

Alien Abduction:
Inference to the Best Explanation and the Management of Testimony

PETER LIPTON

In this paper I explore the prospects of applying inference to the Best Explanation (TIBE – sometimes also known as ‘abduction’) to an account of the way we decide whether to accept the word of others (sometimes known as ‘aliens’). IBE is a general account of non-demonstrative or inductive inference, but it has been applied in a particular way to the management of testimony. The governing idea of Testimonial IBE (TIBE) is that a recipient of testimony (‘hearer’) decides whether to believe the claim of the informant (‘speaker’) by considering whether the truth of that claim would figure in the best explanation of the fact that the speaker made it.

I begin by introducing the general idea of IBE as an account of many non-demonstrative inferences. Then I make some general remarks about the epistemology of testimony, and especially about the problem of accounting for the way we decide whether to believe what we are told. Third, I suggest how IBE could be used to account for this process of testimonial discrimination, the idea of TIBE. In the final two sections I consider the obvious objection to TIBE: it is either trivial or false.

Inference to the Best Explanation in general

According to the explanationist, that is a supporter of IBE, explanatory considerations are a guide to inductive inference. We decide which of the competing hypotheses the evidence best supports by determining how well each of the competitors would explain that evidence (Harman 1967; Lipton 2004). Many inferences are naturally described in this way. Seeing the ball next to the broken vase, I infer that my children have been playing catch in the house, because this would be the best explanation of what I see. Darwin inferred the hypothesis of natural selection because, although it was not entailed by his diverse biological evidence, natural selection would provide the best explanation of it. When astronomers infer that a galaxy is receding from the earth with a specified velocity, they do this because the recession would be the best explanation of the observed red-shift

of the galaxy's characteristic spectrum. When the detectives infer that it was Moriarty who committed the crime, they do so because this hypothesis would best explain the fingerprints, blood stains and other forensic evidence. Sherlock Holmes to the contrary, this is not a matter of deduction. The evidence, as normally construed, will not entail that Moriarty is to blame, since it always remains possible that someone else was the perpetrator. Nevertheless, Holmes is right to make his inference, since Moriarty's guilt would provide a better explanation of the evidence than would anyone else's.

IBE can be seen as an extension of the idea of *self-evidencing* explanations, where the phenomenon that is explained in turn provides an essential part of the reason for believing the explanation is correct. For example, the speed of recession explains the red-shift, but the observed red-shift may at the same time be an essential part of the reason astronomers have for believing that the galaxy is receding at that speed. Self-evidencing explanations exhibit a curious circularity, but the circularity is benign. The recession is used to explain the red-shift and the red-shift is used to determine the recession, yet the recession hypothesis may be both explanatory and well-supported. According to IBE, this is a common situation: hypotheses are supported by the very observations they are supposed to explain. Moreover, on this model, the observations support the hypothesis precisely because it would explain them.

IBE thus partially inverts an otherwise natural view of the relationship between inference and explanation. According to that natural view, inference is prior to explanation. First we must decide which hypotheses to accept; then, when called upon to explain some observation, we will draw from our pool of accepted hypotheses. According to inference to the best explanation, by contrast, it is by only by asking how well various hypotheses would explain the available evidence that we determine which hypotheses merit acceptance. In this sense, inference to the best explanation has it that explanation is prior to inference. Relatedly, IBE reconstrues the relationship between knowledge and understanding. Knowledge that is in a natural sense prior to understanding why: when we ask why some phenomenon occurred we already know that it did but don't yet understand why. But according to IBE, we ask why questions not just in order to gain new understanding of things we already know, but also to gain new knowledge. In that sense, why-questions have an instrumental function: they are a tool for the acquisition of new knowledge.

Although it gives a natural account of many inferences in both science and ordinary life, the model clearly needs further development. What, for example, should be meant by *best*? It is sometimes taken to mean likeliest or most plausible, but inference to the likeliest explanation would be a disappointingly uninformative model, since the main

point of an account of inference is to say what leads one hypothesis to be judged likelier than another, to give the symptoms of likeliness. A more promising approach construes *best* as *loveliest*. On this view, we infer the hypothesis that would, if correct, provide the greatest understanding.

IBE should thus be construed as inference to the loveliest explanation. Its central descriptive claim is that loveliness is a guide to likeliness, that the explanation that would, if correct, provide the most understanding, is the explanation that is judged likeliest to be correct. This at least is not a trivial claim, but it faces at least three challenges. The first is to identify the explanatory virtues, the features of explanations that contribute to the degree of understanding they provide. There are a number of plausible candidates for these virtues, including scope, precision, mechanism, unification and simplicity. Better explanations explain more types of phenomena, explain them with greater precision, provide more information about underlying mechanisms, unify apparently disparate phenomena, or simplify our overall picture of the world. But analyzing these and other explanatory virtues is not easy, and it also leaves the other two challenges. One of these is to show these aspects of loveliness match judgments of likeliness, that the loveliest explanations tend also to be those that are judged likeliest to be correct. The remaining challenge is to show that, granting the match between loveliness and judgments of likeliness, the former is in fact our guide to the latter.

Testimony in general

Testimony is here understood very broadly: it is tantamount to assertion. On this standard reading, our reliance on testimony is enormous: almost all our beliefs are based at least in part on the word of others (Coady 1992). The flow of new testimony is strong and steady, and life is too short for us to explicitly assess everything we are told before we believe it. At the same time, we obviously cannot believe everything we are told: for one thing, that would mean believing contradictions. So it is natural to give a ‘default-trigger’ model of testimony. According to a simple version of such a model, in most contexts the hearer simply accepts what she is told, without engaging in any conscious evaluation or inference. When I ask someone the time and she tells me, I believe her with scarcely more deliberation and scarcely less confidence than I believe what I seem to see in front of my eyes.

There are however in the case of testimony diverse conditions that may trigger the hearer to switch from default into evaluative mode, where he pauses to consider whether he should believe what he has been told. This is to some extent analogous to situations where we hesitate to believe our eyes, though the switch into evaluative mode is much

more common in testimonial than in perceptual cases. The obvious triggers include cases where what is claimed contradicts firmly held beliefs, cases of contradictory testimony, and cases where there is reason to worry about incompetence or insincerity. What interests me is what hearers do once they have entered evaluative mode. How do they then decide whether to believe what they have been told? That is the problem of the management of testimony, the problem that is the focus of this essay.

Testimonial inference to the best explanation – TIBE – is the solution to this problem that I want to begin to develop and defend, but first I should flag a number of complications about the way we react to testimony, mostly to suggest that they don't matter for my purposes. To begin with, I should confess that the simple version of a default-trigger mechanism that I have just given may well be too simple. For example, we really need a kind of double-default model, where in addition to the many cases where we simply accept what we are told, there are also many cases where we simply suspend judgement, without bothering to assess the speaker, again because life is too short and the proffered information just doesn't matter enough to us. Again, maybe it better to think of even the original default cases where we seem immediately to accept what we are told as having an inferential aspect. Thus even in the mundane case of asking the time, my acceptance is not so blind as to lead me to believe that it is midnight when the sun is blazing, whatever the speaker says. As Elizabeth Fricker (1995) has suggested, even in what I would classify as default cases we seem to engage in a kind of plausibility monitoring that is at least closely related to inference. (The debate over the extent of inference in testimony and other core issues in the epistemology of testimony is well discussed in Kusch 2002.) But for my present purposes these issues probably don't matter, because what interests me here is the structure of explicit inferences about testimony, not how common such inferences are or how similar they are to less explicit processes.

This focus on the triggered cases of careful deliberation does however raise another question. Suppose that the speaker tells me something, and after careful deliberation I decide to believe him, and moreover, my belief is fortunate enough in this case to count as knowledge. Is it right to count this as a case of *testimonial* knowledge? Perhaps there has now been too much processing for that expression to be appropriate and it should rather be reserved for the default cases. One way of motivating this thought might be to suggest that testimonial knowledge, properly so-called, is the result of the transmission of knowledge that is testimony's primary function, but that this is not what is happening in the triggered inferential cases, precisely because they are inferential. The inferential process generates a warrant for what the speaker says that will at least often be different

from the warrant the speaker herself enjoys. The process is generative rather than purely transmissive, and hence the resulting knowledge the speaker enjoys is not testimonial knowledge.

If the default cases are purely transmissive while the trigger cases are not, then one might mark this difference by reserving the expression ‘testimonial knowledge’ for the former. On the other hand, even in the trigger cases where the speaker comes to know what has been asserted, the fact that the speaker said what she did remains an essential part of the warrant that the hearer has for that knowledge. That is a reason to use the expression ‘testimonial knowledge’ to cover trigger as well as default cases. And there is an even more telling reason to do this it turns out that even default cases are not purely transmissive. And this is a claim with which I have considerable sympathy.

On first hearing, a transmission model of testimony is very plausible. It is natural to suppose that certain cognitive mechanisms, such as perception and inference, have the job of generating new information in the epistemic community, while testimony has the quite different job of spreading the wealth around, of transmitting what others have learned by other means (such as perception and inference). Of course nobody will deny that the generation of new knowledge is a common result of this transmission – if I tell you P and you already know that if P then Q, then you may come to know Q, which could well be something nobody in the community knew before. So clearly the use of testimonial knowledge gained in default mode may result in the generation of new knowledge. But this common case does not rule out the claim that when a hearer comes to know P in virtue of accepting in default mode the speaker’s assertion that P, then there is never any generation vis a vis P. Yet this too is false, as Jennifer Lackey (1999) has shown. For it may be that the speaker failed to know P because of a defeater, while the speaker nevertheless comes to know P by accepting what he has been told, because he does not suffer from that (or any other) defeater to P. An example of the sort of case Lackey discusses would be one where you see Tim Williamson on the street and tell me that he is in town, and I uncritically accept your testimony. I thus come to now that he is in town, even if you do not because, unbeknownst to me, you have been told falsely that there is someone around who looks very much like him. More generally, a speaker’s belief may fit slightly too poorly with her other beliefs to qualify as knowledge for her, while it fits sufficiently better with the other beliefs of the hearer to count as knowledge for him.

Where does this leave us? What Lackey cases show is that if there is a divide between generative and merely transmissive uses of testimony, this does not line up with the distinction between default and trigger cases, since her cases may be both default and generative. So generativity does not seem a good reason to refuse the title ‘testimonial

knowledge' to the trigger cases. Or so it seems to me; but we do not have to take a decision on this matter for present purposes, since my focus in this essay is not on the extension of testimonial knowledge, but rather on the cognitive methods we bring to trigger cases, whether the outcome ought ever to be counted as testimonial knowledge or not.

In cases of deliberation, the hearer will often have to make a second-order inference, an inference about the reliability of the speaker's own inference. Moreover, he must do this without having the speaker's evidence himself, though he does have the information (crucial, as TIBE will have it) that the speaker made the claim in question. The issue is how the hearer goes about making this second-order inference. Any adequate account of this process must meet a condition I have already mentioned: it must show how the hearer can assess the speaker's testimony without have independently to determine the truth value of what was said. We are trying to account for how our dependence on testimony is managed, not how it is avoided. Of course sometimes the effect of hearing someone say P is just to encourage us to find out for ourselves whether P, making no use of the speaker's testimony, but these are not the cases that interest us here. We want an account of how we manage testimony, not how we occasionally do without it.

Testimonial Inference to the Best Explanation

TIBE applies the general IBE scheme of inference to the particular task of assessing testimony. It is distinctively a form of IBE, because it has it that testimonial inference is an inference to the best explanation of the relevant evidence. And it is distinctive from other forms of IBE, because it takes the central datum to be explained not to be some natural phenomenon such as the red-shift of galactic light, but rather the fact that the speaker said what she did. It is an abductive inference from the fact of utterance to the fact uttered. The governing idea of TIBE is that when we are in evaluative mode, we infer that what we are told is true when its truth is part of the best explanation of the fact that the speaker said it (Elizabeth Fricker 1995, 404-5; Peter Lipton 1998, 23-31). As Nick Jardine (forthcoming) puts it, 'the basic criterion for the assessment of testimony is whether or not the production of the testimony is best explained by an account which implies its truth'.

In presenting this account, I want to focus on the evidence provided by the fact that the speaker said what she did, both because this is the primary datum according to TIBE and because this feature helps to distinguish the account from various alternatives. But I take it that TIBE allows for other types of evidence that clearly play a role in testimonial inference. For example, what is to be explained is often not just that the speaker said

what she did, but that *she* said it, and that she said it in the way she did, for example in a tone of exaggerated earnestness. There will also be evidence that has nothing to do with either the speaker or her present utterance. Thus the fact that she has been so reliable on these matters in the past encourages me to trust her this time. And clearly the decision whether to believe what one is told will have some dependence on the prior probability one assigned to what is asserted, which can itself be based on all sorts of evidence. But according to TIBE, it is the fact of assertion that is at least typically primary. Thus when I am told that all the Cambridge trains have been cancelled, but am unsure whether to believe the speaker, what I ask in the first instance is not why the trains were cancelled, but why the speaker claimed that they were. TIBE thus emphasises the role of the fact of assertion as evidence for whether the speaker should be believed, and the importance of determining whether the asserted fact follows from the best explanation of that evidence.

A couple more examples of TIBE in action may help. I need a new tire for my bike and I ask the person in the bike shop what type of tire I should get. He says that the best bet for me is the Roadrage XT, because it is hard-wearing and has good puncture-resistance. Should I believe him? Well, why did he recommend that tyre? Perhaps only because it would make him the most profit. Or perhaps he genuinely believes that it is my best bet, but he is incompetent to judge, having been brainwashed by the Roadrage sales representative. On reflection, I decide that neither of these possibilities – dishonesty or incompetence – supports a very good explanation of what he said. After all, I have gone to this person for my bicycling needs for years, and he has always treated me well. He has plenty of other brands of tyre in stock, and some of them are considerably more expensive than the Roadrage. So I infer that the reason he said what he did is because he knows that this tyre is the right one for me, from which it follows that what he said is true. So I believe him. By contrast, consider the following case. A man rang my doorbell and claimed that my rain gutters are loose. Should I believe him? They look fine to me, I know that he hasn't been up on the roof to inspect them properly, and I am further discouraged by the fact that he wants me to pay him today to fix them tomorrow. So I infer that the reason why he said what he did is not because he knows that my gutters do need work (though perhaps they do), but because he is hoping to make a fast buck. This unkind explanation does not entail that what he said is true, so I don't believe him.

TIBE is compatible with the idea that our default attitude towards what we are told is simple acceptance. According to TIBE, the switch from default to trigger is a switch from treating testimony as a non-inferential source of knowledge to treating it as evidence upon which an inference is based. At the same time, if you prefer a more thoroughgoing inferentialist picture of testimony, TIBE can be adapted to fit. Thus one might say that

even in cases where we seem uncritically to accept what we are told – say in the train case where the speaker is the station-master – we are still making an inference back to the asserted fact: the station-master said that the trains are cancelled because they are cancelled. As we have already seen, the question of whether all testimonial acceptance is inferential is interesting, but not crucial in order to begin to assess the merits of TIBE, the central claim of which is that where there is testimonial inference, this at least often takes abductive form where what is explained is why the assertion was made.

It is crucial to an account of the management of testimony that it make sense of indispensability of the evidence of the fact of assertion to the inference to the fact asserted. We want an account of testimonial inference, not testimonial avoidance. So it is a signal attraction of TIBE that although it has the hearer deciding whether to trust the speaker by determining whether what the speaker said is likely to be true, it avoids the circle of needing to find out independently whether what the speaker said is true before the hearer can decide whether to trust her. We do not need to know whether what the speaker told us is true in order to construct a potential explanation of why she told us, and the hearer is not determining the truth-value of what the speaker said independently of the speaker's utterance, since the fact of utterance is the hearer's central piece of evidence. In cases where the hearer ends up believing what the speaker said on the basis of TIBE, although he believes on the basis of his own inference, he still would not have believed it if the speaker hadn't said it.

TIBE accounts in a unified way for both incompetence and mendacity, since speakers' misapprehension and intent to deceive may each figure in explanations of why they said what they did. It also accounts for the diversity of factors that enter into testimonial inferences, including facts about speakers, the content of what they say and the manner and context in which they say it. And because TIBE is a species of a general account of inference (IBE), it is able to integrate diverse evidence, both testimonial and non-testimonial. As we have seen, this is important because the evidence on which we base our decision whether to believe what we are told is highly diverse, ranging from the fact that we were told to the personality, status and background of the speaker, to facts about her trajectory through the world, to facts about the diverse topics on which she may pronounce.

As Steven Shapin (from whom I have taken the happy expression 'management of testimony') has observed, the standard epistemological discussions in this area seems to involve a curious combination of maxims and countermaxims (Shapin 1994, ch. 5). Thus the fluency of the speaker's testimony is can be taken to count either for or against her honesty and competence. Hesitant delivery may be a sign of mendacity or of

incompetence; on the other hand it may be a sign of sincerity and care. As Hume says, ‘we entertain a suspicion concerning any matter of fact, when the witnesses...deliver their testimony without hesitation, or on the contrary, with too violent asseverations’ (1748, 171). TIBE seems to give a natural account of these sorts of factors, and of their curious maxim-countermaxim structure. For example, as Jardine (forthcoming) has observed, countermaxims arise in part because testimony is ‘reflexive’: a duplicitous speaker wants to sound credible and so will mimic the maxims. Positing such a clever strategy of deception may be part of the best explanation of the speaker’s utterance, and so lead us not to accept what we are told.

The final feature of TIBE that I will flag in this brief advertisement is that it is an account that supports a generative as opposed to a merely transmissive view of testimony, by showing how the hearer generates novel warrant for the testimony he ultimately accepts. In the triggered cases, the hearer is not simply co-opting whatever warrant the speaker had; the hearer is generating his own warrant, and one based on evidence different from that used by the speaker. She does not believe what she says because she said it. Seeing the fact of utterance as evidence fits well with a generative picture of testimony, a picture which is independently attractive.

Is Testimonial Inference to the Best Explanation Trivial?

The disjunction every philosopher dreads is: trivial or false. In this section I address the triviality worry about TIBE, turning to falsity in the next. One way of motivating the worry about triviality is to ask a quasi-Popperian question: what kind of epistemic behaviour would the account exclude? If virtually any possible form of inference could be given a TIBE gloss, then the account is not very informative. An account of inference should say something about how we do *not* think.

The triviality worry can come in at two levels. It may be a worry about the general IBE scheme, or a more specific worry about the application of this scheme in testimonial inference. One legitimate source of the general worry is the fact that so little has been said about what makes one explanation lovelier than another. For as I have noted, if IBE reduces to Inference to the Likeliest Explanation, then we seem to be coming close to saying as a little as ‘inquirers infer the explanation they judge most likely to be correct’. This is reminiscent to Dr Pangloss’s explanation of why smoking opium makes people sleepy by appeal to its dormative virtue. Inference to the Likeliest Explanation seems itself an unlovely explanation of inference.

There are a number of replies to the general worry. First, even the thin Inference to the Likeliest Explanation account may not be entirely devoid of content, since it is not

obvious that inferences are from evidence to explanations, even when our interest is not in explanation per se. The explanationist's claim is that we ask why not just because we seek understanding of what we already know to be the case, but also because we seek new knowledge. This claimed instrumental use of explanation in the service of inference is plausible but not trivial. Second, to say that we infer to the loveliest explanation may be to make a substantive claim even if no analysis of loveliness is given. Third, it may be possible to say something informative about loveliness, a project I have begun elsewhere (Lipton 2004).

The specific worry is that even if IBE in general makes a substantive claim, TIBE is relatively trivial. One form this worry might take is that, just because it is a species of IBE, TIBE fails to tell us what we want to know, namely what is distinctive about testimonial inferences. In reply one might say that the fact that TIBE subsumes testimonial inference under a more general account is a virtue, and that at the same time TIBE does treat testimonial inferences as special, because of the unusual nature of what is to be explained: the fact that the speaker said what she did. It is not trivial to say that the core evidence for assessing whether the speaker's utterance is true is the fact that the speaker made it. Indeed even if we take TIBE in the thin form of inference to the likeliest explanation of the utterance, the distinctiveness of the explanandum gives the account substantive content.

A different form of the specific worry is that TIBE does not exclude anything, because however we went about making testimonial inferences, these could always be construed as inferences to explanations of the fact of assertion. But this worry is unwarranted. To see why, here are three possible accounts of testimonial inference that are incompatible with TIBE. According to the first, the inference we make in trigger cases is based exclusively on evidence entirely independent of the testimonial episode. I have repeatedly observed this is not a good picture of testimonial management, but only of testimonial avoidance, My present point however is just that this account is incompatible with TIBE, which requires that the inference be based on that fact. So TIBE is not trivial.

A second alternative model of testimonial inference is one that appeals to facts about the testifier to the exclusion of the content of what is asserted, such as Shapin's account of the management of testimony in the context of the Royal Society of seventeenth-century England. Shapin's thought is that hearers decide what to believe by determining the social status of the speaker. For the Royal Society at that time, he claims that the fundamental testimonial maxim was to believe all and only what gentlemen assert. Now in TIBE too, status may play a role in testimonial inference, insofar as the speaker's role helps to explain why she said what she did. And this role may of course be truth relevant,

as Shapin emphasises. He claims that the high social cost in gentlemanly culture of being found to have lied prompted a gentleman to report only what he knew. From a TIBE perspective, this sort of information may figure in the explanation of why the gentleman said what he did and entails that what he said was true. Nevertheless, Shapin's status account of testimonial inference is strikingly different from TIBE. The status account appeals to far fewer factors than TIBE; most strikingly, it does not appeal to the content of the speaker's utterance. We are to ask who the speaker is, but not why she said what she did. On the TIBE account we believe some things a particular speaker says but not others, whereas on the status account it is all or nothing.

The third contrastive account I will offer is a simple appeal to track record. The idea here is that we decide whether to believe the speaker by considering the past claims he has made where we know independently the truth value of what he said (cf. Hume 1748, sec. 10). Although this is quite different from the status account, the track record account exhibits some similar contrasts to TIBE. Thus while a speaker's track record is surely sometimes relevant to testimonial inference, the track record account gives a much thinner account of testimonial inference than TIBE, which considers diverse kinds of evidence apart from truth value of past utterances, and involves many other considerations, especially the situation, beliefs and intentions of the speaker. The track record account also seems to have the consequence that the hearer gives the same credibility to everything a particular speaker says. And here again the source of this lack of discrimination arises from the lack of focus on the content of what is asserted, which plays a key evidential role in TIBE.

Just to complicate things a little, it should be noted that it is possible to develop both the status and track record accounts so they show some sensitivity to content. Thus in the case of the status account one might capitalise on the fact that a single speaker has multiple social statuses, and develop the account so that which status is relevant to trust is content sensitive. In this way, the same speaker could come out trustworthy on some matters but not others. (I owe this suggestion to Jim Secord.) As for the track record account assessments of reliability could be indexed to utterance type, so that a single speaker is registered as having a strong track record say about the time, but a poor one about the weather. (Tony Coady (1992) considers this option and the difficulties it faces.)

The contrast between these accounts and TIBE remains sharp. In neither of them is the fact that the speaker said what she did evidence for the truth of what she said. For this reason, there is a sense in which TIBE is nearly the opposite of these other two accounts. As I have indicated, TIBE can take account of status and track record. Nevertheless, in the case of the status account, what is said is taken to be true because the speaker is

trustworthy; according to TIBE, it is rather that the speaker is trustworthy because of the evidence that what she said is true. And whereas for the track record account, the hearer needs first to certify that the speaker is reliable, and only then accepts what she says; for TIBE, the first thing is to determine whether what the speaker said is likely to be true, from which the hearer may go on to infer that the speaker is reliable. These contrasts bring out a substantial claim that TIBE makes about testimonial inference, the claim that judgements about the truth of a particular questioned utterance may be prior to judgements about the general reliability of the informant, even when we have no independent access to the truth value of that utterance.

Is TIBE False?

So TIBE is not at any rate trivial. But if it is not trivial, it is probably false. So let us consider some plausible counterexamples. What about evidence for believing the speaker that plays no role in explaining the speaker's utterance? Jardine (forthcoming) gives a nice type of case in this category: evidence for the reliability of the speaker that is subsequent to the testimony in question. I am more inclined to trust what someone tells me about the finer points of squash technique if I subsequently see that he is a brilliant player, but presumably his subsequent behaviour does not play any role in explaining his earlier testimony. Similarly, evidence that describes events prior to the utterance may give reason to trust what is said without helping to explain why it was said. Thus, as we have already seen, the speaker's track record of past testimony that has been checked independently seems directly relevant to the assessment of his current testimony, yet at least often it does not help to explain it.

But evidence that does not explain the fact of utterance is not for that reason a counterexample to TIBE. The account does not require that evidence for the truth-value of what is said must explain why it was said; it is enough if that evidence is itself also explained by what explains the utterance. Thus the speaker's impressive track record may be well explained by her competence and honesty, which also explains why she said what she did on the occasion in question. I have emphasised the cardinal evidential role that the fact of utterance plays, according to TIBE, but the account must not be taken to exclude other relevant evidence. Like the fact of utterance, that other evidence must be explained, but it does not have to explain why the speaker said what she did.

A second natural worry is that we may be able to explain why the speaker said what she did in terms that do not entail the truth of what she said, even though we end up believing her. You tell me that there is an excellent production of *A Midsummer Night's Dream* on tonight in the King's Fellow's Garden, and I believe you, and I can explain why

you told me this because I know that you know that I am a great fan of Shakespeare comedies. But this perfectly adequate explanation obviously does not entail either that the play will be on tonight at King's or that it will be excellent. But TIBE can account for this phenomenon too. On the one hand, we may have a default case where we nevertheless generate a non-truth entailing explanation. TIBE does not need a truth-entailing explanation in default cases. On the other hand, we may be a trigger case. Here again we need not worry, once we remind ourselves that the same phenomenon may enjoy several compatible explanations. The explanation I gave of your utterance is fine, but there are also other explanations (or other explanatory factors) that do entail the truth of what is entailed, such as explanations that appeal to your knowledge of the Cambridge thespian scene and to your ability to judge plays. The point of TIBE is not that an explanation of an accepted assertion must entail that what was said is true; it is rather that when hearers make an inference to the truth of what is said, they will do so on the basis of an inference to such a truth entailing explanation. They are free to generate other, non-truth entailing explanation as well.

So TIBE can account for evidence apart from the fact of assertion, and it can acknowledge that not all explanations of the assertion will determine its truth-value. Nevertheless, I am not confident that a strict version of TIBE is leak-proof.

But there is a third worry that may be more problematic, and this concerns the inferential role of some of the hearer's background beliefs. Accepting testimony is not always a matter of the hearer adding to his web of belief: often it also entails revising that web, giving up things he formerly believed. Nevertheless, incompatibility between what the speaker says and some of the hearer's deeply held beliefs is often a reason for rejecting the testimony. Yet here there is no obvious explanatory link to the fact of utterance. Similar remarks apply to cases not of contradiction between testimony and background belief, but to contradictions between different speakers' testimony. If our speaker is contradicted by another speaker's testimony, this provides reason not to believe, but the second speaker's testimony may bear no explanatory relation to the first speaker's testimony. But here too the defender of TIBE has some kind of reply, since she can say that negative evidence, whether from background or from contradictory testimony, will be registered by making a truth-entailing explanation less attractive and so less likely to be inferred. For example, the background beliefs that contradict the testimony may also contradict the truth-entailing explanation of that testimony (indeed would always contradict it, if all explanations were deductive).

I hope I am not being complacent. A simplistic version of TIBE that took the fact of utterance to be the only evidence bearing on testimonial inference is bound to be

inadequate, and even the more permissive version that I have presented may still fail fully to account for all the factors that go into our decision whether to believe what we are told. But we are still left with the core idea that a central mechanism for assessing testimony involves an alien abduction, an explanatory inference as to why the speaker said what she did. The claim is not trivial, but it may nevertheless be correct.

Acknowledgements

My thanks to audiences at the Department of History and Philosophy of Science at Cambridge and at the EPISTEME conference on testimony held at Rutgers for very helpful questions and comments about alien abduction. I am especially grateful for the help of Alvin Goldman, Martin Kusch, Nick Jardine, Jennifer Lackey, Simon Schaffer, and Jim Secord.

References

Coady, C.A.J. (1992) *Testimony* (Oxford: Oxford University Press).

Fricker, E. (1995) "Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony", *Mind* **104**: 393-411.

Harman, G. (1965) "The Inference to the Best Explanation", *Philosophical Review*, 74, 88-95.

Hume, D. (1748) *An Enquiry Concerning Human Understanding*, T. Beauchamp (ed.), Oxford: Oxford University Press (1999).

Jardine, N. (forthcoming) 'Historical Testimony'.

Kusch, M. (2002) *Knowledge by Agreement* (Oxford: Oxford University Press).

Lackey, J. (1999) "Testimonial Knowledge and Transmission", *The Philosophical Quarterly* 49: 471-490.

Lipton, P. (2004) *Inference to the Best Explanation*, 2nd edition (London: Routledge).

Locke, John (1700) *An Essay Concerning Human Understanding*, 4th Edition, P.H. Nidditch (ed) (1975) (Oxford: Oxford University Press).

Shapin, S. (1994) *A Social History of Truth* (Chicago: University of Chicago Press).