

Lehrer's dynamic theory of knowledge

Hans Rott

Department of Philosophy

*University of Regensburg**

Version of 22 January 2002, 24:50

Abstract. This paper argues that problems of epistemology and problems of belief change are tightly interwoven and that a successful analysis of knowledge depends on a proper solution of the problems of belief change. Following Keith Lehrer, we assume that what an agent knows at a certain instant is dependent upon how he is inclined to change his beliefs in time. However, two of Lehrer's fundamental concepts, personal justification and undefeated justification, turn out to be problematic. I argue that systematic solutions to the problems must have recourse to a well-developed theory of belief change. While this argument tries to establish that the theory of (the static concept of) knowledge needs help from the theory of the dynamics of belief, it is finally argued that the reverse strategy should be followed in studies of belief change where the crucial epistemological distinction between foundationalism and coherentism proves useful for the analysis of belief representation and revision.

Keywords: Epistemology, theory of knowledge, justification, logic of belief change

1. Introduction

Philosophers must not be allowed to confuse epistemic and doxastic concepts. It is their duty to clarify the subtle interconnections between knowledge and belief. As this is too formidable a task for a single paper, I will not develop an epistemological theory of my own, but rather focus on Keith Lehrer's influential theory of knowledge as elaborated in his 1990 book *Theory of Knowledge*. This book represents only one stage of the development of Lehrer's epistemology. It is the successor of, and shows considerable overlap with, a book with the title *Knowledge* published by the same author in 1974. The basic structure of the 1990 definition of 'knowledge' is later retained in Lehrer (1997, Chapter 2) and duplicated in an analogous definition of 'wisdom'. The recent second edition of the *Theory of Knowledge* (Lehrer 2000) presents a concept of knowledge that is much simplified as compared to the one of the first edition.¹ The present paper, however, is based mainly on the more 'dynamic' 1990 version of Lehrer's book. Its purpose is, first, to draw attention to some problematic features of Lehrer's account and, second, to argue that a proper understanding of knowledge does not

* 93040 Regensburg, Germany, hans.rott@psk.uni-regensburg.de

¹ Cf. the remarks made at the beginning of Section 4 below.



only require an understanding of belief *simpliciter*, but in addition a thorough understanding *of the dynamics of belief*.

The main point of this paper is to show that a good theory of belief revision is necessary for a proper development of a theory of knowledge. I shall later also argue that the theory of belief revision can profit from a study of the concepts that have evolved in recent epistemology. We ought to have a basic picture of what knowledge is, and of how knowledge can be obtained, when we search for principled solutions of the problems of belief change. In Rott (2001), I exploit the fundamental distinction between foundationalist and coherentist accounts of knowledge that has played a central role in the epistemological literature of this century.² I want to argue that this distinction applies more properly to theories of *doxastic* states than to theories of *epistemic* states. It is unfortunate that philosophers have tended to focus on knowledge without attending equally closely to the—seemingly less problematic—notion of belief.

Another distinction is not going to be treated in this paper. Externalist accounts of knowledge argue that people or animals need not be able to justify their knowledge, but rather hold that there must be some reliable (causal, counterfactual, nomological) connection between the knower and the things known. Indeed it makes perfect sense to say that, for instance, the dog knows where in the garden its favourite bone is buried, but I want to focus on an internalist and what Lehrer (1990, pp. 4, 36) calls ‘characteristically human sort of knowledge’. To my mind, the verb ‘to know’ is ambiguous and its variant meanings reflect different and conflicting intuitions—a fact that has caused much unnecessary dispute in recent epistemology. I do not claim to cover all uses of ‘to know’.

2. Epistemology, knowledge representation and revision

Arguably, the representation of pieces of knowledge should not differ from the representation of beliefs. From a first-person perspective there is little if anything that allows one to distinguish between mere belief and real knowledge, and it is doubtful whether we should expect the representation of belief and knowledge to represent more than what is accessible to the reasoner or reasoning system itself.

Now let us suppose, for the sake of argument, that questions concerning the representation of knowledge and belief have been answered to everybody’s satisfaction. Many interesting questions are still left open.

² See, for instance, Lehrer (1990), Plantinga (1990) or Sosa (1980).

Solutions regarding the concept and the representation of knowledge do not automatically give answers to questions concerning the dynamics of belief and knowledge. How is 'knowledge', or better: alleged knowledge, revised in the light of new evidence? How should it be revised? Agents are fallible, and what they consider to be knowledge quite often turns out to be false—and hence not to be proper knowledge at all. At this stage the focus of attention gets shifted from knowledge to belief, and normative questions are seen to become increasingly important. Epistemology, the philosophy of mind and psychology analyse and describe what knowledge is and how it is obtained and represented in human beings. Knowledge representation has to do with normative issues in so far as there are many different approaches to representing information, some of which are 'better' and some of which are 'worse'—with respect to what we demand from the relevant 'knowledge systems' (e.g. computational tractability, efficiency, comprehensiveness, reliability, transparency). The change of alleged knowledge, that is, the change of belief and acceptance systems, is intrinsically beset by normative problems as well. It surely is reasonable to ask for an 'ethics' of belief and acceptance. But it seems that this question can be set aside when we talk about knowledge. In so far as 'knowledge' is perceived to be in need of revision, however, it is perceived to be inferior to real knowledge, and should rather be accorded the status of belief, opinion, prejudice or some similar sort of doxastic (rather than epistemic) term. Problems of the ethics of belief arise in so far—and perhaps only in so far—as problems of a genuinely dynamic sort arise. The ethics of belief in the traditional understanding is primarily concerned with the question when it is rational or justified to *adopt* some new belief, thus involving an act of belief *acquisition*.³ The additional problem of belief change is that we have to face the question when to *eliminate* or *replace* which of the previously held beliefs, thus involving acts of belief *dislodgement*.

Prima facie, it appears that the questions posed by epistemology, knowledge representation and 'knowledge' revision are, though related, clearly separable. In any case, there does not seem to be a close connection between the theory of knowledge and the theory of belief revision. It is the thesis of the present paper that precisely this is illusory. Even if we do not contemplate any problems of the 'intermediate' field of knowledge representation, we can find very close interdependencies of epistemology and belief revision. More precisely, I want to illustrate that

³ For the exciting ethics of belief debate that took place in the nineteenth century, see the anthology edited by McCarthy (1986).

- (i) the analysis of knowledge requires a proper solution of the problem of belief revision
- (ii) the analysis of belief revision should not be conducted without a proper understanding of the concepts and categories that have been used in study of knowledge

We must actually restrict claim (i) considerably because the following considerations will be based on the particular epistemological theory of Keith Lehrer. I take Lehrer to be following in great detail a trail that was started by Plato, in a beautiful passage in one of his earlier dialogues:

True opinions too are a fine thing and altogether good in their effects so long as they stay with one, but they won't willingly stay long and instead run away from a person's soul, so they're not worth much until one ties them down by reasoning out the explanation. . . . And when they've been tied down, then for one thing they become items of knowledge, and for another, permanent. And that's what makes knowledge more valuable than right opinion, and the way knowledge differs from right opinion is by being tied down. (*Meno* 97e–98a, Plato 1994, p. 69)

Most of the following considerations will not depend on the details of Lehrer's theory; in fact I shall offer a few non-trivial improvements on some of his definitions. But I want to base my arguments on the overall architecture of Lehrer's undertaking. If Lehrer were completely misguided, then what I say about the relation between epistemology and the theory of belief revision might equally well be mistaken.

Claim (ii) needs to be qualified as well. The analysis of belief revision is not dependent on features that distinguish genuine knowledge from mere belief. It is rather dependent on the structure and the formation of beliefs as they are relevant in the theory of knowledge. What I have in mind above all is the fundamental distinction between *foundations* and *coherence* theories of knowledge. This distinction happens to have come to the fore in the theory of knowledge, but it may just as well be placed in a theory of belief.⁴ It is primarily concerned with the inferential relations between various beliefs, that is to say, with the internal structure of our belief systems. The contrast lies in the answer to the question whether there is such a thing as a *belief base*, that is, a distinguished set beliefs that are not in need of an inferential justification by other beliefs and that taken together inferentially justify all the remaining ('derived') beliefs. All of this can be dealt with in a

⁴ Wolfgang Spohn has pointed out to me that this point is fully explicit in *BonJour* (1985).

general theory of belief; nothing requires to refer this topic to the theory of knowledge. Claim (ii) can still be upheld, given the fact that many relevant aspects of the structure of beliefs have as a matter of fact come out most clearly in epistemological discussions.

3. Central concepts of Lehrer's theory of knowledge

In this section we unroll the central parts of Lehrer's theory of knowledge and show how they are rooted in problems and questions which belong to the theory of belief change. The following presentation is based on the summary in Lehrer (1990, pp. 147–149). The time parameter t which does not play any interesting role in Lehrer's theory will be removed.

For a long time it has been thought in philosophy that knowledge is justified true belief. The short and famous article by Gettier published in 1963 has made it clear that this analysis is inadequate. Let us have a look at one of Gettier's counterexamples.

Suppose that Smith has very strong evidence for

ϕ : Jones owns a Ford.

Suppose further that Smith is totally ignorant of Brown's whereabouts. Still he can (and does) correctly infer from ϕ that

ψ : Either Jones owns a Ford, or Brown is in Boston.

χ : Either Jones owns a Ford, or Brown is in Barcelona.

ξ : Either Jones owns a Ford, or Brown is in Brest-Litovsk.

By pure coincidence, and entirely unknown to Smith, Barcelona happens to be the place where Brown actually is. However, Jones does *not* own a Ford. Now, χ is a *true justified belief*, since ϕ is a justified belief and χ may be logically derived from ϕ , and the second disjunct of χ is true. However, it would be utterly counterintuitive to say that Smith *knows* that χ is true, because he believes that χ is true 'for the wrong reasons'. He has just been lucky that the right belief occurred to him.⁵

⁵ Actually luck is not even needed. If we do not require that a belief manifests itself in an occurrent event in the believer's mind, but may reside hidden in an implicit theory of his or hers, we may assume that the beliefs of an agent are logically closed. Then any justified false belief ϕ gives rise to an infinite set of justified true beliefs—all disjunctions $\phi \vee \phi'$ where ϕ' is any arbitrary truth. So Smith is not just lucky, but his false belief that Jones owns a Ford will automatically generate infinitely many justified true beliefs like χ .

Gettier's example has led many epistemologists to the conclusion that knowledge is *more* than justified true belief. They do not discard this venerable definition of knowledge, but supplement it by a fourth clause. Lehrer's suggestion is the following:

DK S knows that ϕ if and only if

- (i) S accepts that ϕ ,
- (ii) it is true that ϕ ,
- (iii) S is completely justified in accepting that ϕ , and
- (iv) S is completely justified in accepting that ϕ in a way that does not depend on any false statement.

Although Lehrer (1990, pp. 10–11) separates acceptance from belief, I do not think that it would make a big difference if we substituted belief for acceptance in clause (i).⁶ Similarly, clause (ii) is not very controversial in the theory of knowledge. The essential parts of the definition are the last two clauses which appeal to the problematic notions of justification and dependence. I have substituted in clause (iv) a formulation taken from Lehrer (1990, p. 18) for the formulation in his official summary in Lehrer (1990, p. 147). In the latter he requires S to be completely justified in accepting that ϕ 'in a way that is not defeated by any false statement'. This statement seems to be slightly screwed up. What Lehrer means, I think, is that S is completely justified in accepting that ϕ in a way that cannot be defeated by pointing out that the justification relies on a false statement. The above phrasing of clause (iv) expresses this more accurately than Lehrer's official formulation.⁷

Lehrer further characterizes the concept of knowledge in a two-layered strategy.⁸ First he develops the notion of personal justification which is based on an agent's subjective acceptance system. In a second step the purely subjective standpoint gets transcended by several

⁶ According to Lehrer, we sometimes believe ϕ for the sake of felicity or the pleasure of believing so, but we would not accept ϕ for these reasons. Lehrer says that there is always a potential conflict between the 'ancient system of perceptual belief' which is the 'yield of habit, instinct, and need' and the 'truth-seeking ... scientific system of acceptance'. While the former is an 'automatic input system', the latter is 'capable of ratiocination'. (Lehrer 1990, pp. 113–114) In this paper, I work on the simplifying assumption that all doxastic attitudes are 'aiming at' truth: If an agent either believes or accepts that ϕ , this means that he *holds ϕ true*.

⁷ Lehrer (1990, p. 138) apparently thinks that the two formulations are 'equivalent'.

⁸ Compare in particular Lehrer (1990, pp. 141–152) and (1997, 28–45).

operations on the agent's current belief set that could help him to approximate the whole truth.⁹

3.1. PERSONAL JUSTIFICATION AND THE COMPARATIVE REASONABLENESS OF ACCEPTANCE

All of Lehrer's considerations are based on the notion of an acceptance system which is defined as follows.¹⁰ We slightly adapt the notation.

D1 A system X is an *acceptance system* of S if and only if X contains just statements of the form, S accepts that ϕ , attributing to S just those things that S accepts with the objective of accepting that ϕ if and only if ϕ .

This definition could be amended a little by taking into account that on Lehrer's account, there is only one acceptance system for each agent at a certain time. So it seems that D1 should better start like that: 'A system X is *the* acceptance system of S if and only if ...'

Lehrer is one of the main advocates of a coherence theory of knowledge. According to this approach, all justification comes from coherence with a given acceptance system X . There is no justification *simpliciter*, only justification on the basis of some X .

D2 S is *justified in accepting* ϕ on the basis of system X of S if and only if ϕ coheres with X of S .

What is needed now is of course an elucidation of 'coherence-with-a-system'. Lehrer defines it in terms of the relations of competing, beating and neutralizing of propositions.

D3 ϕ *coheres* with X of S if and only if all competitors of ϕ are beaten or neutralized for S on X .¹¹

Interestingly, Lehrer does not seem to require that ϕ belongs to X for D2 and D3. The function that takes an acceptance system X and yields back the set of all sentences cohering with X , or equivalently,

⁹ These operations are carried out only hypothetically, 'for the sake of argument' in fictitious test dialogues with an omniscient sceptic. In reality, the agent's state of mind remains unchanged.

¹⁰ This is the term used in Lehrer (1990); in Lehrer (1997, pp. 25–29) and Lehrer (2000) the leading part is taken by the 'evaluation system' which includes not only a person's accepted propositions (geared to truth) but also her preferred propositions (geared to merit).

¹¹ Definition D3 on p. 148 of Lehrer (1990) actually starts as follows: ' S is *justified in accepting* ϕ on the basis of system X of S if and only if ...' From definition D2 and the surrounding text, however, it is obvious that this is a misprint and the formulation given above is the intended one.

of all sentences S is justified in accepting on the basis of X , may be interpreted as an inference operation in the sense of Rott (2001).

- D4 ψ *competes* with ϕ for S on X ¹² if and only if it is more reasonable for S to accept that ϕ on the assumption that ψ is false than on the assumption that ψ is true, on the basis of X .
- D5 ϕ *beats* ψ for S on X ¹³ if and only if ψ competes with ϕ for S on X , and it is more reasonable for S to accept ϕ than to accept ψ on X .
- D6 χ *neutralizes* ψ as a competitor of ϕ for S on X if and only if ψ competes with ϕ for S on X , but $\psi \wedge \chi$ does not compete with ϕ for S on X , and it is as reasonable for S to accept $\psi \wedge \chi$ as to accept ψ alone on X .

Intuitively, definitions D4 and D5 are not beyond reproach. Let us assume, for instance, that it is extremely reasonable to accept that ϕ , and that ψ weakens the reasonableness of accepting ϕ just a little bit, by a more or less negligible degree. Should we say that ψ ‘competes with’ ϕ ? In what sense would we speak of a competition? ψ may be about quite another subject matter than ϕ , so there cannot be much rivalry or conflict between the two. And again, if it is still a little more reasonable to accept ϕ than to accept ψ , should we say that ϕ ‘beats’ ψ ? Surely in the situation just described it would make good sense to accept both ϕ and ψ (assuming that it is reasonable to accept ψ in the first place), even if ϕ beats ψ in Lehrer’s sense. We are ready to accept sentences that weaken each other a little, and we tend to connect such sentences by ‘although’. If we say, ‘They get along together very well, although they have no interests in common’, in symbols ‘ ϕ although ψ ’, we do accept both ϕ and ψ . It does not matter that ψ ‘competes with’ ϕ in Lehrer’s sense, and it does not matter whether ϕ ‘beats’ ψ or not.

The last three definitions, D4 – D6, all lead us to a comparative concept ‘reasonable-to-accept’, relativized to a given acceptance system X . Before turning to a brief discussion of that concept, we finish off the first, subjective part of Lehrer’s analysis of knowledge. For personal justification, it is just the agent’s current acceptance system which is the basis for judgements of coherence.

- D7 S is *personally justified in accepting* that ϕ if and only if S is justified in accepting that ϕ on the basis of the acceptance system of S .

¹² Some renaming of Lehrer (2000): ‘ ψ is an *objection* to ϕ for S on X ’.

¹³ Lehrer (2000): ‘objection ψ to ϕ is *answered* for S on X ’.

Now of course everything hinges on what is meant by ‘reasonable-to-accept’—again, relativized to a given acceptance system. Surprisingly, Lehrer does not say very much about this, but he considers it as an advantage that his primitive term of reasonableness is open to many different interpretations. In the few pages he devotes to the topic (Lehrer 1990, pp. 127–131), however, he recommends to employ *cognitive decision theory*.¹⁴ Lehrer suggests to identify the degree of reasonableness of accepting a hypothesis ϕ with its *expected epistemic utility*:

$$r(\phi) = p(\phi) \cdot Ut(\phi) + p(\neg\phi) \cdot Uf(\phi)$$

where $p(\phi)$ and $p(\neg\phi) = 1 - p(\phi)$ are the probabilities of ϕ being true or false respectively, and $Ut(\phi)$ and $Uf(\phi)$ are the positive or negative utilities of accepting the hypothesis ϕ , when ϕ is true or false, respectively. The utility $Ut(\phi)$ is supposed to reflect the informativeness of ϕ , and possibly other virtues such as ϕ 's explanatory power, simplicity, or pragmatic value, and the advantage of conserving existing beliefs (Lehrer 1990, p. 131).

It should be noted that this absolute degree of reasonableness of accepting is not quite sufficient for what we need in order to understand the foregoing definitions. In the definition of competition, Lehrer appeals to the reasonableness of accepting ϕ *on the assumption that ψ is true or false*. What we need, then, is something like expected *conditional* epistemic utilities $r(\phi|\psi)$ and $r(\phi|\neg\psi)$, and it is left unspecified how we can get them. There is no problem with the well-known concept of conditional probabilities,¹⁵ but it is not quite clear whether the utilities of accepting a hypothesis ϕ *conditional on accepting ψ or $\neg\psi$* should be thought of as different from the plain, unconditional utility of accepting ϕ . But if the utility of accepting a sentence is dependent on accepting some other sentence, should not the utility of accepting a sentence be sensitive to the context of acceptance, that is, to the acceptance system X as a whole?

This question leads us to a problem that is both more general and more important. As indicated in definitions D4 – D6, the reasonableness to accept a statement may be—and probably should be—relative to the

¹⁴ The most eminent advocate of cognitive decision theory is Isaac Levi (1967; 1984). For a critical voice, see Weintraub (1990). In Lehrer (1997, in particular pp. 30–35), the contribution of cognitive decision theory to personal justification has vanished and its role is taken over by the principle of trustworthiness, a kind of rationality principle that is geared solely to the acquisition of truth. Cognitive decision theory is still present, however, in the second edition of the *Theory of Knowledge* (Lehrer 2000, pp. 144–148).

¹⁵ Let us assume that $p(\psi)$ and $p(\neg\psi)$ are positive, so that conditionalizing by either ψ or $\neg\psi$ is not beset with the problem of an ill-defined division by zero.

acceptance system of the agent. But the above definition of $r(\phi)$ does not reflect this. It is rather an absolute measure of reasonableness. This is quite contrary to the coherentist's aim of evaluating *systems* of hypotheses rather than *single* hypotheses taken in isolation. Lehrer can counter this objection by saying that the utility functions Ut and Uf depend on the current acceptance system X . But then one may ask whether it is illuminating to base an analysis of 'coherence of ϕ with a system X ' on an unexplained notion of 'utility of accepting ϕ on the basis of system X '.¹⁶

Appealing to cognitive decision theory suggests that the acceptance of a proposition is a matter of decision. Such an assumption is at least controversial. But in this respect Lehrer's replacement of 'belief' by 'acceptance' is a prudent move. It is certainly much more plausible to say that an agent decides to *accept* something than that he decides to *believe* something. Another potential point of criticism is that it need not be (objective?, subjective?) *probability* that is taken into account when the reasonableness of accepting certain hypotheses gets assessed. Perhaps *plausibility*, a notion with different formal characteristics, is an equally suitable candidate. This objection, too, could be countered, by pointing out that decision theory is simply a theory based on probabilities.¹⁷

3.2. VARYING IMPERATIVES FOR COHERENCE

The most serious questions for Lehrer's account seem to be the following. How can we be sure that coherence in his sense, explained by means of a complicated mechanism of competing, beating and neutralizing, which in turn is based on a decision-theoretic criterion, will give us a coherent acceptance set?¹⁸ If every element of an acceptance set coheres with that very set (as it presumably should), can we be sure that the acceptance set is consistent?¹⁹ Will the acceptance set be closed under logical consequences? Both consistency and closure are

¹⁶ In his brief discussion of expected epistemic utilities, Lehrer shows little awareness of the fact that Ut and Uf may or should depend on X , if reasonableness of acceptance is to be relative to the current acceptance system. Lehrer contrasts $Ut(\phi)$ with ϕ 's (objective or subjective) probability and links it to ϕ 's truth, but not to the acceptance of other sentences. He does not say anything about Uf .

¹⁷ Theories of plausibility of the kind I have in mind are offered, amongst others, by Grove (1988), Rescher (1976), Shackle (1961), and Spohn (1988). They would have to be supplemented by a qualitative decision theory.

¹⁸ The distinction between relational coherence (coherence as a relation) and systemic coherence (coherence as a property of a system) is discussed in connection with Lehrer's theory by Olsson (1999).

¹⁹ This problem for Lehrer's theory has also been treated by Olsson (1998).

themselves requirements of—inferential—coherence. If all of Lehrer's central definitions are finally based on a numerical degree of reasonableness, why not take that very same degree as the sole arbiter of acceptance, without the taxing detour via definitions D2 – D7?

All these questions call for a fresh and systematic look at the principles that are involved in Lehrer's concept of coherence.

First and foremost, we have coherence according to Lehrer's own theory. The corresponding imperative is this: *Accept precisely those sentences the competitors of which are beaten or neutralized!*

Second, we have seen that Lehrer's theory is grafted on top of cognitive decision theory which of course brings along its own standards of coherence. Lehrer does not address this point. An attempt to make it explicit reveals that there are several ways to go. The following idea seems initially plausible: *Accept those sentences that promise the greatest expected cognitive utility!* A moment's reflection, however, shows that this idea is premature. Why should we reject sentences of positive (or at least non-negative) expected utility for the sole reason that there are still more useful ones? Don't all sentences with positive r -values contribute to the overall expected utility? Thus the right imperative seems to be this: *Accept all sentences with positive (or non-negative) expected cognitive utility!*²⁰ But perhaps this is again mistaken. It may be wrong to suppose that the utilities of individual sentences simply add up to yield the utility of the whole body of beliefs. This would mean that we cannot rely on the above imperatives, because conditional utilities (see above) may be very different from the unconditional ones. For instance, I may entertain two alternative hypotheses, both with high expected epistemic utility, which contradict each other. Accepting either one of them seems reasonable, but accepting both of them would lead to inconsistency which is not particularly useful. Thus one should take into account interactions between the individual beliefs, and regard beliefs as constituting a system that may be assessed *only* holistically. This is very much in the spirit of coherentists anyway, who argue that only the corpus of belief or knowledge taken as a whole is a proper unit of epistemic appraisal. The degree r of reasonableness of acceptance then must not be applied to individual sentences, but to *sets* of sentences, and corresponding imperative reads thus: *Accept those sets of sentences that promise the greatest expected cognitive utility!*

A third concept of coherence is the logical or inferential one. I shall discuss its implications in detail later, but it is expedient to anticipate

²⁰ This precept is plausible only if one assumes, as Lehrer apparently does, that the expected utility of *rejecting* a hypothesis is zero. An alternative idea would be to accept just those propositions ϕ which have a degree of reasonableness that exceeds a contextually fixed threshold value.

the main points already here. A set of sentences is inferentially coherent if it is consistent and closed under consequences. The corresponding imperative is: *Accept all the logical consequences of what you accept, but avoid accepting contradictions!*

In Rott (2001, Chapters 7 and 8), it is shown that the second and the third concepts of coherence are compatible with one another. But even if one is ready to grant Lehrer's mechanics of justification in the sense of his definitions D3 – D6 some plausibility in itself, it is doubtful whether it can be made to cohere with other concepts of coherence. One has to face the fact that different coherence criteria may conflict with one another, and decide which of these criteria are the most justified ones. As of 1990, Lehrer's theory is a hybrid of at least three different intuitions.

3.3. UNDEFEATED JUSTIFICATION AND JUSTIFICATION GAMES

In order to arrive at knowledge one has to go beyond the actual acceptance system of an individual agent. In a sense very close to the passage of Plato's *Meno* quoted in Section 2 above, knowledge must be stable under criticism. In Lehrer's 'justification games' the part of the critic is taken over by an omniscient 'sceptic'. Let us now have a look at the formal definitions that take Lehrer from merely personal justification to complete and indeed indefeasible justification.

D8 A system V is a *verific system* of S if and only if V is a subsystem of the acceptance system of S resulting from eliminating all statements of the form, S accepts that ϕ , when ϕ is false.

As in the case of definition D1, it would be preferable to say that the V mentioned in D8 is *the* verific system of S , since it results from the unique acceptance system of S (at a given time) by just cutting out the false beliefs. The verific system is the basis for verific and complete justification:

D9 S is *verifically justified in accepting* that ϕ if and only if S is justified in accepting that ϕ on the basis of the verific system of S .

D10 S is *completely justified in accepting* that ϕ if and only if S is personally justified in accepting ϕ and S is verifically justified in accepting ϕ .

Complete justification in this sense, however, does not solve the Gettier problem. Smith's belief that Jones owns a Ford need not depend on any false belief. Let us suppose that Jones told Smith that he has a Ford, showed him papers stating that he, Jones, owns a Ford, and

always drives a Ford on his way from his home to his office. All this is believed and known (in some pre-theoretical sense) by Smith, and justifies his belief that Smith does in fact own a Ford. So the reason for Smith's not knowing that sentence χ above is true is not that he accepts false sentences, but rather that he is not aware of all true sentences that are relevant to the case. Notice that it is not enough to know *some* relevant facts since a biased selection of true facts may be utterly misleading and turn the agent away from some other truths. Lehrer suggests to solve this difficulty by looking at what he calls the 'ultrasystem' of an agent. Here are his definitions.

D11 *S* is justified in accepting that ϕ in a way that is undefeated if and only if *S* is justified in accepting ϕ on the basis of every system that is a member of the ultrasystem of *S*.

D12 A system *M* is a member of the *ultrasystem* of *S* if and only if either *M* is the acceptance system of *S* or results from

- eliminating one or more statements of the form, '*S* accepts that ψ ', when ψ is false,
- replacing one or more statements of the form, '*S* accepts that ψ ', with a statement of the form '*S* accepts that not ψ ', when ψ is false,
- or any combination of such eliminations and replacements in the acceptance system of *S*

with the constraint that if ψ logically entails χ which is false and also accepted, then '*S* accepts that χ ' must also be eliminated or replaced just as '*S* accepts that ψ ' was.

Before moving on to the criticism of Lehrer's, we should mention his key result: 'Knowledge reduces to undefeated justification, a just reward for our arduous analytical efforts.' (Lehrer 1990, p. 149) Clearly, undefeated justification implies personal justification: the actual acceptance system of *S* is a member of the ultrasystem; it also implies verifical justification: the sceptic can make *S* eliminate all his false beliefs, thus effecting a transition to the agent's verifical system;²¹ finally, it also implies truth: if ϕ were false, the sceptic could make *S* eliminate ϕ .

Lehrer's theory of knowledge can be called a *dynamic* one because we can think of the repeated operations mentioned in D12 as (potential) steps in a journey through the space of belief states.

One may wonder why Lehrer requires for the undefeated justification of ϕ that *every* member of the ultrasystem must support ϕ . It would

²¹ We neglect the possibility that the removal of a false belief may tear along some true beliefs.

seem sufficient that S is *ultimately justified*. By this we mean that for every member M of the ultrasystem of S there is another member M' of the ultrasystem of S which improves on M and on the basis of which S is justified in accepting ϕ . That M' *improves on* M means, of course, that M' can be reached from M by some combination of ‘truth-conducive’ eliminations and replacements of the kind specified in definition D12. This concept seems more adequate since even if S knows that ϕ , pre-theoretically understood, a mischievous sceptic may well advance an impressive battery of true facts speaking *against* the truth of ϕ , so that S loses his confidence that ϕ is true. Only later in his conversation with the omniscient sceptic, when S comes to know more about the truth, will he regain his old true and justified belief. Although the correct belief would be dropped on the receipt of true but misleading information, we may consider it to constitute knowledge, since one can later learn that this information has in fact been misleading. Temporary doubts about ϕ should not count, so it seems, as long as all potential paths of the ultra justification game finally lead to ϕ ’s acceptance.²² Lehrer himself seems to agree with that when discussing his Grabit example:²³

Suppose I see a man, Tom Grabit, with whom I am acquainted and have seen often before, standing a few yards from me in the library. I observe him take a book off the shelf and leave the library. I am justified in accepting that Tom Grabit took a book, and, assuming he did take it, I know that he did. Imagine, however, that Tom Grabit’s father has, quite unknown to me, told someone that Tom was not in town today, but his identical twin brother, John, who he himself often confuses with Tom, is in town at the library getting a book. Had I known that Tom’s father said this, I would not have been justified in accepting that I saw Tom Grabit take the book, for if Mr. Grabit confuses Tom for John, as he says, then I might surely have done so, too. (Lehrer 1990, p. 139)

Lehrer summarizes the lesson to be drawn from this example as follows: ‘I may be said to know that Tom Grabit took the book despite the fact that, had I known what his father said without knowing about

²² Could there be an eternal wiggling of the acceptance value of ϕ in response to the sceptic’s challenges? Not if we neglect the possibility that a truth-conducive change can make S drop truths (as we decided to do in footnote 21) and if we set aside questions of infinity. In such a context, the sceptic has the means to make S accept the true and complete theory about the world which, of course, cannot be further improved.

²³ A similar example about barns and papier-mâché facsimiles originally due to Carl Ginet is discussed in Nozick (1981, pp. 174–175) and Bach (1984, pp. 40–41).

his [the father's] madness, I would not know whether it was Tom who took it.'

In contrast to Lehrer's undefeated justification, ultimate justification no longer implies that S is justified on the basis of his current personal or verific acceptance systems. But what we have been looking for is an *objectified* notion of justification, which can be conjunctively added to the (at least partially) subjective notions or personal and verific justification. In particular, we cannot dispense with verific justification if we want to end up with the right analysis of Gettier type examples. In the example discussed on page 5, Smith is both personally and ultimately justified in believing that either Jones owns a Ford or Brown is in Barcelona (ξ). What prevents him from knowing this is that his belief that Jones owns a Ford is false, so he is not verifically justified to believe that ξ .²⁴ Ultimate justification is not sufficient for knowledge.

Since every theory can be improved by a transition to *the one true and complete theory about the world*,²⁵ ultimate justification reduces to justification on the basis of that theory. On the one hand, it does not seem objectionable to call for that theory as the final arbiter of knowledge. On the other hand, I cannot see why the whole truth must be coherent, in Lehrer's or in any other but the purely logical sense. Shouldn't we try to avoid stipulating that the one true and complete theory is coherent, because that would mean basing epistemology on a questionable metaphysics? Similarly, I do not see any intuitive reason why every truth should be justified, on the basis of the true and complete theory. If this is right, then it is *impossible*, by Lehrer's own definition of knowledge as well as by the definition using ultimate justification, that an agent will know the whole truth. Shouldn't we try to avoid this conclusion? This suggests that undefeated justification may not be necessary for knowledge, and even ultimate justification may not be.

But let us stop with these cosmic speculations now and return to more definite matters again. In Lehrer (2000, especially pp. 153–154, 168–169), there are no replacements any more, and there is no talk of strong corrections. In this new account, the ultrasystem is closer

²⁴ This argument depends on the assumption that if ξ is not part of an acceptance system, then it cannot be justified on the basis of that system. Strictly speaking, Lehrer does not make that assumption; compare Lehrer's definitions D2 – D6 and especially my comment on D3 above.

²⁵ Saying this actually steps beyond Lehrer's account, which does not provide for the possibility of knowledge *expansion* through the sceptic—which marks an important difference between Lehrer's concept of replacements and the usual understanding of the concept of revision. The uniqueness involved in talking about 'the one true and complete theory about the world' is of course relative to the language used, which I assume as given.

in essence to what was called the verifical system earlier,²⁶ and undefeated justification is similar to verifical justification, i.e., justification on what remains when everything false is eliminated from the person's acceptance system.

Unfortunately, Lehrer does not tell the reader why he has changed his earlier definitions and given up on the idea that not only eliminations, but replacements, too, may be prompted by the sceptic. It is not clear whether he just considers it as a simplification of his former account or whether he thinks that the new edition of his book actually corrects an inadequacy of his former account. I presume that the reason lies in the problem of misleading information. In cases like the Grabit example, receipt of information about what Tom Grabit's father said (without information about the father's mental state) would probably have done away my acceptance that Tom Grabit stole the book, even though this seems to be a bit of genuine knowledge. So it appears that according to Lehrer, we should not admit replacements or additions to our stock of accepted propositions when testing for knowledge. This is an interesting argument, but its validity may well be doubted. Misleading effects cannot only be achieved by adding truths but also by removing falsehoods. There are other cases of a similar structure where in fact no genuine knowledge seems to be involved. Lehrer (2000, p. 160) discusses a 'newspaper example' originally due to Harman (1973). I do not think, however, that the distinguishing criterion offered does the job that it has been assigned by Lehrer. In the newspaper example, the subject's justification is said to depend on "unstated" facts about the newspaper's trustworthiness. But then, why are we not to assume, for instance, that the justification of our belief that Grabit stole the book depends on the unstated fact that no-one has offered any evidence about look-alike suspects? If we can't rule out this, we don't seem to know that Grabit stole the book, according to Lehrer's definition, and the strategy of using only eliminations rather than both eliminations and replacements does not help.

In sum, then, I cannot see that there is an epistemologically significant difference between the sceptic's removing errors and his supplying new truthful information. Sometimes misleading evidence that one does

²⁶ I neglect here, perhaps uncharitably, the fact that Lehrer's second edition uses richer 'evaluation systems' instead of the 'acceptance systems' of the first edition. The presentation of Lehrer (2000) follows that of Lehrer (1990) more closely than is warranted by its contents. While the first edition has the ultrasystem as a genuinely new and complex system, in the second edition the ultrasystem is nothing more than the pair consisting of the original system and the verifical system. Introducing the term 'ultrasystem' for this entity seems a bit pompous, but it is just a reflection of how the book came into being.

not possess may block the claim to knowledge, just as wrong beliefs that one does possess may do.²⁷ Whenever the omniscient sceptic succeeds in making us abandon a belief as a result of an improvement of our belief set, this is strong indication that the belief had not been a piece of knowledge. Knowledge, so it seems, should be stable in any kind of critical, truth-directed dialogue. For this reason I will stick to the richer 1990 definition of ultrasystems.

4. Lehrer's logical constraints for eliminations and replacements, and how to improve them

We have seen that *eliminations* and *replacements* of accepted propositions are of paramount importance for Lehrer's approach to the theory of knowledge. Eliminations and replacements are respectively called 'weak corrections' and 'strong corrections' in Lehrer (1997, pp. 45–49). These operations are very close to the operations of contraction and revision as they are known in the theory of belief revision.²⁸ As Lehrer places no constraints on the structure of acceptance systems (see Definition D1), there seems to be no need for him to apply non-trivial change operations to acceptance systems.²⁹ May we not just eliminate a false sentence ψ by simply dropping the statement 'S accepts that ψ ' from the system X , and similarly, may we not simply substitute the statement 'S accepts that $\neg\psi$ ' for the the statement 'S accepts that ψ ' in X , when a false sentence ψ is to be replaced by its negation?

The answer is, 'No'. Effortless eliminations and replacements are excluded by the *logical constraint* stated in Lehrer's definition D12: 'if ψ logically entails χ which is false and also accepted, then " S accepts that χ " must also be eliminated or replaced *just as* " S accepts that ψ " was.'³⁰

²⁷ This is reminiscent of the idea that the justification for ϕ does not only consist in the presence of reasons for ϕ , but also in the absence of reasons against ϕ . The point is given pride of place in nonmonotonic reasoning in the tradition of Doyle (1979).

²⁸ I have mentioned an important difference between 'replacements' and 'revisions' in footnote 25: Replacements substitute $\neg\phi$ for some previous belief ϕ , while revisions include expansions of belief sets by sentences about which there had not been any opinion before.

²⁹ If this were true, then Lehrer would turn out to be a foundationalist (at least, a foundationalist in the sense of Chapter 3 of Rott 2001).

³⁰ My italics. Compare footnote 11 on p. 194 of Lehrer (1990): '... with the constraint that if ψ logically entails χ , which is false and also accepted, then " S accepts that χ " must also be eliminated or replaced *in the same way as* " S accepts that ψ " was.' (Again, my italics) I am reading the phrases "in the same way as" (in Lehrer's

These are the only constraints Lehrer enters into his official definitions, but in the running text he acknowledges more constraints of a similar kind. They are in a sense complementary to the ones we just mentioned. Let ψ again be the false sentence to be eliminated or replaced by its negation $\neg\psi$. While the first group of constraints concerns (false) sentences *implied by* ψ , the second group deals with (necessarily false) sentences *implying* ψ . The first group of constraints is forward-looking, the second group is backward-looking. Here is the quotation from Lehrer (1990, p. 141):

The sceptic ... may require the claimant to eliminate anything the claimant accepts that is false, and the claimant must eliminate the specified item from his acceptance system *and at the same time* eliminate anything he accepts that logically implies the eliminated item. Or the sceptic may require the claimant to replace anything the claimant accepts that is false with the acceptance of its denial *and at the same time* replace anything that logically implies the replaced item with acceptance of its denial. (My italics)

In order to make the discussion of *Lehrer's logical constraints* more easily surveyable, I will now give shorter, semi-formalized formulations.

- (FE) *Forward Elimination.* If ψ is to be eliminated and $\psi \vdash \chi$, then item χ , if false, must be eliminated.
- (FR) *Forward Replacement.* If ψ is to be replaced by $\neg\psi$ and $\psi \vdash \chi$, then item χ , if false, must be replaced by $\neg\chi$.
- (BE) *Backward Elimination.* If ψ is to be eliminated and $\chi \vdash \psi$, then item χ (which is false) must be eliminated.
- (BE) *Backward Replacement.* If ψ is to be replaced by $\neg\psi$ and $\chi \vdash \psi$, then item χ (which is false) must be replaced by $\neg\chi$.

I have not mentioned here that only χ s that are accepted should possibly be eliminated or replaced. This should be self-evident.

I am not going to discuss the falsity conditions in (FE) and (FR) which presume an impartial, objective, omniscient supervisor for the eliminations and replacements in the ‘ultra justification game’. I want to reflect on the logical constraints only in so far as they are accessible to the agent himself. In doing so, I want to avoid the assumption that the agent is or should be omniscient; however, I do embrace the idealizing assumption that he or she is capable of drawing all and only valid logical inferences that can be drawn on the basis of their acceptance systems.

footnote 11) and “just as” (in D12) as indicating that a replacement of ψ by its negation should enforce a replacement of χ by its negation.

The point I want to make is that Lehrer's logical constraints are too simplistic. Being based on consequence relations between single sentences, they in effect consider the accepted items *in isolation* rather than as items in an acceptance *system*. The constraints do not really address the logical coherence of a belief with all its surrounding beliefs. It is important to take into account the *context of the remaining accepted items* when formulating logical constraints for eliminations and replacements. Without any claim that these are 'the right' constraints, the following ones are certainly more adequate in that they show some sensitivity to the context in which beliefs are situated. We keep on using the variable ' ψ ' for the false sentence that is to be eliminated or replaced by its negation, and give both a formulation that is close to Lehrer's own statements and a slightly more formalized version.

(FE⁻)/(FR⁻) *Forward Elimination/Forward Replacement.*

If ψ is an *essential* premise for the logical derivation of χ and χ is false, then ' S accepts that χ ' must also be eliminated. In symbols:

If ψ is to be eliminated or replaced by $\neg\psi$, and $F \cup \{\psi\} \vdash \chi$ for some set F of accepted items which are not eliminated, but $F \not\vdash \chi$ for any such F , then item χ (if false) must be eliminated.

(FR⁺) *Forward Replacement.*

If the addition of $\neg\psi$ creates a *new implication* of ξ (and ξ is true) then ' S accepts that ξ ' must be added to the acceptance system. In symbols:

If ψ is to be replaced by $\neg\psi$, and $F \cup \{\neg\psi\} \vdash \xi$ for some set F of accepted items which are not eliminated, then item ξ (if true) must be added, if it is not already accepted.

(BE⁻)/(BR⁻) *Backward Elimination/Backward Replacement.*

If F is a set of accepted premises that logically implies the eliminated or replaced item ψ , then for at least one (false) member χ of F , ' S accepts that χ ' must be eliminated. In symbols:

If ψ is to be eliminated or replaced by $\neg\psi$, and $F \vdash \psi$ for some set F of items, then at least one member of F (which is false) must be eliminated.

These conditions correct a number of counterintuitive features of Lehrer's constraints. In the case of a replacement of ψ by $\neg\psi$, there is no reason to *replace* by their negations all the χ s that are either critically implying ψ or essentially implied by ψ . It is enough that such χ s are *eliminated*. Let us temporarily employ the following concepts of critical and essential implication. A sentence χ *critically implies*

another sentence ψ in the context of a set F , if F implies ψ , $F \setminus \{\chi\}$ does not imply ψ and χ is one of ‘the weakest’ or ‘most vulnerable’ elements of F . A sentence χ is *essentially implied* by another sentence ψ in the context of a set F , if F implies χ , but $F \setminus \{\psi\}$ does not imply χ . Critically implying and essentially implied χ s are to be eliminated when ψ is just eliminated without being replaced by its negation. And in this respect, there is no reason at all to think that replacements present problems any different from those presented by contractions.

However, a replacement occasions one kind of adjustment of the acceptance system which cannot be initiated by an elimination. Since there is an addition of $\neg\psi$ to the old acceptance system, it is reasonable to require that any new derivations that are made possible by this new item should in fact be made, and the results be added to the acceptance system. This is condition (FR⁺) which has no counterpart in Lehrer’s definitions.

The most important respect in which the above constraints improve upon Lehrer’s constraints is that they pay attention to the fact that ψ or any of the χ s mentioned are part of a system. The F s mentioned in the constraints represent the contexts in which the respective tests for implications have to be made. The contexts are subsystems of the original acceptance system of the agent. By making the contexts explicit, we raise important new questions: Which items may be included in those F s that figure in the forward-looking constraints? Which items should be excluded from those F s that figure in the backward-looking constraint? More generally, how do we know which items in an acceptance system survive the process of elimination or replacement? Lehrer should give answers to these questions if his theory of knowledge is to be considered complete, but he fails to do so. It is precisely the theory of belief change that addresses these questions and offers a variety of ways to answer them.

The constraints (BE⁻) and (BR⁻), which concern sets F that imply the false belief ψ that has to be given up or replaced, leave much room for choices. It requires the agent to give up at least one (false) belief in F , but it does not tell us *which* belief or beliefs ought to be given up. The constraint does not fully determine what to do, but leaves us the freedom to choose. But how are the choices which members to eliminate from the acceptance set to be made? Even if one cannot enumerate concretely the cognitive values that govern our decisions what to give up, it is possible to inquire into the logic of ‘coherent’ or ‘rational’ choices involved in belief change (Rott 2001).

A choice-theoretic perspective has already come up in Lehrer’s suggestion to explicate the notion of personal justification. We saw that he ends up with a decision-theoretic explication of the comparative

notion of reasonableness-of-acceptance. According to this approach, it is more reasonable to accept ϕ than to accept ψ just in case ϕ has greater expected epistemic utility than ψ . But this is a comparison between two single items of belief only, and it is not clear how it can be extended to a criterion for assessing whole systems of belief. In fact it is not at all evident that the *logical* constraints that Lehrer mentions on his way from personal to undefeated justification mix well with the *decision-theoretic* advice given in the case of personal justifications. It has to be shown that a combination of logical and choice-theoretic constraints is indeed possible. In Rott (2001), I have tried to show this independently of Lehrer's particular idea that go for maximal or at least non-negative expected epistemic utility. Using choice-theoretic methods, one can maintain the idea that acceptance systems resulting after some change ought be inferentially coherent, that is, consistent and closed with respect to a given logic. Choice and logic can indeed coexist in harmony.³¹ It is unclear whether Lehrer's coherence mechanics in terms of competing, beating and neutralizing is compatible with such an account, but what seems to be clear is that both logic and the theories of choice and decision are better motivated and more principled than Lehrer's coherence mechanics. In case of conflict, they might therefore take priority.

5. Epistemology and belief change – a symbiotic relationship

In this paper, I have used concepts and ideas borrowed from belief revision theory to elaborate on an important contemporary account in the theory of knowledge. But the symbiosis between the two research areas may equally well be viewed from the opposite perspective. In this concluding section, I want to give an indication of how concepts and ideas developed in epistemology can help to structure and interpret much of the work that has been in belief revision theory.

One of the most relevant distinctions for belief revision is that between *foundationalist* and *coherentist* approaches in epistemology. Lehrer (1990, p. 13) characterizes the fundamental difference between foundationalist and coherentist views of knowledge as follows.

According to *foundationalists*, knowledge and justification are based on some sort of foundation, the first premises of justification. These

³¹ Another fundamental idea which is not mentioned in Lehrer's account is that doxastic changes should be *conservative*, i.e., that they should incur only *minimal* changes to the previous acceptance system. On this idea, and its role in belief change theories, compare Rott (2000).

premises provide us with basic beliefs that are justified in themselves, or self-justified beliefs, upon which the justification for all other beliefs rests.

Coherentists argue that justification must be distinguished from argumentation and reasoning. For them, there need not be any basic beliefs because all beliefs may be justified by their relation to others by mutual support. [My italics, HR]

Ernest Sosa (1980, pp. 23–24) makes essentially the same point in more metaphorical terms:

For the *foundationalist*, every piece of knowledge stands at the apex of a pyramid that rests on stable and secure foundations whose stability and security does not derive from the upper stories or sections.

For the *coherentist* a body of knowledge is a free-floating raft every plank of which helps directly or indirectly to keep all the others in place, and no plank of which would retain its status with no help from the others. [My italics]

It is important to see that the categorical distinction between foundationalism and coherentism can more properly be applied to theories of belief than to theories of knowledge, since no reference is made in the drawing of this distinction to the question of whether the beliefs in question are actually true.³² There is also a second, more ‘dynamic’ sense in which the distinction becomes relevant, and that is when it comes to the modelling of the dynamics of belief.³³ So there are two questions that separate foundationalists from coherentists.

- (1) Can a distinction between basic beliefs and derived beliefs be validly drawn?
- (2) And if so, are changes of beliefs made primarily on the base level or on the level of ‘coherent theories’?

It is clear what the foundationalist’s and the coherentist’s answers will look like. The former affirms while the latter denies the first question. In response the second question, the former would say ‘on the base

³² As Lehrer realizes very clearly, this objective question has to be linked to the subjective question of coherence in a separate step.

³³ There has been an ongoing controversy over the coherentism-vs.-foundationalism issue in the belief revision literature for more than a decade, see Nebel (1989), Gärdenfors (1990), Doyle (1992), Nayak (1994), del Val (1994; 1997), Hansson and Olsson (1999), Bochman (2000; 2001) and Rott (2001). Particularly influential as a mediator between epistemology and belief change was Harman (1986).

level', while the latter, lacking a distinguished base level, must opt for the level of coherent theories. It is necessary to work out in greater detail the concepts and distinctions on which these answers are based. In Rott (2001) I have tried to provide a framework that helps us to understand the issues involved and to characterize two fundamentally different *perspectives* on the process of belief revision. These perspectives turned out to be related to, but not to be identical with the dichotomy between foundationalism and coher. I did not go as far, however, as Hansson and Olsson (1999) who argue that the coherence theory in the epistemologist's sense is trivialized in the context of the coherence perspective on belief revision.

When transferring the idea of 'foundations' of knowledge to the area of belief change, it is not particularly important whether the basic beliefs are true, let alone infallibly true. Neither is it important that they are justified to such a high degree that they may be regarded as certain. In the current theories of belief change, belief bases are not supposed to carry any of these connotations. Basic beliefs are distinguished from derived beliefs only by the fact that they are somehow 'given', either explicitly or as things that are taken for granted. Givenness is not at all supposed to imply indefeasibility here. Still, as I tried to show in Rott (2001), these categories coming from epistemology can be exploited for an illuminating analysis of the dynamics of belief.

Acknowledgement

This paper is a revised version of the second chapter of Rott (2001). I want to thank Gordian Haas, Keith Lehrer, Erik Olsson, Wolfgang Spohn and all the other participants of the 2000 Konstanz Workshop on Lehrer's epistemology for many helpful comments, even though they will certainly be dissatisfied with what I made out of them.

References

- Bach, Kent: 1984, 'Default Reasoning: Jumping to Conclusions and Knowing When to Think Twice', *Pacific Philosophical Quarterly* **65**, 37–58.
- Bochman, Alexander: 2000, 'A Foundationalist View of the AGM Theory of Belief Change', *Artificial Intelligence* **116**, 237–263.
- Bochman, Alexander: 2001, *A Logical Theory of Nonmonotonic Inference and Belief Change*, Springer Verlag.
- BonJour, Laurence: 1985, *The Structure of Empirical Knowledge*, Harvard University Press, Cambridge, Mass.

- del Val, Alvaro: 1994, 'On the Relation Between the Coherence and Foundations Theories of Belief Revision', *AAAI'94 – Proceedings of the National Conference of the American Association on Artificial Intelligence*, pp. 909–914.
- del Val, Alvaro: 1997, 'Non-Monotonic Reasoning and Belief Revision: Syntactic, Semantic, Foundational, and Coherence Approaches', *Journal of Applied Non-Classical Logics* **7**, 213–240.
- Doyle, Jon: 1979, 'A Truth Maintenance System', *Artificial Intelligence* **12**, 231–272.
- Doyle, Jon: 1992, 'Reason Maintenance and Belief Revision: Foundations vs. Coherence Theories', in Peter Gärdenfors (ed.), *Belief Revision*, Cambridge University Press, Cambridge, pp. 29–51.
- Gärdenfors, Peter: 1990, 'The Dynamics of Belief Systems: Foundations vs. Coherence Theories', *Revue Internationale de Philosophie* **44**, 24–46.
- Grove, Adam: 1988, 'Two Modellings for Theory Change', *Journal of Philosophical Logic* **17**, 157–170.
- Hansson, Sven O., and Erik J. Olsson: 1999, 'Providing Foundations for Coherentism', *Erkenntnis* **51**, 243–265.
- Harman, Gilbert: 1973, *Thought*, Princeton University Press, Princeton.
- Harman, Gilbert: 1986, *Change in View*, Bradford Books, MIT Press, Cambridge, Mass.
- Lehrer, Keith: 1990, *Theory of Knowledge*, Routledge, London.
- Lehrer, Keith: 1997, *Self-Trust. A Study of Reason, Knowledge, and Autonomy*, Oxford University Press, Oxford.
- Lehrer, Keith: 2000, *Theory of Knowledge*, second, substantially revised edition, Westview Press, Boulder.
- Levi, Isaac: 1967, *Gambling With Truth*, Knopf, New York.
- Levi, Isaac: 1984, *Decisions and Revisions. Philosophical Essays on Knowledge and Value*, Cambridge University Press, Cambridge.
- McCarthy, Gerald D. (ed.): 1986, *The Ethics of Belief Debate*, Atlanta Scholars Press, Atlanta.
- Nayak, Abhaya C.: 1994, 'Foundational Belief Change', *Journal of Philosophical Logic* **23**, 495–533.
- Nebel, Bernhard: 1989, 'A Knowledge Level Analysis of Belief Revision', in Ronald Brachman, Hector Levesque and Ray Reiter (eds.), *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, San Mateo, Cal., pp. 301–311.
- Nozick, Robert: 1981, *Philosophical Explanations*, Belknap Press, Harvard University Press, Cambridge, Mass.
- Olsson, Erik J.: 1998, 'Competing for Acceptance: Lehrer's Rule and the Paradoxes of Justification', *Theoria* **64**, 34–54.
- Olsson, Erik J.: 1999, 'Cohering With', *Erkenntnis* **50**, 273–291.
- Plantinga, Alvin: 1990, 'Justification in the 20th Century', *Philosophy and Phenomenological Research* **50**, Supplement Vol., 45–71.
- Plato: 1994, *Meno*, in Jane M. Day (transl. and ed.), *Plato's Meno in Focus*, Routledge, London.
- Rescher, Nicholas: 1976, *Plausible Reasoning*, Van Gorcum, Assen, Amsterdam.
- Rott, Hans: 2000, 'Two Dogmas of Belief Revision', *Journal of Philosophy* **97**, 503–522.
- Rott, Hans: 2001, *Change, Choice and Inference*, Oxford University Press, Oxford.
- Shackle, G.L.S.: 1961, *Decision, Order and Time in Human Affairs*, Cambridge University Press, Cambridge.

- Sosa, Ernest: 1980, 'The Raft und the Pyramid: Coherence Versus Foundations in the Theory of Knowledge', *Midwest Studies in Philosophy* **5**, 3–25. Reprinted in E.S., *Knowledge in Perspective*, Cambridge University Press, Cambridge 1991, pp. 165–191.
- Spohn, Wolfgang: 1988, 'Ordinal Conditional Functions', in William L. Harper and Brian Skyrms (eds.), *Causation in Decision, Belief Change, and Statistics*, Vol. II, Reidel, Dordrecht, pp. 105–134.
- Weintraub, Ruth: 1990, 'Decision-Theoretic Epistemology', *Synthese* **83**, 159-177.