

KNOWLEDGE, SCEPTICISM AND COHERENCE¹

Keith Lehrer

University of Arizona and University of Graz

I defend a coherence theory of knowledge.² It is a reply to the sceptic and at the same time an admission that we cannot prove the claims of the sceptic to be in error. We have much to learn from the sceptic, including what knowledge is. I shall articulate, however briefly, a coherence theory of knowledge. I shall then explain what sort of concession to the sceptic it permits. The theory is based on the assumption that we are fallible, that is, that we can err, no matter how well justified we might be in what we accept. Coherence yields a theory of justification, not of truth, but an adequate match between coherence and truth is all that is required to yield knowledge. Let us see how.

Knowledge is based on what we accept as true and on the truth of what we accept. But the acceptance of something true does not suffice for knowledge, for we may be unjustified in what we accept and yet accept something true. We may, in fact, proceed by the most irrational methods and fallacious reasoning and still be lucky enough to accept something true. My own view of acceptance is that it differs from belief in constituting a positive evaluation of belief at a metamental level of evaluation. Thus, acceptance, unlike mere belief, constitutes the best efforts of a person to obtain truth and avoid error. But our best efforts can go awry. What, then, must we add to convert our acceptance of something true to knowledge of it?

The answer is justification, but justification is the place where the sceptic dwells. We must enter, nonetheless. The sceptic raises objections to what we accept, whether it concerns tables, persons, galaxies or neutrinos. And the sceptic can be expected to raise doubts, hyperbolic doubts from Descartes or more mundane doubts from everyday life. The sceptic might appeal to a powerful demon, if demonology attracts her, or a malicious scientist manipulating our brain in a vat, if science fiction attracts her. Or she might appeal to our dreams, the illusions of sense, or even our familiar lapses of memory. How should I reply? What can I appeal to in order to argue that I know the sceptic is wrong in her hyperbolic and mundane doubts?

I can only appeal to what I accept. Where else would I seek premises to meet the objections of the sceptic if not the premises that I accept? I must either appeal to what I accept to meet the objections or put my hand over my mouth in silence. But the voice of reason speaks forth in terms of what I accept. There is no demon, and no one manipulating my brain either, at least not with a computer and electrodes. It is, the voice of reason tells me, more reasonable to accept that there is a table before me, that I am in a room full of people, than that the sceptical fancies are true. I face an either-or. I either appeal to what I accept or concede invincible ignorance. But I do not accept that it is reasonable that I concede, so I take the other route and meet the sceptical objections in terms of what I accept. Now if I can meet the hyperbolic and mundane objections of sceptics, including most saliently the sceptic within me, pertaining, for example, to the claim that I see a table before me, then I have at least a personal justification for accepting that I see a table before me. I am personally justified in accepting that. This sort of justification is coherence with my background system of states of acceptance, my acceptance system. Personal justification is coherence with my acceptance system. Coherence with my acceptance system is the ability to meet sceptical objections on the basis of it.

It is important to consider the nature of coherence and the acceptance system. My acceptance system consists of states of acceptance, states that might be described by statements of the form—I accept that *p*. Thus, the acceptance system does not consist of the thing accepted, namely, *p*, but, instead my acceptance of it, of *p*. The reason concerns the reasonableness of what I accept. My acceptance of what I accept, *p*, for example, makes the acceptance of *p* reasonable for me if I am trustworthy and, hence, reasonable in what I accept. Drop away my acceptance, and reasonableness drops away with it. My acceptance system, assuming I am trustworthy and reasonable in what I accept, allows me to meet objections in two ways. I may answer the objection by beating the objection. This means I can, by appeal to what I accept, conclude that it is more reasonable for me to reject the sceptical objection than to accept it³. Alternatively, I can meet the sceptical objection by neutralizing the objection. This means I can, by appeal to what I accept, conclude that the sceptical objection taken in conjunction with something I accept renders the objection irrelevant because the conjunction is as reasonable as what the sceptic alleges. The states of my acceptance system provide the replies to the sceptic. They exhibit the way in which what I claim to know coheres with the acceptance system. The acceptance system consists of states that may be used as a map of replies to objections.⁴ These replies exhibit coherence with the system and yield personal justification on the basis of it.

The sceptic, who has her doubts about her, will note that personal justification may be completely founded on error. Indeed, if her doubts are based on truth, then my personal justification is based on error and defeated by it. Justification defeated by error is useless to convert anything into knowledge. If, however, what I accept to meet the sceptical objections is true, then my justification is undefeated by error. Undefeated justification of something I accept is what is

required for the conversion to knowledge. It exhibits the needed match of coherence and truth.

To require a complete match between what I accept and what is true to achieve undefeated justification would be unrealistic and unnecessary, however. It will suffice for undefeated justification to restrict the acceptances used to meet skeptical objections to the subset of my original acceptances that are true, that is, to t-acceptances. We may think of my original system of acceptances combined with a system of t-acceptance, the t-system, as constituting an *ultrasystem*, which adds the t-system to the acceptance system.⁵ Undefeated justification is justification that acknowledges what is accepted in the original acceptance system but restricts the acceptances used to meet objections to achieve undefeated justification to members of the t-system. Undefeated justification is, therefore, a kind conditional justification because it is conditional on the t-system. The point of requiring undefeated justification based on the ultrasystem to acknowledge original acceptances, even though acceptances used to meet objections are restricted to the t-system, is to restrain undefeated justification in terms of the original acceptance system. In particular, if the person accepts something false, say *f*, which precludes personal justification of a target proposition, *p*, on the basis of the acceptance system, then acknowledgment of the acceptance of *f* in the ultrasystem precludes undefeated justification of *p*. This insures that undefeated justifications will be a subset of personal justifications and not surreptitiously introduce new justifications by the elimination of error.⁶

Thus something a person accepts is justified on the basis the ultrasystem if and only if all objections can be met, that is, beaten or neutralized by appeal to the t-acceptances acknowledging the existence of other acceptances of the ultrasystem. The existence of all acceptances of the ultrasystem must be acknowledged, but the content of the original acceptances is bracketed off as unusable for purposes of meeting skeptical objections unless it is also the content of a t-acceptance. Justification for accepting something is undefeated if and only if it is justified on the basis of the ultrasystem. Undefeated justified acceptance is knowledge.

The advantages of the coherence theory of knowledge just briefly sketched are profound, for they permit me to solve all the problems of knowledge. The sceptic who proposes that I do not know of the thoughts or feelings of others, of the mental life of others, for they may be robots, is met with the reply that they are not, that is more reasonable for me to accept that others have thoughts and feelings on the basis of my acceptance system than to accept that they are robots. The sceptic about the past who says I do not know anything about the past because the world might have sprung into existence a minute ago just as it now is, receives the answer that it is more reasonable for me to accept that the world has existed for a very long time on the basis of my acceptance system than that it came into existence five minutes ago. Other sceptical doubts receive similar answers based on my acceptance system. Sceptical doubts about the *a priori*, for example, are answered in the same way, by the affirmation of the reasonableness of accepting what we do rather than sceptical hypotheses of deception. The answers to the

sceptic based on coherence with our background system provide a unified theory of empirical and *a priori* knowledge. Moreover, these answers are based on a systematic account of ourselves and capacities articulated within the acceptance system. If I am right in what I accept to meet the objections to a sceptical claim, whether about the empirical or *a priori*, then my coherent personal justification will be undefeated and convert into knowledge. I will know that the sceptic is wrong.

It is important to notice that this reply to the sceptic admits the *possibility* of the truth of the sceptical hypotheses at the same time that it denies their actual truth. What they show, the sceptical hypotheses, is that we are fallible. We can be deceived. The account of justification as undefeated coherence is an argument to the effect that it suffices for knowledge that we are not deceived in what we accept to meet the objections of the sceptic, though we must admit the possibility of the truth of those hypotheses and, therefore, our fallibility. If we are deceived in the possible way the sceptic imagines, we shall fail to detect the deception. This concession of our fallibility amounts to the admission of a gap between personal justification and truth. The personal justification that we have based on our acceptance system cannot guarantee the truth of what we accept. We can be fully justified in accepting something on the basis of our acceptance system and yet be in error.

It is this fallibility that has given importance to the coherence theory. There is an objection to the coherence theory that has long been thought to be decisive against it. It is called the *isolation argument* or the *isolation objection*. It is simple enough to express. Any system of acceptances may be isolated from reality. If my acceptance system is isolated from reality, then it may be a system of mistakes, or one great systematic mistake, and, therefore, not a suitable basis for knowledge. This objection is one that must be met, but the recognition of fallibilism, which says, in effect, our most fully justified acceptances may be false, reveals that the problem is a problem for any theory of knowledge and is not specific to the coherence theory. If it is a fact, which I think it is, that our most fully justified beliefs may be false, then any theory faces the isolation objection. Our most fully justified beliefs may be false, and so every theory of knowledge, not simply the coherence theory, must face the isolation argument.

What then is the answer to the isolation objection given by the coherence theory? What is the answer to the objection that our evidence, however convincing, may be deceptive? The answer offered by the coherence theory is that personal justification requires that these objections can be answered by the acceptance system. To put the answer in the first person, I accept that I am not isolated from reality, that I am not deceived, that my faculties are connected with reality and are not fallacious. It is part and parcel of any justification that I have for any specific thing that I accept about the world that I am not isolated from the world and that my evidence about the world is not deceptive. Of course, my accepting that I am not isolated or that my evidence is not deceptive, does not make it true. If what I thus accept is not true, then my justification is defeated and I am ignorant. If, on the other hand, it is true that I am not isolated or deceived in my evidence, then

personal justification as coherence converts into knowledge. I have called this argument the *transformation argument*. Any sceptical objection must be met by what I accept for me to be personally justified in accepting some specific claim. So, if what I accept to meet the sceptical objections is true, then personal justification transforms into knowledge by the acceptance of truths that meet the sceptical objections.

The voice of the sceptic is not yet quieted, however. For she will reply that even if I turn out to be right in what I accept, my acceptance is but a mere assumption. She may continue by claiming that if what I accept is a mere assumption, then it is a mere surd of reason, something that is unexplained and irrational. I may, she concedes, argue for things that I accept on the basis of other things I accept, but I shall in this manner either be led into a regress or reach a point at which I claim that what I accept just is reasonable without argument and be left with an unexplained surd of reason.

To reply to the sceptic, I must argue that the things I accept are reasonable while avoiding the regress and the surd. But how can I do it? There is a way, and I shall follow it. But let me indicate at the outset that my goal is modest. I only wish to argue that what I accept is something that it is more reasonable for me to accept than not to accept. Given the existential choice, to accept or not to accept, I wish to argue that when I do accept, that is more reasonable.

How can I argue that I am reasonable in even this modest way, in a way that falls short of being fully justified in what I accept but may serve as the systematic basis for converting acceptance that is justified and undefeated by error into knowledge? The answer has been suggested above and must now be made explicit. I have acknowledged our fallibility. Even Descartes found no infallible starting point, though he thought he did. He assumed that he doubted and thought. This was his starting point. But an extreme materialist would reject that claim arguing that what Descartes said was thought was nothing more than neural activation in the brain, and, in scientific truth, there is no room for thought. The extreme or eliminative materialist may be wrong, but the hypothesis is possible, and it shows that there is no starting point that carries an infallible guarantee of truth.

A positive answer to the question of whether I am reasonable in what I accept which avoids the regress and surd must accept our fallibility and acknowledge that we must proceed without any guarantee of truth. To acknowledge that I am fallible, however, is compatible with the accepting that I am, nevertheless, worthy of my trust in what I accept. I may have a capacity to be trustworthy in what I accept even though I am fallible and have no guarantee of success. I do not have to be perfect in order to be trustworthy or to be worthy of my own trust. As an analogy, I might hire a guide in a foreign city, Istanbul, for example, whom I consider to be trustworthy and worthy of my trust to guide me through the city, even though I know that even the best guides in such a complicated city are not perfect and sometimes lose their way.

If I am to answer the sceptic without or within, I must ask myself whether I should accept that I am trustworthy, at least for myself, that is, whether I am

worthy of my trust. If the answer is that I am not, then I have nothing I can say and must be silent. But, if, on the contrary, I accept that I am trustworthy, I can argue that I am reasonable in what I accept while avoiding the sceptic and the surd. Let us see how the argument runs.

I begin with the principle

A. I am trustworthy in what I accept.

This principle does not entail that I am trustworthy in everything that I accept, for I am fallible in my trustworthiness, but it does provide a reason, a kind of inductive one from the existence of a capacity to the successful exercise of it. Now suppose I accept that p , that I see a table. I may conclude from A that I am trustworthy in accepting that p , that I see a table. If I am trustworthy in what I accept, then I am reasonable in what I accept. The argument could be amplified, but the point is a simple one. My trustworthiness serves the objectives of reason, and if I am trustworthy in the way I serve the objectives of reason in what I accept, then I am reasonable to accept what I do. The simple form of the argument is that if I am trustworthy in accepting something, then I am reasonable in accepting it.

But how does this argument avoid the regress and the surd when the question is raised by sceptic of why I am reasonable to accept the principle A itself, that I am trustworthy in what I accept? I could argue for my trustworthiness by consideration of other things I accept and my success in attaining truth, but that way a regress threatens, whatever the merits of such arguments in supporting the principle. Must I just insist that the principle A of my trustworthiness just is reasonable without explanation of why and embrace the surd? Fortunately, that is not necessary. The reason is that principle A explains why it is reasonable to accept principle A. I accept that I am trustworthy in what I accept, and if I am trustworthy in what I accept, then I am reasonable in accepting that I am trustworthy in what I accept. My trustworthiness in what I accept explains why I am reasonable in accepting that I am trustworthy in what I accept. In short, just as principle A explains why it is reasonable for me to accept the other things that I accept, so it explains why it is reasonable for me to accept principle A itself.

The foregoing reply depends on the truth of the claim that I am trustworthy. It is not a proof of the truth of that principle. In fact, I cannot prove the truth of that basic principle on which the reasonableness of my acceptance system depends. I would need to use principle A to prove the reasonableness of accepting it and would argue in a circle. To argue in a circle proves nothing. If, however, I am, in fact, trustworthy as I accept, my trustworthiness explains my reasonableness in accepting that I am trustworthy in what I accept. Explanation and proof divide. The truth of principle A suffices for the explanation of why it is reasonable to accept it and for the avoidance of the surd, but not for a proof of the truth of the principle.

So, proof cannot be circular but explanation may be. The reason is a simple and familiar one. When we come to explain things, assuming our explanations are finite, we either end with some principle which is unexplained, a kind of explan-

atory surd, or some principle must explain not only other principles but itself as well. We must choose between the surd and the loop. The advantage of the loop is that nothing need be left unexplained. Those who seek to maximize explanation will prefer the loop, as I do, but I have no proof that anyone should seek to maximize explanation in philosophy or anywhere else. The preference for leaving nothing unexplained and entering the loop of explanation as a result is one I act upon in developing my philosophy. I do not pretend to offer any proof that one ought to proceed in this way, but there is no fallacy in the explanatory loop. Of that I am confident. There is no fallacy in maximizing explanation by application of a loop.

Externalists such as Dretske, Goldman and most recently Plantinga⁷ among others have sought to avoid the loop by arguing that epistemic conditions are satisfied as the result of some relationship between belief and something external to it. I have argued against such theories elsewhere and do not wish to repeat my argument here⁸. My argument rests on the idea that if a person does not consider herself to be worthy of her trust in what she believes, then her belief is not knowledge, though it may be useful information, no matter how the belief is related to external matters. If, on the other hand, a person accepts that she is worthy of her trust in what she accepts and is in a position to meet objections to this claim, then she need only be correct in what she accepts to attain knowledge. I do not deny that some relationship of what a person accepts to something external to it is necessary for knowledge, for knowledge requires the truth of what one accepts, but I do deny that such an external relationship suffices. It cannot suffice without consideration of whether one is worthy of one's own trust in what one accepts and an understanding of what that requires. The loop of trustworthiness turns acceptance upon itself to explain reasonableness and justification from within the loop.

How is my trustworthiness in what I accept related to the truth of what I accept? Does my trustworthiness in what I accept consist of a high rate of success in accepting the truth? One might hope that one has a high frequency of obtaining truth in what one accepts when one is trustworthy in what one accepts, but I do not see any guarantee of it. In the first place, being trustworthy is a matter of how one changes what one accepts, and the ways or methods one uses to change in order to correct errors in what one accepts. If one is trustworthy in the way one changes one's ways, then one is trustworthy in what one accepts because of that even if one is not now very successful in obtaining truth. Moreover, one's trustworthiness depends on being trustworthy in the evaluation of trustworthiness of others on whom one depends. If one is trustworthy in evaluating others and changing what one accepts on the basis of such evaluations, one may be trustworthy in what one accepts even if one is not very successful in obtaining truth before one thus changes what one accepts.

Secondly, in science we are often trustworthy in what we accept even if we are more often wrong than right. The methods of science are aimed at the acceptance of powerful and comprehensive theories on the basis of scientific methods, but the frequency of truth in such matters cannot be expected to be high. We proceed by correcting our errors without any guarantee of truth or even a high

truth frequency. If one asks how often we must be right in order to be trustworthy in what we accept, the answer is only this. We must be right a trustworthy amount of the time. It is a mistake to think of our trustworthiness as simple high truth frequency reliability. Trustworthiness is a matter of being worthy of trust and turns irreducibly on the notion of what is worth accepting and what methods are worth using. There is an irreducible element of epistemic value coiled at the center of the life of reason.

My acceptance of my trustworthiness, concerning what I accept and how I reason, is a keystone of my acceptance system. Is it a foundation? That is the wrong metaphor. A keystone in the arch of acceptance fits into the top of the arch and holds the arch together as it is supported by the other stones in the arch. The principle of my trustworthiness is a keystone of my acceptance system and sustains coherence within it. The reasonableness of what I accept confirms the use of what I accept to meet the objections to specific claims yielding coherence with the acceptance system and personal justification based upon it. When justification is undefeated, knowledge results.

Do I know that I am trustworthy in what I accept? I do because my justification is undefeated. Do I know that I know? I accept that I know, and my justification is again undefeated. So I know that I know. Does my justification guarantee that I am trustworthy in what I accept? It does not. Does it guarantee that my justification is undefeated? I have no guarantees. I cannot prove to the sceptic that I am trustworthy or that my justifications are undefeated by error. But I can know and know that I know, for I do not need proof or a guarantee of truth to know or know that I know. With that knowledge, I rest fallibly content with my reply to the sceptic. I cannot prove that she is wrong, but I know she is wrong.

Notes

1. An earlier version of this paper was presented to the National Academy of Science of Hungary under the title "Knowledge and Scepticism" and was published in Hungarian in the *Proceeding of the Hungarian National Academy of Science*.
2. My most complete articulation of coherence theory is contained in *Theory of Knowledge*, (Boulder and London: Westview Press and Routledge, 1990), Chaps. 6–9, and my most recent formulation is contained in *Self-Trust: A Theory of Reason, Knowledge and Autonomy*, (Oxford: Oxford University Press, 1997), Chap. 2. In the latter work, I consider the background system to be an evaluation system consisting of preferences, reasonings and acceptances instead of just the latter. But only the sub-system of the evaluation system consisting of the acceptance system is relevant here.
3. This is a minor alteration of my earlier work concerning the notion of beating a competitor. Earlier I defined beating a competitor as being more reasonable to accept the target proposition than the competitor, but that allows for the possibility that it might be more reasonable to withhold on both the target proposition and the competitor than to accept either.
4. I am indebted to Frank Jackson for the idea of a map of replies. Cf. David Braddon-Mitchell and Frank Jackson, *Philosophy of Mind and Cognition*, (Oxford and London: Blackwells, 1996), 187–195.

5. This is a new conception of the ultrasystem. I formerly supposed that members of the system contained just states of acceptance, some original members of the acceptance system and other replacements of original acceptances of things that were false by acceptances of their denials. For explanation, see next note.
6. I have altered my conception of members of the ultrasystem to include only original acceptances that are true while retaining other acceptances but disallowing their use to meet skeptical objections. My reason for the change is to overcome an objection raised by Peter Klein and John Pollock to the effect that deleting false acceptances would have the untoward result that our justification for accepting that we accept that p would be defeated in case p was false.
7. Fred Dretske, *Knowledge and the Flow of Information*, (Cambridge, Mass.: MIT Press, 1981), Alvin Goldman, *Epistemology and Cognition*, (Cambridge, Mass.: Harvard University Press, 1986), Alvin Plantinga, *Warrant and Proper Function*, (New York: Oxford University Press, 1993).
8. In *Theory of Knowledge*.