

Contrastivity and indistinguishability

Adam Morton and Antti Karjalainen

There are many contrastive constructions in natural languages. "I would prefer Canadian rather than American dollars", "She hopes that the flowers are delivered to her apartment rather than to her office", "Copernicus knew that the sun rather than the earth is the center of the solar system", "Descartes could explain why a ball approaching a wall at 45 degrees reflects at 45 rather than 60 degrees." The last two of these have received a fair amount of philosophical attention. The whole list deserves attention.

We give a general description of a class of contrastive constructions, attributing attitudes such as knowledge, belief, hope, or fear. They all involve a systematic ignoring of possibilities which are treated as if they were indistinguishable from others, which allows language to describe fittings between the cruder contents of our minds and the richer contents of the world. They allow us to make quite detailed ascriptions without claiming too much.

focussing on the intended senses Most contrastive idioms are ambiguous. "Sally knows that the burglars will come in through the window rather than the door" usually says more than "Sally knows that if the burglars were to come in through either the window or the door they would choose the window", though that is one meaning that it can express in a suitable context. For the "if window or door then window" sentence doesn't express any kind of knowledge that the burglars will come in through the window. Yet there is a concept of contrastive knowledge, discussed in a number of recent papers (see Johnsen 2001, Morton and Karjalainen 2003, Sinnott-Armstrong 2004, and Schaffer 2005) according to which knowing that p rather than q does express a kind of knowledge of p . To begin to home in on the general class of interpretations of contrastive constructions in question, consider some features of the two concepts that have come under philosophical scrutiny, contrastive knowledge and contrastive explanation (for contrastive explanation see Garfinkel 1980, Lipton 1991.)

(a) They do not reduce to non-contrastive variants. Suppose we say that Sally knows that her motorcycle weighs one hundred rather than two hundred kilos, but does not know that it weighs one hundred rather than one hundred and one kilos. Then we are ascribing a kind of imprecise knowledge to her of the weight of her

motorcycle, focusing on its being one hundred kilos. We are not saying that if she knows one thing then she does or does not know another. Similarly, an explanation of why a ball reflects from a wall at 45 degrees rather than 60 degrees which is not an explanation of why it reflects at all rather than shattering or piercing the wall, is, as several authors have convincingly argued, not to be understood as a complex of non-contrastive explanations of individual events. So the relevant senses are irreducibly two-place.

(b) The argument places are implicitly propositional. They attribute relations between pairs of propositions, whose roles are typically very different. So "she knows that the burglars will come in by the window rather than the door" is to be taken as a contraction of "she knows that the burglars will come in by the window rather than that the burglars will come in by the door". And "she hopes that the flowers will be delivered to her apartment rather than to her office" is to be taken as a contraction of "she hopes that the flowers will be delivered to her apartment rather than the flowers will be delivered to her office."

(c) The second argument place is not redundant. "René knows that he exists rather than not-existing" says no more than that René knows that he exists. "Martin understands why there is something rather than nothing in the fridge" says no more than that Martin understands why the fridge is not empty. Nearly all contrastive idioms can be used in ways that relate a person just to one of the propositions mentioned. But those are not the uses of these constructions that are of interest to us.

These three characteristics are only a beginning. They don't give necessary and sufficient conditions. But they go some way to isolating the contrastive concepts that are essential components of our thinking. In terms of them we can see that some "rather than" constructions are *less* likely to be of the intended kind, to be "essentially contrastive" as we shall put it, than others. For example "prefers that" most often expresses a two-place comparative rather than contrastive relation. "Believes p rather than q" is another example. When we say that someone believes that her car keys are in her backpack rather than in her jacket we usually mean just that she believes that her car keys are in her backpack and also believes that they are not in her jacket. That is not to say that we have no use at all for an essentially contrastive concept of belief. We return to the question below.

Contrastive idioms are closely related to a variety of other constructions. Usually when one uses a *wh*- construction, such as "knows who", "explains why" "has a belief about where", what one

says can be paraphrased in contrastive terms. For example "Kaija knows where the money is hidden" means "Kaija knows that the money is hidden at some location rather than any other location in class C" where C is specified by the context. In fact, *wh-* constructions allow a very general way of expressing contrastive ideas. When we say that, for example, Kaija knows that the money is in the basement rather than in the attic, we usually mean that Kaija knows that the money is in the basement rather than in each location from a wide but far from universal range of locations. So we can rephrase the belief that is communicated as

Kaija knows which proposition in P is true: the money is in the basement

where P is a set of propositions describing the location of the money or more idiomatically and loosely, leaving P implicit

Kaija knows that where the money is, is in the basement

or

Kaija knows where the money is: in the basement

Generalizing from the terminology of contrastive explanation, we will refer to the proposition that is the main object of the attitude as the *target* of the attitude, and to the whole set of propositions from which it is chosen and contrasted as the *foil*. We will sometimes write "person s has attitude A to target p with foil P" as " $A_s(p:P)$ ".

Contrast and logic A basic purpose of contrastive constructions is to avoid unwanted claims. When we say that René explained why the ball reflected at 45 degrees rather than 60 degrees, we are avoiding saying that René explained why the ball did not shatter or go through the wall. But this aim conflicts with closure principles for knowledge and explanation: the idea that, for example, if a person knows that p and appreciates that q is a logical consequence of p, then the person knows that q. For "did not shatter" follows logically from "reflected at 45 degrees", so if we apply the analogous closure principle to contrastive explanation we are forced to say that René explained why the ball did not shatter. And this would deny the whole point of the idiom. Similar collapses await contrastive knowledge and other contrastive idioms.

Closure principles raise subtle and controversial issues. It is far from obvious what to say about them in the case of non-contrastive knowledge, let alone in the full range of contrastive idioms. So anything we say is tentative and experimental. First, we will ignore the issue of how to model the fact that people have limited logical abilities, so they do not see many consequences of their thoughts and can fail to see contradictions in what they think. We suspect that

contrastivity has something to contribute to this problem. We will only discuss the issue of modelling situations where it is natural to ascribe a state of mind to someone by relating them to a particular sentence, but intuitively wrong, to relate that person in the same way to a logically equivalent sentence, as a result *not* of the person's limited logical capacity, but of something intrinsic about knowledge or understanding or hope. The standard example is knowledge. We should take a person to know from long experience that she has hands. "I have hands" is logically equivalent to "I have hands and I am not a brain in a vat". But we might be much more reluctant to say that the person knows that she has hands and is not a brain in a vat, just on the basis of memory and a glance before her. That does not seem to be enough to refute metaphysical skepticism. James Hawthorne has pointed out that these issues arise throughout our knowledge-attributions. (The issues here stem from Dretske, see Dretske 1981, 200. They have been taken furthest by Hawthorne, see Hawthorne 2004.)

This situation raises semantic and epistemic problems. The semantic problem is how to understand the p-place of "S knows that p". Propositions and facts seem ruled out, since knowing that p would then mean knowing that q, for all equivalent q. The epistemic problem is that very often we do extend our knowledge by following chains of logical consequence. Whether or not it is part of the meaning of "knows", it seems that as a very general fact we do know what we see follows from our knowledge. But if sometimes we cannot extend our knowledge in this way, we ought to ask when we can and when we cannot.

This is where contrastivity enters. When a person knows that she has hands she knows that her arms terminate in hands rather than in stumps or flowers. She does not know that what she takes to be her arms are terminated by hands rather than by programs in the vat computer or ideas in the mind of God. Similarly, it would make sense to say of someone that she knows that she is staying home this summer rather than going on a cruise, while denying that she knows that this lottery ticket rather than this other one will or will not win. So it would be very nice to be able to argue in a non-question-begging way that knowing that one has two hands rather than stumps, or knowing that one will spend the summer at home rather than on a cruise, entails that one does not have three hands rather than stumps, or that one will spend the summer within ten miles of one's house rather than on a cruise, but does not entail that one has any number of hands rather than computer-induced hand-simulation, or that this ticket rather than that ticket will win the lottery.

The crude idea is that deduction can extend contrastive knowledge only when it works within the set of propositions implicit in the contrast. Suppose that you know which proposition in P is true: you have two hands rather than stumps. Let P be {you have n hands – $n = 0, 1, 2, \dots$, you have flowers/heads/claws/prostheses/... growing from your wrists}. Then if realize that “you have two hands” entails some proposition q , which is a Boolean combination of propositions in P , and realising it you continue to believe that you have two hands, we should draw the conclusion that you know that q . And not otherwise.

This formulation is too strong in that it restricts the range of propositions over which deduction can extend one’s knowledge to Boolean combinations of members of P , when in fact more general logical combinations seem permissible. The more general formulation below should take care of this problem. It is also too restricted, in that it is directed at contrastive knowledge, while we would like to say something illuminating about a range of contrastive constructions. It is not in fact obvious how plausible deductive closure conditions are for “fears that p rather than q ” or “hopes that p rather than q ”. It is easier to motivate conditions based on logical equivalence. If p and q are logically equivalent, then they can be taken as representing the same proposition. Anyone who believed one of them and was unconfused about logical relations would believe the other. So the natural condition to explore for an arbitrary attitude “s A’s that p rather than that q ” will be something like “if s A’s that p out of class P and s realizes that p is logically equivalent to q , then s A’s that q out of class P ”. Here too restrictions of eligible propositions q will be required. (The contemporary philosopher who has thought most about this is Jonathan Schaffer. See Schaffer, forthcoming.)

Distinguishability When any propositional attitude is contrastive, it is likely to conflict with deductive relations. That is, if person bears attitude A to p rather than q , then there is likely to be some logical consequence of p , r , such that the person does not bear A to r rather than q . We want to explain this as a general phenomenon, rather than a quirk of knowledge, and so we will give a general framework for contrastive attitudes, and then try to show in terms of this which deductive relations are more and less problematic.

Attitudes, such as knowledge, fear, hope, and understanding, can be taken as relations to propositions. Take propositions to be sets of very finely individuated possible worlds. (For a discussion of the advantages and disadvantages of this approach see Stalnaker 1984. We are confident that the core of what we say could be transferred into any of a number of other ways of understanding the objects of the attitudes.) For each possible world w , each person s , and each

attitude A there is a large class of other worlds, call it $I_A(w, s)$, which the person cannot distinguish from it. This notion of indistinguishability is relative to the attitude A. s may for example be able to distinguish two worlds in imagination, but not in evidence: she may be able to think of many differences between the worlds, but no evidence that she can get will tell her which one she is in. When the attitude is understanding or explanation, then the worlds that are not distinguished from the worlds of interest could also be described as invisible to explanation: the explanatory scheme simply does not take them into account.

For each set of worlds p there is another set of worlds, which we can also write $I_A(p, s)$ which is the set of worlds which are A-indistinguishable from worlds in p . Formally $I_A(p, s)$ is $\{w: \exists w' (w \in I_A(w', s) \ \& \ w' \in p)\}$. Intuitively $I_A(p, s)$ is the proposition whose world-extension the person could not distinguish from p , with respect to A. It is tempting to require that when a person has A to p rather than q then $I_A(p)$ and $I_A(q)$ be disjoint. But in fact this is wrong, as is shown by cases of overlapping indistinguishability. When we express contrastive knowledge in everyday discourse we usually finesse this point by choosing a "rather than" proposition q such that nearby worlds in which it is true are distinguishable by the knower on the basis of evidence available to her from nearby worlds in which the known proposition is true. And we usually pick for the known proposition p a fairly general proposition, which holds in a range of worlds indistinguishable from actuality. So it is as well for us to state explicitly that when we say that a person has attitude A to p rather than q , we intend that p and q are members of indistinguishability classes which may overlap and which may contain many members besides p and q . Thus when a person knows that the window ledge is one metre long rather than two metres long, this is consistent with her not knowing that it is one metre long rather than 101 centimetres long, and also consistent with the fact that she could not tell – suppose that looking is the only means available to her- a ledge of two metres long from one of 201 centimetres. (It is no accident that these definitions are like those in Williamson 1992, developed further in Williamson 2000; issues of contrastivity and issues of precision are closely related.)

The indistinguishability relation I coarse-grains the totality of possible worlds. Many physical devices and perceptual systems do the same, in that they will give the same output in any of a range of worlds. Visual perception for example, cannot distinguish a dog from a perfect holographic image of a dog. But a person may know that it is an image rather than a dog before her. A person may be unable to know that she is in a classroom rather than a life-long computer-

generated illusion containing a classroom episode. But she may be able to fantasise that she is in a life-long illusion rather than a classroom. The coarse-graining is relative to the attitude.

It is easy now to see why deductive relations cut across contrasts. Given a set of worlds p , and a set $I_A(p, s)$ A -indistinguishable from p for s , there will always be a consequence of p which is within $I_A(p, s)$. For the consequences of p are all the sets of worlds of which p is a subset. And these include sets including all the worlds in p and excluding some in some proposition indistinguishable from p . Moreover they include propositions from every indistinguishability class, since they include "p or r" for any r whatsoever. As a result logical consequence will always lead from propositions that a contrastive relation holds between to propositions it does not. That is, it will if we make the following assumptions:

- (a) If $I_A(p, r)$ and $A_s(p \text{ rather than } q)$ then not $A_s(r \text{ rather than } q)$
- (b) For some h, v , and r : $A_s(h \text{ rather than } r)$ and $I_A(h, v)$
- (c) If $A_s(p \text{ rather than } q)$ and p entails r then $A_s(r \text{ rather than } q)$

We then get the unwanted pair

- (d) $A_s(p \text{ rather than } q)$ (from (b) and (c))
- (e) not $A_s(p \text{ rather than } q)$ (from (a) and (b))

So what deductive relations do preserve contrastivity? There is an obvious suggestion to make. The suggestion is best understood with an image in mind. In the space of all possible worlds imagine the worlds that are indistinguishable for a person as a plane clinging to the range of worlds between which she can distinguish. Then we want logical equivalence to govern the extension of an attitude only insofar as it is restricted to this plane. More formally, we can express this as follows.

We make use of the point made above that a contrastive proposition normally alludes to a greater range of contrasts than is specified by a single foil. So consider a proposition $A_s(p:P)$. P is a set of propositions, thus a set of sets of worlds. Let W_p be the set of all worlds which are members of propositions in P . Let P^2 be the set of all subsets of P , i.e. the set of all propositions true in worlds in W_p . Then the natural conditions to impose are:

- (Equivalence) if q is logically equivalent to p and $q \in P^2$ then $A_s(q:P)$
if and only if $A_s(p:P)$
- (Entailment) if $A_s(p:P)$ and q is a logical consequence of p & $q \in P^2$
then $A_s(q:P)$

It seems also consistent with this general strategy to require

- (Uniformity) if $A_s(p:P)$ and $P' \subseteq P^2$ then $A_s(p:P')$

As we noted, Equivalence is more plausible than Entailment as a condition on contrastive attitudes in general.

Does Entailment give the desired answers in the special case of contrastive knowledge? The following follow from Entailment plus natural assumptions about the situation of the person *s*:

if *s* knows that the physical objects she sees at the end of her arms are hands, and "these are hands" entails "these are not stumps" then *s* knows that the physical objects she sees at the end of her arms are not stumps.

if *s* knows that where she'll be next summer is at home and "*s* is at home" entails "*s* is not on a cruise" then *s* knows that where she'll be next summer is not on a cruise

And the following do not follow

if *s* knows that the physical objects she sees at the end of her arms are hands, and "these are hands" entails "*s* is not a brain in a vat" then *s* knows that she is not a brain in a vat.

if *s* knows that where she'll be next summer is at home and "*s* is at home" entails "*s* did not win the lottery" then *s* knows that where she'll be next summer is not somewhere a lottery would take her.

We are understanding these sentences in the way described earlier in this paper, so that e.g. "*s* knows that the physical objects at the end of her arms are hands" is taken as "*s* knows what physical objects are at the end of her arms: hands." We are not going to give formalizations of these sentences, together with the background assumptions, and proofs that these consequences really do follow. The main reason, besides a concern for the patience of the reader, is the fact that a full treatment requires assumptions about multi-premise closure, that is, the logical consequences of sets of contrastive statements, involving different *P* and thus different basis sets P^2 . That is a task for a separate and more technical treatment.

How fine can the grain be? Different attitudes require different indistinguishability relations. It is plausible to assume that there is a limit to the fineness with which any person can distinguish possibilities. But it is not clear where that limit lies.

Consider some candidates. Indistinguishability with respect to imagination seems as if it might be the most basic kind of distinguishability. After all, if you cannot even imagine what the difference between *p* and *q* is then you can hardly know that *p* rather than *q*, or explain why *p* rather than *q*. A problem with this suggestion is that logical inference can carry us towards distinctions that we cannot grasp imaginatively. You cannot imagine the difference between a phase in the initial expansion of the universe taking 10^{-6} seconds and its taking 10^{-8} seconds. But you might be able to manipulate a physical theory that explained why the figure was 10^{-6}

rather than 10^{-8} . We might stipulate that this constitutes a kind of imagination, but in so doing we would have shifted the finest grain of contrastive content to something else, linguistic or symbolic expressibility.

So consider expressibility in symbols. This too may not be the finest conceptual grain, since one may be able to imagine distinctions between possibilities that one cannot put into words. The conductor may know or hope that the oboe soloist will manage to produce this timbre rather than that one, without being able to describe the essential contrast between timbres.

Indistinguishability with respect to belief combines both imaginability and symbolic expressibility. You can believe anything that you can put into an intelligible symbolic form, and you can have beliefs with contents that outrun the symbols available to you, as when you believe that the oboe solo will sound one way rather than another. Of course this attests largely to the width and vagueness of the concept of belief: it can apply to any content that we reason with or use as a basis for action. So indiscriminability with respect to our beliefs might provide the finest grained partition of the space of all sets of possible worlds available to us. This is a basic reason why deductive closure is problematic for knowledge: we can distinguish in belief possibilities that our evidence and other grounds for knowledge cannot distinguish between.

This is also a reason why belief is very often *not* contrastive. We do not regularly say "I believe that p rather than q", meaning this as an essentially contrastive attitude. In fact we would not understand a contrastive assertion of them. For there to be a point in saying "I believe p rather than q" rather than simply "I believe that p" there would have to be some contrast with some "I do not believe p rather than r". But this would require q and r to be indistinguishable with respect to belief. And a person could never be in a position to say this, since once she had distinguished q from r they would be distinguishable.

So it is tempting to think of belief as a non-contrastive concept. But the third person situation is rather different. The most evident differences are when there are large cognitive differences between the ascriber of beliefs and the individual to whom they are ascribed. Suppose for example that we are describing Reilu the sniffer dog at Helsinki airport. Reilu can be convinced that there is cocaine in a suitcase. And her reaction is very different from if she smells plastic explosive, something else she has taught to be very interested in. The cocaine smell might be caused by some synthetic drug, though. So from Reilu's point of view, conceptually and perceptually, cocaine and plastic explosive are distinguishable but cocaine and the synthetic drug

are not. The possibility of the synthetic cocaine-like drug as something distinct from cocaine is not something that figures in Reilu's thinking. As a result the following dialogue would be perfectly natural.

Jaako: Look, Reilu's discovered something, from her behaviour it could be cocaine. That's what she thinks it is.

Kaija: We've had a lot of synthetic stuff passing through here recently. Some of it's legal, unfortunately. Does she think it is not that?

Jaako: No, there's no difference for her. Cocainy-stuff is cocainy-stuff for Reilu. She believes it's cocaine rather than plastic explosive, though.

So when the ascriber can distinguish possibilities that the believer cannot, it makes sense to say "believes p rather than q". There are cases where one human being can say this of another, either because of greater cognitive development, or cultural difference, or individual quirks. An example of the last would be when we find out that someone cannot understand the difference between two attitudes, say irony and sarcasm. She understands the difference between either and say exaggeration, understatement or literalness, but she has no idea how the subtly related pair differ, and uses them interchangeably. We might say of such a person that she thinks someone making a joke was being ironical rather than exaggerating, though she did not think she was being ironical rather than being sarcastic.

Essentially contrastive attributions of belief are thus not at all impossible. We use them largely in describing the beliefs of agents whose thinking discriminates between different possibilities than we do. These uses are perfectly compatible with the assumption that for any agent distinguishability with respect to belief provides the finest grain of discrimination, from which the indistinguishabilities with respect to other attitudes may be got by coarsening.

That assumption is plausible. But it is not completely obvious. Consider how fear, for example, might link to a finer grid of possibilities. This would require there to be two possibilities, p and q, such that it was not possible for a human being to believe that p but not q, yet possible to fear that p rather than q. It would not be possible for a person to articulate their fear in words, since once verbalised the feared possibilities would become potential objects of belief. And for the same reason it is not possible to give explicit examples of such possibilities. But we can raise the question as follows. There must be aspects of the universe that are not discriminable by human thought, given the limited nature of our cognitive capacities after only a few million years of evolution. Is it so clear that there could not be such aspects which are discriminable by

fear, hope, or other emotions, so that we can fear or hope for one thing rather than another, even though we cannot articulate the contrast between what we feared or hoped for and what we did not? Such inarticulable objects of emotion might be expressed in religious rituals, works of art, or metaphysical contortions of language.

Conclusion: contrastive constructions: We have not defined contrastive knowledge, fear, hope, understanding, or belief, or given definite criteria for distinguishing them from their non-contrastive variants. And we have said nothing about a very fundamental question: when one uses an attitude-terms non-contrastively is one always tacitly making a contrastive claim, though often one where the contrast class is too obvious to need stating, or vacuously wide? It would be tempting to answer Yes, but we have no pretensions to have made that case. Any defence of that rather radical claim, which would have important consequences both in epistemology and philosophy of mind, is likely to make use of the idea that an assertion "s As that p rather than q" is usually best taken as saying that s has A to target p against a more general foil P, a set of propositions of which q is just an example.

What we have done is to present a general framework in terms of which a variety of contrastive attitudes can be understood. If we are right, when one attributes a contrastive attitude to a person one is relating a person to a nexus of possibilities. The full set of possibilities is richer than she can articulate with the conceptual resources associated with that attitude, and so the attribution only makes sense relative to a coarse-graining of it, a way of making some distinctions within it insignificant. The language of contrastive attitudes is one of our central ways of describing that coarse-graining.

University of Bristol
University of Alberta

BIBLIOGRAPHY

- 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.

- Johnsen, Bredo (2001), 'Contextualist Swords, Skeptical Plowshares', *Philosophy and Phenomenological Research* 62: 385-406.
- Lipton, Peter (1991) *The Inference to the Best Explanation*, London: Routledge.
- Reiber, Steven (1988) "Scepticism and contrastive explanation", *Nous* 32, pp. 198-204.
- Schaffer, Jonathan (2005), 'Contrastive Knowledge', in Tamar Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology* 1. Oxford: Oxford University Press.
- Sinnott-Armstrong, Walter (2004), *Pyrrhonian Skepticism* Oxford: Oxford University Press.
- Stalnaker, Robert (1984) . *Enquiry* Ithaca N.Y.:Cornell University Press.
- Williamson, Timothy (1992) "Inexact Knowledge", *Mind* 101 no 402, 217-242.
- Williamson, Timothy (2000) *Knowledge and its limits* Oxford: Oxford University Press.