens to fit the data is supported to some specific degree. Instead, a standard practice begins with a hypothesis to be tested and some data that are consistent with it. We then have to consider whether the data are the result of the processes postulated by the hypotheses or have some other origin. In the very simplest case one considers the alternative, ‘null’, hypothesis that the data are produced by some random factors which have an assumed distribution. Statistical techniques enter at this point, to estimate how likely it is that the data were produced by the null hypothesis, rather than by the hypothesis being tested. Suppose that the conclusion is that the likelihood that the data are produced by random factors is very small. Then we can take ourselves to have evidence that the hypothesis *rather than* the alternative accounts for the data. And, all going very well, this may be part of our grounds for accepting the hypothesis.¹⁰

Suppose this is right, and that evidence is generally contrastive. This certainly does not show that knowledge has to be contrastive. For one thing, it is a subtle and controversial matter what has to be added to truth and adequate evidence to produce knowledge. And of course we are not claiming that all knowledge is contrastive. What it does show is that a central component of the concept of knowledge, that of support by adequate evidence, is often contrastive. When true beliefs are well supported by contrastive evidence, and the contrast has definite limits, we often want to call them knowledge. But it is knowledge with a proviso – that some possibilities have not been excluded. One way of capturing this fact is to let the knowledge ascription inherit the contrastivity of the evidence. For example Hellenistic astronomers had evidence that discriminates a ball-shaped earth from a flat one. They did not have evidence that discriminates a ball-shaped earth from one that is spherical except for indentations at the poles. So a Hellenistic astronomer who believed the earth to be a ball would have had a true belief, which his evidence partially discriminated from some but not all false alternatives. Most of us would be reluctant to say that the astronomer knew that the earth was ball-shaped rather than ball-shaped-with-indentations, for these

9 Many factive idioms have specific contrastive variants. See where, understand how, tell where, remember how, when, where. In fact, some non-factive idioms have factive variants which link to knowledge via a wh-tag. For example ‘tell’ is non-factive: you can tell lies or speak from ignorance. But ‘tell where’ is factive. If you tell someone where the gold is hidden then the gold is hidden there, and in fact you can’t tell someone where the gold is unless you know that it is there. One reason we don’t say ‘believes where/why’ is that the result would be factive, and thus no longer contrast with ‘knows where/why’.

10 See Hacking (1965), Mayo (1996), and almost any statistics text. Chapters 8 to 10 of Spiegel (1961) will do nicely.

reasons. Under suitable conditions most of us would be happy to say that he knew that the earth was a ball rather than flat. Would we say that he knew, simpliciter, that the earth is a ball? Some would and some would not, and many would change their tune from one conversation to another. But all would agree that the astronomer is in a state much like knowledge, which is based on evidence that discriminates some but not all alternatives.

Considerations like these will not do more than reveal a gap into which contrastive knowledge can be fitted. They will not show that this is the only way of filling the gap. They also support the idea in a different way. If evidence were not ever contrastive, then it is hard to see how beliefs whose claim to knowledge consists in their relation to evidence could ever be contrastive. (The contrasts would have had to appear miraculously from a combination of non-contrastive evidential force and non-contrastive belief.) But evidence is contrastive, so there is no objection from this quarter.

The contrastivity of evidence remains when we think in terms of more sophisticated statistics. (But if you don't want to know why, the first paragraph of this section gives the point: statistical evidence usually favours one hypothesis rather than another, but not necessarily rather than a third.) Sometimes the hypothesis to be tested is itself probabilistic and is a source of the probabilities used to estimate likelihoods. Sometimes we make much more general null hypotheses, in particular loosening assumptions about the distribution of alternative causes. But none of these complications affect the main points. The hypothesis is tested against alternatives. The alternatives represent only some of the ways the hypothesis could be wrong, being chosen for inherent plausibility and statistical docility. And the result is not a measure of the absolute confirmation of the hypothesis itself but of its relative confirmation in a context in which the alternatives are its appropriate rivals. As long as these points hold, the data provide evidence that the hypothesis rather than any of a set of specific alternative accounts for the data.

The greatest threat to this attitude comes from Bayesian statistics. For on the official Bayesian presentation the effect of evidence is to change the absolute degree of belief of a hypothesis, from which its standing relative to alternatives is derivative. And in fact for official Bayesianism the standing of a hypothesis relative to every conceivable alternative is well defined, so that the contrast implicit in 'rather than' becomes vacuous.¹¹ But in real Bayesian life this orthodoxy is considerably relaxed. It has to be, given the practical tensions for mere human beings between two Bayesian axioms. The first is the core idea that you need to start from somewhere. You need prior probabilities for every hypothesis to be supported or undermined. The second, practically conflicting idea is that you must have a single probability assignment over all the propositions you can conceive.

No gaps, no double values. So we must have prior probabilities for every hypothesis we are considering, and we must have the same prior probabilities for

¹¹ See Howson and Urbach (1989). A simplified introduction is in chapter 9 of Morton (2002).

every such situation. So to give probabilities of any set of incompatible but not impossible hypotheses (the moon is made of Cheddar, the moon is made of Gruyere, the moon is made of Camembert, ...) we will have to make very fine distinctions between vanishingly small numbers. (It is actually slightly more probable that the moon is made of Cheddar than that it is made of Camembert, given its stability under gravitational forces.) And we will have to have assignments of probability to hypotheses that we have never seriously considered.

In practice no Bayesian lives up to these ideas. The quantifiers are taken the other way round: in every evidence-assessing situation there is an assignment of prior probabilities to hypotheses that makes sense in the situation. (Plus a rule of thumb that it's a bad idea to vary the assignments too drastically from occasion to occasion.) So on any occasion we will use a partial assignment which gives no probabilities at all to many hypotheses that are irrelevant to the problem at hand, and which gives zero probability to many hypotheses which are not strictly impossible but whose likelihood can be disregarded for practical purposes. Only in this way can we assess the evidence for things we are really interested in without having to struggle through the complications induced by non-zero probabilities for brains in vats, Martian abductions, and angelic interference. And only in this way can we produce the prior probabilities we need to assess the evidence for totally new hypotheses that have never been considered before.¹²

Seen this way, the difference between the Bayesian and traditional attitudes to statistical evidence is not so great. On both attitudes when we evaluate the evidence for a hypothesis, we must take it to be one of a set of hypotheses that are serious contenders. The set leaves out hypotheses we are not taking seriously and hypotheses we have not yet thought properly about. What we learn from analysing the evidence is that it does give reason for preferring one hypothesis to one or several alternatives. We will then have reasons for rejecting sensible skeptical worries.

(Perhaps the subjects are picking up cues from the experimenter's tone of voice"; "Perhaps some intergalactic dust is absorbing the light'.) But uninteresting, irrelevant, unserious, or unorthodox alternatives need not be ruled out. This is also one reason for the tremendous power of hypotheses formulated in quantitative terms. If a hypothesis says that a parameter has a specific value, or that a function has a specific form, then it is easy to formulate sets of alternatives to it.¹³

At a later more sophisticated, or perhaps more rhetorical, stage we may be able to convert the evidence-based discrimination between hypotheses into a degree of support for a single hypothesis. We may be able to say that it is reasonable to believe *H*, or even that *H* is now well established. It is obviously very controversial how often when we do this we have simply decided to ignore the many alternative hypotheses between which and *H* no available evidence discriminates. The important point now is not whether this is ever not the case, but just that any such

12 See chapters 9 and 10 of Williamson (2000).

13 See Glymour (1980), chapter 5, where the emphasis is on testing one claim in a theory rather than another, when a quantitative formulation allows their contributions to be discriminated.

stage is based on a prior assessment in which we establish that there is evidence for believing *H rather than H'*.

Fourth habitat: *explanation*

Very often we hold a belief because of its explanatory power. Explanation is often thought to be contrastive: we explain why this rather than that happened, sometimes by appeal to the fact that one fact rather than another is the case.¹⁴ (In fact to explain is, intuitively, to have or transmit a particular form of *wh*-knowledge, knowing why.) So, if knowledge sometimes arises from an inference to the best explanation and if explanation is sometimes contrastive, then it is attractive to suppose that our knowledge of an explanatory hypothesis should often be contrastive.

Or, to emphasize the aspect that we argue for below, limits of the contrastive width of our knowledge of the fact explained should produce limits on the contrastive width of any explanations of it, and thus of our knowledge of the explanatory facts.¹⁵

Now it must be said right away that on some standard accounts of contrastive explanation it is perfectly possible to have contrastive explanation without contrastive knowledge.¹⁶ In particular, consider the account due to David Lewis as reformulated by Peter Lipton.¹⁷ According to this account an explanation says why an event *e* occurred rather than another possible event *e** when it specifies a cause of *e* and the non-occurrence of any event that would be related to *e** as the cause of *e* is to *e*. Contrastive explanation, on this account, selects among all the causes of an event those whose influence makes the difference between that effect rather than another. That the *x* key was pressed explains why the character *x* rather than the character *y* appears on the screen. For if *y* had appeared the causal history would not have involved the pressing of any key that actually was pressed. That the *x* key was pressed does not explain why the character *x* appeared rather than the screen's going blank. For that could have happened as a result of a computer fault that had no connection with any key's being pressed, or as a result of a memory overload which could have been caused by the pressing of any key.

There is no contradiction, on this account, in saying that someone knows that *c* has occurred, knows that it explains why *e* rather than *e'*, and knows that *e* has

14 See the works referred to in footnote 1.

15 Of course, one can argue that the inference to best explanation does not give knowledge, following van Fraassen in (1989). But van Fraassen's arguments can be read as arguing instead for a particular interpretation of the contrast space of knowledge resulting from IBE. IBE, on this interpretation, in telling us that a hypothesis may be taken as empirically adequate, gives us knowledge that this objective situation rather than any alternative with different empirical consequences is the case. (But not rather than empirically indistinguishable alternatives.)

16 This point was made clear to us by Peter Lipton.

17 See Lipton (1991) p. 42. Neither Lipton nor we would intend this to be precise, let alone a definition of contrastive explanation. Lipton alludes to John Stuart Mill's Method of Difference, "a cause must lie among the antecedent differences between a case where the effect occurs and an otherwise similar case where it does not".

occurred, but that all three items of knowledge are non-contrastive. In particular, an epistemic agent that could discriminate all possible worlds, and whose knowledge was thus entirely non-contrastive, could give contrastive explanations. Such an agent would be able to know that *c* is central to the causal story of why *e* rather than *e'*, and not central to the story of why *e* rather than *e''*. But of course we humans fall very far short of being about to discriminate all possible worlds. So the important question is whether the explanatory knowledge of creatures with limited discriminatory powers has to be contrastive.

An immediate connection to make is that if one can explain why one event occurred rather than another then one is usually in a position to know that the latter did not occur. (If I know that the reason the *x* rather than the *y* appeared on the screen was that I pressed the *x* rather than the *y* key, then I am in a position to know that the *y* did not occur.) So if it is impossible for a person to know that *e* rather than *e'* then she cannot explain why *e* rather than *e'*. (If my key-pressing explains why an *x* rather than a *y* appeared, and I know which key I pressed, then even if my sight is too feeble for me to tell by sight which key appeared, I can tell in terms of my key-pressing.) So if there are possibilities that human beings are incapable of discriminating from actuality, we can expect all our explanations to inherit a contrastivity that respects this limitation. A more general claim along the same lines would be that when one does not know that *e* rather than *e'* one cannot explain why *e* rather than *e'*. And to the extent that inference to the best explanation is the source of beliefs of the kind that are doing the explaining, one would not know facts which explain why *e* rather than *e'*. So one might suspect that when one does not know that *e* rather than *e'* one cannot know that *c* rather than *c'*, where *c* explains *e* and *c'* would if true explain *e'*. The suspicion can be sustained, if we make a couple of assumptions.

Assumption 1: if a person knows that *c* rather than *c** then she can discriminate *c* from *c**.

Assumption 2: if a person can discriminate *c* from *c** and if *c*'s occurring rather than *c** would explain to her why *e* occurred rather than *e**, then the person can discriminate *e* from *e**.

Now suppose that a person is in possession of an explanation of the form '*c* rather than *c** explains why *e* rather than *e**'. Then by assumption 1 she can discriminate *c* from *c'*. So by assumption 2 she can discriminate *e* from *e**. So when she cannot discriminate *e* from *e** she cannot know that *c* rather than *c**, where the fact that *c* rather than *c** would explain why *e* rather than *e**. So, when failure to discriminate is at the root of a restriction in the contrastive force of knowledge of some proposition, as it very often is, then it will generate a restriction in the contrastive force of knowledge that explains that proposition.¹⁸

This argument applies to any conception of explanation in terms of which the

18 A small qualification: the acquisition of an explanation can increase one's discriminatory capacities. So the knowledge that comes from explanation cannot be finer-grained than the knowledge that we have of the facts it explained *after we have managed to explain them*.

relation 'c rather than c* explains why e rather than e*' can be defined in a way satisfying the two assumptions.¹⁹ The Lewis-Lipton conception is one such. The relation holds when c explains why e rather than e* and c* is the alternative event mentioned in the definition. (See above.) It can also be reconstructed on a Hempelian account of explanation, when c describes initial conditions from which e can be deduced and c* describes conditions from which e* can be deduced, using the same set of laws. We expect that on a variety of ways of understanding explanation it will be possible to reconstruct the argument. For, to put the general point very loosely and intuitively, explanation provides ways of discriminating one possibility from another, so that we cannot expect the knowledge that comes from explanation to be finer-grained than our knowledge of the facts that it explains.²⁰

Putting the pieces together: the capacities of limited agents

To end the paper, we will try to say something about contrastive knowledge in general. There have been several hints above about the general pattern involved, when we were discussing one or another context in which it is natural to use a 'knows p rather than that q' idiom. In several places we referred to agents' limited capacities to discriminate one possibility from another. We get contrastive knowledge when we can discriminate some but not all possibilities. Consider a human or animal agent in an environment and equipped with reactions that are sensitive to some features of that environment. When the environment exhibits the features the agent reacts. Salmon are sensitive to the smells of their native rivers; scientists are sensitive to evidence supporting hypotheses; sniffer dogs are sensitive to particular drugs and explosives. But just as fundamental as sensitivity is failure of sensitivity. Salmon cannot discriminate water from their native stream from water that has the same mineral and organic constituents. (Chemically identical water is the salmon equivalent of a Cartesian deceiver.) Scientists cannot discriminate between possibilities that would produce all the same observations for human observers. Scientists, all humans in fact, also have conceptual limitations that make us unable to articulate the differences between some possibilities, and unable to understand what observations would differentiate others whose distinctness in principle we can grasp. (Of course, being human we find it easier to assert that such limitations must exist than to produce very helpful examples of them.)

19 If as suggested above being able to explain something is a kind of wh-knowledge, knowing why, then all explanation why e rather than e* has the implicit form 'e rather than e* because c rather than c*. Just as knowing where is knowing that the location is one place rather than another, knowing why should be knowing that the reason is one cause rather than another.

20 It is worth comparing what we say here with the interesting account in Reiber (1998). Reiber takes as a basis for knowing that p, that p is the best explanation of the fact that the person believes that p. For him explanation is contrastive, but the knowledge in question is non-contrastive but relative to a context which is a function of the contrast space of the explanation. We doubt that this works, as an explanation of 'knows that p' for reasons that Reiber himself gives in his last section. But much of what he says can be reformulated in our terminology.

Given the limitations on agents' capacities to discriminate, all the simple facts about the reliability of agents' information are contrastive. Given that the truth lies in one of a pair, agents can often with considerable reliability give a reaction that indicates which one. That is too weak to constitute any kind of knowledge, though, since the truth may lie in neither one, and if it does, the agent's realization that it is one of the pair may be purely accidental. (You have it on good authority that if there is a choice between Istanbul and Karachi your company will send you to Karachi. But they could send you to Rome or Los Angeles. You consult your horoscope which leads you to believe you will be living near a mosque, so you conclude it is to be Istanbul or Karachi, and thus that it is to be Karachi. You do not know that you will be sent to Karachi rather than Istanbul.²¹) On the other hand contrastive knowledge does not need the central feature of non-contrastive knowledge, which is that the agent should be sensitive to the absence of all possibilities except for the actual one and a certain hard-to-define and possibly contextually determined class of possibilities which are somehow irrelevant. As David Lewis²² puts it 'S knows that P iff S's evidence eliminates every possibility in which not-P – Psst! – except for those possibilities that we are properly ignoring.' It is as if s knows that p when s knows that p rather than that q for all q – except those that don't count.

To say that s knows that p rather than that q is thus to say something stronger than that if p or q were the case s would reliably believe that p, and weaker than that s knows that p. It is *something* like: s is sensitive to its environment in such a way that some process that will result in either p or q will reliably induce a reliable discrimination between p and q. In many cases this will amount to 's can eliminate every possibility in which p is false because q is true instead'. In some cases, when the contrast is very narrow, contrastive knowledge will fall far short of non-contrastive, and in others, when the contrast is very wide, it will seem hardly necessary to mention the contrast since the person has noncontrastive knowledge or something very close to it. The most interesting cases lie in between, when the contrast is wide enough to exclude many relevant possibilities but narrow enough that some significant ones are not excluded.²³ (For example Reilu with respect to cocaine, explosives, and artificial cocaine-odor, or the Hellenistic astronomer with respect to ball-shaped, flat, and indented earths.)

21 Examples like this show that knowing that p rather than q does not reduce to knowing (p or q) or knowing (if p or q then p.) You can know that if it is Istanbul or Karachi then it is Istanbul, without knowing that it is Istanbul rather than Karachi. Note that you can know that if the robin is in the tree or on the ledge then it is on the ledge, without knowing where the robin is, so not knowing that it is on the ledge rather than some other places.

22 See Lewis (1996).

23 We are grateful to a referee for making this clear to us. Again a comparison with contrastive explanation is helpful. Some contrastive explanations come cheap: it is trivial to explain why the tides come roughly every twelve hours rather than roughly every millisecond. But that should not obscure the non-trivial contrast structure of serious explanations that say for example why the tides come every 12.4 hours rather than every 10 (and which may not explain why they come every 12.4 rather than every 12.2.)

That is as near as we will come to a definition. Incomplete as it is, it fits well with two more features of contrastive knowledge. The first is the lesser relevance of considerations about defeaters or ignorable possibilities. In the territory of occasional fake barns a person can know that what she is looking at is a barn rather than a castle or a silo, even if she does not know – simpliciter, non-contrastively – that it is a barn. One consequence of this is that contrastive knowledge is not vulnerable to examples of nested opposing defeaters such as those introduced by Robert Brandom²⁴. Suppose that in situation 1 you have very basic ornithological background and are looking at a robin with no unusual complications. You know that there is a robin on the branch. In situation 2 there is a flock of orioles around the tree, though you are not aware of them. If one of them rather than the robin had been on the branch you would have taken it for a robin. You do not know that there is a robin on the ledge. And in situation 3 there is also a crowd of ornithologists surrounding the orioles remarking loudly on their resemblance to robins, though you are no more aware of them than you are of the warblers.

You know that there is a robin on the ledge. And so on. At each stage the conclusion alternates between knowledge and ignorance, and – it has been argued – there is no reason to believe that it has to converge to a limiting value. But the alternations only affect non-contrastive knowledge. At each stage you know that there is a robin on the ledge rather than on the tree. Even at the first stage you do not know that there is a robin rather than an oriole on the ledge. But at all stages you know that there is a robin rather than a crow there.

The other feature of contrastive knowledge that now falls into place is the *de re* aspect of many contrastive and *wh*- idioms. If you know where the robin is then you know *of* the robin that it is in one place rather than a range of others.

Your sensitivity is to the location of the bird and not to other features of it, and is triggered by its coming into your range of vision. There are some situations – for example involving paranoid delusions about stalking robins – which would result in the belief that there is a robin on the branch but which would not be of a kind that reliably invokes the location-tracking sensitivity. In such situations you could lack non-contrastive knowledge that there is a robin on the branch, and still know of that robin that it is on the branch rather than on the tree trunk.

This completes our plea for contrastive knowledge. We hope to have made a case that it is both robust and stringent. It is robust in that the facts in virtue of which it applies to agents are non-contextual aspects of their situation in the environment. It is stringent in that it is not trivially satisfied, and as a result ascriptions of contrastive knowledge can have explanatory force. It is thus a concept that we should take seriously if we are searching for epistemic properties that individuals can have by virtue of functioning as living creatures in a natural environment.²⁵

24 Brandom (1998). See also the discussion of the example in Kornblith (1999).

25 We have had very helpful suggestions from Jonathan Adler, Christopher Hookway, and Peter Lipton. When this paper was near completion we saw a draft by Jonathan Schaffer arguing for similar conclusions.

CV Antti Karjalainen is finishing his PhD dissertation at the University of Bristol. His work is on Epistemological Contextualism with a special focus on the relations between the skeptical paradox and the Gettier-problem.

CV Adam Morton is professor of philosophy at the University of Oklahoma. His book *The Importance of Being Understood: Folk Psychology as Ethics*, has recently appeared with Routledge.

References

- Boër, Steven E and William G Lycan (1986) *Knowing who*, Cambridge MA: MIT Press.
- Brandom, Robert (1998) "Insight and blindspots of reliabilism", *The Monist*, **81**(3), 371-392.
- DeRose, Keith (1995) "Solving the skeptical problem", *Philosophical Review* **104**, 1-50
- Dretske, Fred (1981) "The pragmatic dimension of knowledge", *Philosophical studies* **40**, 363-378.
- Dretske, Fred (2000) "Epistemic Operators", in *Perception, Knowledge and Belief*, New York: Cambridge University Press, 30-47.
- Garfinkel, Alan (1980) *Forms of Explanation*, New Haven: Yale University Press.
- Ginet, Carl (1992) "Causal theories in epistemology", in Jonathan Dancy and Ernest Sosa, eds. *A companion to epistemology*, Oxford: Blackwell.
- Glymour, Clark (1980) *Theory and evidence*, Princeton: Princeton University Press.
- Goldman, Alvin (1976) "Discrimination and perceptual knowledge", *Journal of Philosophy* **73**, 771-91.
- Goldman, Alvin (1986) *Epistemology and Cognition*, Cambridge MA: Harvard University Press.
- Gordon, Robert (1988) *The Structure of the Emotions*, Cambridge UK: Cambridge University Press.
- Hacking, Ian (1965) *The logic of statistical inference*, Cambridge UK: Cambridge university Press.
- Hanson, N. R. (1958) *Patterns of Discovery*, Cambridge UK: Cambridge University Press.
- Hintikka, Jaakko (1992) "Different constructions in terms of 'knows'", in *A companion to epistemology*, Jonathan Dancy and Ernest Sosa (eds.), Oxford: Blackwell Publishers, 99-104.
- Howson, Colin & Peter Urbach (1989) *Scientific Reasoning: the Bayesian approach*, La Salle, Ill.: Open Court.
- Hookway, Christopher & Tobias Grimaltos (1995) "When deduction leads to belief", *Ratio* **8**, 24-42.
- Kornblith, Hilary (1999) "Knowledge in humans and other animals", in *Philosophical Perspectives* **13**, *Epistemology*, James E Tomberlin (ed.), Oxford: Blackwell.
- Lewis, David (1996) "Elusive knowledge", *Australasian Journal of Philosophy* **74**, 549-67.
- Lipton, Peter (1991) *The Inference to the Best Explanation*, London: Routledge.
- Mayo, Deborah (1996) *Error and the growth of experimental knowledge*, Chicago: Chicago University Press.
- Morton, Adam (2001) *A guide through the theory of knowledge*, third edition. Oxford: Blackwell.
- Reiber, Steven (1988) "Scepticism and contrastive explanation", *Nous* **32**, pp. 198-204.
- Spiegel, Murray R (1961) *Statistics*. McGraw-Hill.
- Temple, Dennis (1988) "The Contrast Theory of Why-Questions", *Philosophy of Science* **55**, 141-151.
- Williamson, Timothy (1992) "Inexact Knowledge", *Mind* **101**, 217-242.
- Williamson, Timothy (1995) "Is knowing a state of mind?", *Mind* **104**, 533-65.
- Williamson, Timothy (2000) *Knowledge and its limits* Oxford University Press.