

INCONSISTENCY:
THE COHERENCE THEORIST'S NEMESIS?

Mylan ENGEL
Northern Illinois University, DeKalb

In his recent book *Theory of Knowledge*¹, Keith Lehrer defends a coherence theory of epistemic justification according to which a person is justified in accepting that *p* if and only if *p* coheres with the other things that that person accepts. Therein he also avers that "One ingredient in coherence is consistency. A belief logically inconsistent with others fails to cohere with them."^{2,3} This yields two unacceptable results. The first being that one can *never* be justified in accepting an inconsistent set of propositions. The second being that a person who already accepts an inconsistent set of propositions can *never* be justified in accepting *anything*, since every proposition is inconsistent with a set of inconsistent propositions. The latter result is especially worrisome since most (if not all) people do accept some inconsistent propositions. Thus, the coherence theorist faces the following dilemma. If consistency is necessary for coherence, then the aforementioned problems arise. If consistency is *not* necessary for coherence, then it seems that the coherence theorist would have to deny deductive closure in order to prevent other pairs of inconsistent propositions from readily cohering with an already inconsistent acceptance system, a move with counterintuitive consequences of its own.

In what is to follow, I shall argue that consistency is *not* necessary for coherence, at least not if coherence is necessary for justification.

1. Keith Lehrer, *Theory of Knowledge* (Boulder and San Francisco: Westview Press, 1990).

2. *Ibid.*, p. 94.

3. Other coherence theorists tend to concur. See, e.g., Laurence Bonjour's *The Structure of Empirical Knowledge* (Cambridge: Harvard University Press, 1985), and his "The Coherence Theory of Empirical Knowledge" in *Empirical Knowledge: Readings in Contemporary Epistemology*, ed. Paul K. Moser (Totowa, NJ: Rowman and Littlefield, 1986).

I shall also argue that, despite his assertion to the contrary, Lehrer's analysis of coherence does not rule out the possibility of coherently inconsistent acceptance systems. To set the stage for these arguments, section 1 provides an explication of the relevant portions of Lehrer's epistemology. Section 2 deals with minimal inconsistency and coherent acceptance. There we shall see that a person can indeed be justified in accepting a proposition inconsistent with the other things she accepts. We shall also see that Lehrer's theory yields the same result, his assertion to the contrary notwithstanding. The third and final section concerns the problems which arise when inconsistency already infects a person's acceptance system and explores the tenability of coherence without consistency as a source of justification.

1. Lehrer's Monster

Lehrer has dubbed his most recent theory of knowledge "The Monster Theory".⁴ The main tenets of the Monster Theory can be captured formally as follows:

- (MT) *S* knows that *p* iff:
- (i) *p* is true,
 - (ii) *S* accepts that *p*,
 - (iii) *S* is completely justified in accepting that *p*, and
 - (iv) *S*'s complete justification in accepting that *p* is not defeated by any false proposition.

Condition (i) is a standard truth condition where truth is cashed out in terms of correspondence. Condition (iv), while needed to handle the Gettier problem, is irrelevant to the present enterprise and can be safely ignored. This leaves conditions (ii) and (iii). They constitute the heart of the monster and will be the sole focus of the present section.

Lehrer replaces the traditional belief condition with an acceptance condition, to wit, condition (ii). He has two main reasons for

4. This baptism occurred in Lehrer's article "Knowledge Reconsidered", in *Knowledge and Skepticism*, ed. Marjorie Clay (Boulder: Westview Press, 1989).

doing so. First, acceptance is a more dispassionate and voluntaristic doxastic attitude than belief. We often have little control over what we believe. Sometimes we continue to believe a proposition in the face of strong counterevidence, because it is extremely comforting to so believe. In such a situation, we might well accept the falsity of a proposition we cannot help but believe. Other times we find a proposition so painful that we cannot bring ourselves to believe it, despite overwhelming evidence of its truth. Faced with such a situation, we might accept a proposition we cannot bear to believe. In both kinds of cases, the more reflective attitude of acceptance reflects our epistemic stance.

Lehrer's second reason for opting for acceptance over belief is more fundamental. Acceptance is a *purposive* doxastic attitude. When we accept a proposition, we do so for a purpose. Moreover, we might accept a proposition for one purpose which we would not accept for another. For example, an attorney might accept that her client is innocent for the pragmatic purpose of providing her client the best defense possible, even though she would not accept this for the epistemic purpose of gaining truth. Lehrer characterizes the uniquely epistemic kind of acceptance required by condition (ii) as follows:

It is accepting something for the purpose of attaining truth and avoiding error with respect to the very thing one accepts. More precisely, the purpose is to accept that *p* if and only if *p*.⁵

Thus, condition (ii) is to be understood as an abbreviation of the following condition:

(ii') *S* accepts that *p* for the purpose of attaining truth and avoiding error with respect to *p*.

Henceforth, the term 'acceptance' will only be used to refer to this epistemic kind of acceptance. Let us now turn our attention to condition (iii).

Condition (iii) asserts that knowledge requires complete justification, which Lehrer defines as follows:

5. Keith Lehrer, *Theory of Knowledge*, *op. cit.*, p. 11.

(CJ) *S* is completely justified in accepting that *p* iff *S* is personally and verifiably justified in accepting that *p*.⁶

To unpack this definition, let us start by examining personal justification. For Lehrer, personal justification is a form of subjective justification such that justification is exclusively a function of the other things one accepts. Accordingly, *S* is personally justified in accepting that *p* if and only if *p* coheres with *S*'s acceptance system, where *S*'s acceptance system is the set of statements of the form – *S* accepts that *p* – attributing to *S* just those things *S* accepts for the epistemic purpose. Lehrer's definition of coherence appeals to the notion of a competitor, where a competitor is to be understood as follows: *c* competes with *p* for *S* on the basis of *S*'s acceptance system if and only if it is less reasonable for *S* to accept that *p* on the assumption that *c* is true than on the assumption that *c* is false on the basis of *S*'s acceptance system. Lehrer's preliminary definition of coherence will suffice for our purposes. It states that *p* coheres with *S*'s acceptance system if and only if it is more reasonable for *S* to accept that *p* than to accept any competing claim *c* on the basis of *S*'s acceptance system. When *p* is more reasonable for *S* to accept than some competing claim *c* on the basis of *S*'s acceptance system, then *p* beats *c* for *S* on the basis of *S*'s acceptance system. From the foregoing definitions, we may derive the following formal definition of personal justification:

(PJ) *S* is personally justified in accepting that *p* iff it is more reasonable for *S* to accept that *p* than to accept any competing claim *c* on the basis of *S*'s acceptance system.⁷

6. *Ibid.*, p. 135. The required temporal indices which Lehrer includes have been suppressed for the sake of simplicity and will remain suppressed throughout this essay.

7. (PJ) is actually an oversimplification of Lehrer's account of personal justification because it does not allow for the possibility of neutralizing competitors. Occasionally, a competing claim *c* will be at least as reasonable as *p* on the basis of *S*'s acceptance system, yet it will compete with *p* only indirectly. In such a case, some third statement *n* may be able to neutralize *c* as a competitor. This happens when the conjunction *n-c* is as reasonable as *c* alone on the basis of *S*'s acceptance system and *n-c* does not compete with *p*. I have omitted this wrinkle of Lehrer's theory of personal justification because it is irrelevant in the present context.

An example will help motivate the account. Suppose I am in the produce section of a local supermarket standing before an apple display with my eyes open. Consider the following statements.

P: I see an apple.

C: I am looking at a stand filled with wax ersatz apples.

C clearly competes with P because it is less reasonable for me to believe that I see an apple on the assumption that I am looking at a stand filled with wax apples than on the assumption that I am not looking at a stand filled with wax apples. However, on the basis of the other things I accept, it is more reasonable for me to accept P than to accept C. After all, I accept that supermarkets are in the business of selling food, not in the business of deceiving customers. I also accept that supermarkets rarely, if ever, fill their stands with ersatz produce. I may be mistaken in accepting these things, but since I do accept them, P is more reasonable for me to accept than C. If every other competitor is less reasonable to accept than P on the basis of my acceptance system, then P coheres with my acceptance system and I am personally justified in accepting that P. Of course, to be completely justified in accepting that P, I would need to be verifiably justified in accepting that P, as well. Let us, therefore, examine verifiably justification.

Verifiably justification results from coherence with the verifiably system. The verifiably system is the subsystem of the acceptance system which remains after all error has been deleted from said acceptance system. Thus, *S*'s verifiably system is the set of statements of the form – *S* accepts that *p* – attributing to *S* just those *true* things that *S* accepts for the epistemic purpose. Cashing out coherence with the verifiably system in a way analogous to coherence with the acceptance system, we can formally define verifiably justification as follows:

(VJ) *S* is verifiably justified in accepting that *p* iff it is more reasonable to accept that *p* than to accept any competing claim *c* on the basis of *S*'s verifiably system.⁸

8. As with (PJ), (VJ) is an oversimplification because (VJ) fails to take into account the possibility of neutralization. Such an oversimplification is warranted because neutralization has no bearing on the topic of discussion.

Since personal and verifical justification have rather different roles to play in Lehrer's account of complete justification, I shall conclude this section with a brief discussion of their respective roles. To understand the role of verifical justification, we should start by noting that we lack cognitive access to our verifical system.⁹ Lacking access to our verifical system, we are unable to determine when we are verifically justified in accepting a proposition. Since we lack access to our verifical justificatory status, such justification cannot guide us in our epistemic reflections. What, then, is the purpose of verifical justification? Its purpose is to provide a conceptual connection between complete justification and truth. Without such a truth connection, complete justification would lack any underlying rationale. Why, after all, should we be interested in completely justified acceptance, if not for the fact that such acceptance is likely to be true?¹⁰ Thus, the need for a truth connection is clear. Since personal justification alone is unable to provide a truth connection, it must be supplemented with some kind of truth-connected justification. Verifical justification provides the needed supplementation, although it does so in an externalist way which some may find unsatisfying.

It is appropriate that we conclude this section with a discussion of personal justification, since it is the engine which drives complete justification. Personal justification is a normative notion in the sense of being acceptance-guiding. It tells us what we should and should not accept. In short, we ought to accept that p when and only when we are personally justified in accepting that p . Lehrer acknowledges the normative aspect of personal justification,¹¹ and he admits that personal justification must be our sole guide in reasoning and acceptance, for he tells us:

9. It is easy to see that we lack such access to our verifical system. After all, were we aware of any errors in our acceptance system, we would for the purpose of attaining true and avoiding error, stop accepting these erroneous propositions, and so, they would not belong to our acceptance system. Thus, any errors in our acceptance system must be unknown to us, and so, we are unable to determine the membership of our verifical system.

10. Laurence Bonjour makes a similar point regarding epistemic justification in general in his "Can Empirical Knowledge Have a Foundation?", *American Philosophical Quarterly*, Vol. 15, No. 1 (January, 1978), p. 5.

11. Keith Lehrer, *Theory of Knowledge*, *op. cit.*, p. 127.

When I consider whether I am completely justified in accepting something now, I can only determine whether I am personally justified and, if I am, conclude that I am completely justified.¹²

The relevance of these remarks concerning the normativity of personal justification is straightforward. They allow us to ignore verifical justification in the present context, because if a person is personally justified in accepting a proposition inconsistent with her acceptance system, then she ought to accept that proposition. If such a proposition exists, then consistency is not necessary for coherence. The burden of the next section is to demonstrate that such a proposition exists.

2. Coherence, Minimal Inconsistency, and Reasonable Acceptance

This section explores the relationship between coherence and minimal inconsistency. We shall see that a minimally inconsistent set can be coherent, personally justified, and extremely reasonable to accept. Let us start by defining minimal inconsistency. A minimally inconsistent set is an inconsistent set such that all of the members of the set are required to make it inconsistent. Following Lehrer, we may define minimal inconsistency more precisely as follows:

(MI) A set of statements is a minimally inconsistent set iff the set of statements is logically inconsistent and such that every proper subset of the set is logically consistent.¹³

The argument to show that a person can be personally justified in accepting a minimally inconsistent set parallels the famous "Paradox of the Preface". Said paradox arises when we imagine an historian who writes the following disclaimer in the preface of her rather lengthy book: "I have made every effort to ensure that each

12. *Ibid.*, p. 151.

13. Keith Lehrer, "Induction, Rational Acceptance, and Minimally Inconsistent Sets", in *Metamind* (Oxford: Clarendon Press, 1990), p. 99.

historical claim in this book is true and accurate, and I accept of each claim that it is true. Of course, being fallible, I also accept that at least one of these claims is erroneous." This modest historian accepts a minimally inconsistent set. Is she irrational for doing so? Before answering this question, we should note that non-historians face essentially the same problem. Consider a person *S* who accepts the following consistent set of contingent propositions: $\{p_1, p_2, \dots, p_n\}$ where *n* is very large.¹⁴ If *S* is reasonably modest, she too may accept that at least one of the contingent things she accepts is false. To avoid a self-referential truth paradox, this latter acceptance should be construed as the acceptance of the disjunction of the denials of the other things she accepts, to wit,

$$(IMP) \sim p_1 \vee \sim p_2 \vee \dots \vee \sim p_n.$$

IMP is an Inconsistency-Making Proposition for *S*, because if *S* accepts IMP, then the set of things she accepts thereby becomes minimally inconsistent. Should *S* accept IMP? Kyburg and Campbell have both argued that accepting IMP is extremely reasonable for *S* and that, consequently, *S* should accept IMP.¹⁵

Lehrer disagrees. In "Reason and Consistency" he argues that accepting IMP is unreasonable for a person with the ideal of maxiverificity, i.e., the ideal of accepting that *p* if and only if *p*, because accepting IMP is a defeatist policy which guarantees failure.¹⁶ He also contends that the expected value of accepting IMP is negative, since IMP contains so little information, and that since it is unreasonable to accept a proposition with negative expected value, it is unreasonable to accept IMP.¹⁷ I shall examine each of these

14. The set of propositions under consideration is restricted to contingent propositions so as to prevent necessary truths from being unbeatable competitors of IMP, the disjunction of the denials of propositions in the set $\{p_1, p_2, \dots, p_n\}$, for each member of this set competes with IMP as is demonstrated in note 19.

15. Henry Kyburg and Richard Campbell have offered similar arguments to demonstrate the reasonableness of accepting IMP. See Kyburg's "Conjunctivitis" in *Induction, Acceptance, and Rational Belief*, ed. M. Swain; and Campbell's "Can Inconsistency Be Reasonable?", *Canadian Journal of Philosophy*, Vol. XI, No. 2 (June, 1981).

16. Keith Lehrer, "Reason and Consistency" in *Metamind*, *op. cit.*, pp. 159-160.

17. *Ibid.*, p. 160-164.

first it will be useful to calculate the probability of IMP.

The probability of IMP for *S* is a function of the respective probabilities of the other propositions *S* accepts. Consequently, the probability of IMP will vary from cognizer to cognizer and can only be calculated for concrete cases. Let us, therefore, calculate the probability of IMP for an extremely charitable case. Suppose that Sam accepts 10,000 independent contingent propositions, each with a .999 probability of being true.¹⁸ We may calculate the probability of IMP for Sam as follows:

$$\text{prob}(IMP) = 1 - \text{prob}(\sim IMP),$$

where ($\sim IMP$) is the following conjunction:

$$(\sim IMP) p_1 \cdot p_2 \cdot \dots \cdot p_{10,000}$$

The probability of $\sim IMP$ is:

$$\begin{aligned} \text{prob}(\sim IMP) &= \text{prob}(p_1) \times \text{prob}(p_2) \times \dots \times \text{prob}(p_{10,000}) \\ \text{prob}(\sim IMP) &= .999^{10,000} = .0000452 \end{aligned}$$

Thus,

$$\text{prob}(IMP) = 1 - .0000452 = .9999548.$$

Since the probability of IMP for Sam is .9999548, IMP is more probable for Sam than any other contingent proposition he accepts. Is Sam personally justified in accepting that IMP? Intuitively, he is. Moreover, as we shall now see, Lehrer's account of personal justification yields the same result.

Recall that *S* is personally justified in accepting that *p* just in case it is more reasonable for *S* to accept that *p* than to accept any competing claim *c* on the basis of *S*'s acceptance system. One claim

18. The independence requirement is added to allow us to make the needed probability calculations. Such a requirement is not, however, as implausible as it might at first seem. A person might, for example, accept 10,000 propositions each of which predicts the winning ticket in a different lottery. These propositions are clearly independent, and people accept such propositions all the time.

which obviously competes with IMP for Sam is \sim IMP. Since it is much more reasonable for Sam to accept IMP than to accept \sim IMP on the basis of Sam's acceptance system, IMP clearly beats \sim IMP for S. Notice, however, that each individual member p_i , belonging to the set of contingent propositions p_1 through $p_{10,000}$ which Sam accepts, also competes with IMP, because for any such p_i , $\text{prob}(\text{IMP}/p_i) < \text{prob}(\text{IMP})$.¹⁹ But since the probability of IMP (.9999548) is much greater than the probability of any such p_i (.999 by hypothesis), IMP beats every competing claim p_i on the basis of Sam's acceptance system, and so, Sam is personally justified in accepting IMP.

We are even more justified in accepting IMP than Sam is. After all, the probability of IMP for each of us is immensely greater than it is for Sam, for we accept vastly more independent contingent propositions than Sam, many of which are a great deal less probable than .999. To get some idea of just how probable IMP is for us, let us suppose, again charitably unrealistically, that we accept only 100,000 independent contingent propositions, each with a probability of .999. Under such a supposition, the probability of \sim IMP for us would be 3.5385×10^{-44} and the probability of IMP for us would be expressed by a decimal point followed by 43 nines and change. In short, IMP is virtually certain for each of us, especially since even this latter probability assessment underestimates the probability of IMP for each of us. In order for us *not* to be personally justified in accepting IMP, there would have to be some competing claim c which is more reasonable for us to accept than IMP on the basis of our acceptance systems. It is, of course, extremely implausible to think that any such competing claim exists, for the *strongest* competitors of IMP for us are the individual contingent propositions belonging to our respective acceptance systems and such propositions are, presumably, much less probable than IMP is for us. Consequently, even on Lehrer's account of personal justification, we turn out to be personally justified in accepting IMP. Even so, Lehrer offers two arguments against the reasonableness of accepting

19. Each such p_i is negatively relevant to and, hence, competes with IMP, because $\text{prob}(\text{IMP}/p_i) = 1 - .999^{9999}$, whereas $\text{prob}(\text{IMP}) = 1 - .999^{10000}$ and, since $.999^{10000}$ is smaller than $.999^{9999}$, $\text{prob}(\text{IMP}/p_i) < \text{prob}(\text{IMP})$.

IMP. Let us examine these arguments more closely, starting with his argument designed to show that the expected value of accepting IMP is negative.

Lehrer would object to the foregoing reasonability assessments because they treat reasonableness as *exclusively* a function of probability. He contends that a less probable but more informative statement can be more reasonable to accept than a more probable but less informative one,²⁰ and so, he asserts:

...the utility of accepting h when h is true depends on how much h tells us, on how informative h is. Therefore, the reasonableness or expected utility of accepting h is a function of the informativeness of h as well as of the probability of h .²¹

So, in order to calculate the reasonableness of accepting a proposition p , we must appeal to the following formula:

$$r(p) = \text{prob}(p)\text{Ut}(p) + \text{prob}(\sim p)\text{Uf}(p),$$

where ' $r(p)$ ' represents the reasonableness of accepting p , ' $\text{Ut}(p)$ ' represents the utility of accepting p when p is true, and ' $\text{Uf}(p)$ ' represents the disutility of accepting p when p is false. Moreover, Lehrer maintains that IMP "is so uninformative that, even if true, that truth is of negligible interest,²² and so, the positive expectation of accepting IMP when IMP true is minimal, whereas the negative expectation of accepting IMP when IMP is false is rather large. Consequently, in Lehrer's estimation, the expected utility of accepting IMP is negative. Obviously the expected utility of accepting \sim IMP is also negative, and as a result, Lehrer contends that the most reasonable epistemic stance to take with respect to IMP is that of *withholding* IMP (i.e. neither accepting IMP nor its negation), because the expected utility of withholding any proposition is zero.²³ To evaluate this argument we need to calculate the expected utility of accepting IMP.

20. Keith Lehrer, "Reason and Consistency", *op. cit.*, pp. 159-160.

21. Keith Lehrer, *Theory of Knowledge*, *op. cit.*, p. 130.

22. Keith Lehrer, "Reason and Consistency", *op. cit.*, pp. 160.

23. *Ibid.*, p. 164.

It is difficult to assign precise quantitative values to $Ut(IMP)$ and $Uf(IMP)$, but following Lehrer's qualitative assessments of their values, let us assume that for Sam $Ut(IMP) = .01$ and $Uf(IMP) = -.99$ and calculate the reasonableness of accepting IMP for Sam. Substituting the respective values for Sam in the reasonability equation above, we get:

$$r(IMP) = (.9999548 \times .01) + (.0000452 \times -.99) = .0099548$$

Hence, accepting IMP has positive expected value for Sam. Even if we assume that for Sam $Ut(IMP) = .001$ and $Uf(IMP) = -.999$, accepting IMP still has a positive expected value of .0009458 for Sam. Similar calculations would show that accepting IMP has positive expected value for each of us, as well. Since the expected value of accepting $\sim IMP$ is clearly negative and the expected value of withholding IMP is zero, the only alternative with positive expected utility is that of accepting IMP, and therefore, we ought to accept IMP.

Furthermore, there is reason to think that the expected value of accepting IMP is much greater than our calculations indicate. Our calculations were predicated on the assumption that IMP is negligibly informative. However, such an assumption seems unwarranted. When discussing the reasonableness of highly informative, but not very probable claims, Lehrer asserts,

The reasonableness of accepting such claims is influenced by our interest in accepting claims which, if true, are *important general truths about ourselves* and our universe²⁴ (emphasis added).

But notice, IMP does tell us something important about ourselves, if only indirectly. It tells us that at least some of what we accept is false, which *ipso facto* tells us that we are fallible in what we accept, and surely, this latter truth constitutes an important general truth about ourselves. Thus, IMP is more informative than it might at first appear, and so, $Ut(IMP)$ is bound to be greater than .01, which *ceteris paribus* entails that accepting IMP has a positive expected

24. Keith Lehrer, *Theory of Knowledge*, *op. cit.*, p. 128.

utility even greater than that indicated by our calculations. Since the expected utility of accepting IMP is positive rather than negative, Lehrer's first argument is unsuccessful. Let us now consider his second argument.

Lehrer contends that insofar as we are *epistemic* inquirers, we must have as our goal the ideal of maxiverificity, i.e. the goal of accepting a proposition if and only if it is true. This contention is implicit in his account of epistemic acceptance, to wit, acceptance aimed at accepting that p if and only if p . A person striving to accept that p if and only if p is *ipso facto* striving for maxiverificity. Such a person seeks to accept all and only truths. As noted earlier, Lehrer maintains that it is unreasonable for a person seeking to accept all and only truths to accept IMP, because accepting IMP guarantees accepting at least one falsehood. After all, if S accepts IMP, then if IMP is true, at least one of the other contingent propositions S accepts will be false and if all of the other contingent propositions S accepts are true, then IMP will be false. Either way, accepting IMP guarantees failure with respect to the goal of accepting *only* truths. Given that accepting IMP guarantees failure, Lehrer contends that it is perfectly reasonable to refuse to accept it, because:

A quite sensible person, however, might adopt maxiverificity as an ideal and refuse to follow a policy that guarantees failure. Accepting S2 [our IMP] is a defeatist policy with respect to the attainment of this ideal and may be eschewed for that reason.²⁵

There are two problems with this response, one which merely undermines it, another which refutes it. I begin with the former. The ideal of maxiverificity involves two goals, the goal of accepting *all* truths and the goal of accepting *only* truths, and these goals pull in opposite directions. In responding as he does, Lehrer is placing greater emphasis on the latter goal. But notice, not accepting a truth guarantees failure with respect to the maxiverific ideal just as surely as accepting a falsehood does. Since IMP is virtually certain for each of us, our interest in accepting *all* truths dictates that we accept IMP. Thus, a person who weights both of the maxiverific goals

25. Keith Lehrer, "Reason and Consistency", *op. cit.*, pp. 160.

equally can get no guidance as to whether or not to accept IMP from the maxiverific ideal *per se*, since this ideal provides both a reason to accept IMP and a reason not to. However, if we must weight one of the maxiverific goals more heavily than the other, then it seems to me that we ought to weight the positive goal of attaining truth more heavily than the negative goal of avoiding error, since, unlike error avoidance, the truth may set us free.

The second objection is decisive. Recall Lehrer's characterization of the epistemic purpose:

It is accepting something for the purpose of attaining truth and avoiding error *with respect to the very thing one accepts*²⁶ (my emphasis).

Consequently, the fact that accepting IMP guarantees failure with respect to the maxiverific ideal is completely irrelevant to the question of whether or not we should accept IMP. In light of the epistemic purpose as described above, the only relevant considerations are those considerations which have a direct bearing on IMP's truth value. If we have every reason to think that IMP is true and no reason to think that IMP is false, then we ought to accept IMP for the purpose of attaining truth and avoiding error *with respect to IMP*. Given the virtual certainty of IMP for us, we do have every reason to think that IMP is true, and so, by Lehrer's own account of acceptance, we ought to accept IMP.

We may summarize the findings of the present section as follows: First, IMP is extremely probable for us, and as a result, it is extremely reasonable for us to accept IMP, much more so than withholding IMP or accepting \sim IMP. Thus, intuitively, we are certainly justified in accepting IMP. Second, since IMP is more reasonable for us to accept than any of its competitors on the basis of our respective acceptance systems, IMP coheres with our respective acceptance systems, and so, according to Lehrer's theory, we are personally justified in accepting IMP. In this regard, Lehrer's theory yields the intuitively correct result. And finally, since IMP both coheres with and is minimally inconsistent with our respective

26. Keith Lehrer, *Theory of Knowledge*, *op. cit.*, p. 11.

acceptance systems, consistency is not a necessary condition for coherence (at least, not as Lehrer defines "coherence"). It remains to be seen whether coherently inconsistent acceptance systems are adequate to justify other acceptances.

3. Coherence, Minimal Inconsistency, and Deductive Closure

Can a coherence theory, such as Lehrer's, which sanctions the acceptance of minimally inconsistent sets of statements be sustained? To answer this question, let us examine three objections which might be raised concerning such minimally inconsistent acceptance systems.

Objection 1: Inconsistency has traditionally been frowned upon because everything follows from an inconsistent set of premises. This observation, when combined with the widely held assumption that rational acceptance is closed under deduction, leads to the following objection: A person who accepts a minimally inconsistent set S_1 could deduce and, hence, *rationally accept* any proposition, even a logically false proposition, on the basis of an argument which has as its premises every member of S_1 .

Such an objection is, however, fundamentally misguided, for while it is true that everything follows *validly* from an inconsistent set of premises, nothing follows *soundly* from such a set of premises, and only apparently sound arguments should be our guide. In short, rationality requires that we accept the conclusion of a valid argument when and only when we have reason to believe that all of its premises are true. Any argument with inconsistent premises *ipso facto* has at least one false premise and, consequently, provides no rational impetus for accepting its conclusion. Accordingly, Lehrer can thwart objection 1 by restricting deductive closure as follows:

(RDC₁) It is rational for a person S to infer that p on the basis of a minimally inconsistent set S_1 iff S can validly deduce p from a consistent subset of S_1 .

Objection 2: Even with the (RDC₁) restriction on deductive closure, it is still possible to validly deduce inconsistent propositions from a minimally inconsistent acceptance system. To see why, let

us suppose that Sam accepts the following minimally inconsistent set of propositions: $\{p_1, p_2, \dots, p_{10,000} \& \text{IMP}\}$. Call this set " S_1 ". Sam can validly deduce $\sim\text{IMP}$ from S_1 , because using the rule of addition he can validly infer $p_1 \cdot p_2 \cdot \dots \cdot p_{10,000}$ (i.e. $\sim\text{IMP}$) from an argument whose individual premises are p_1, p_2, \dots , and $p_{10,000}$. Since the subset $\{p_1, p_2, \dots, p_{10,000}\}$ is consistent and since $\sim\text{IMP}$ follows validly by addition from this subset, (RDC_1) entails that it is rational for Sam to accept $\sim\text{IMP}$. Thus, Lehrer's coherence theory faces the following problem: According to said theory, it is rational for Sam to accept IMP because IMP coheres with Sam's acceptance system, and it is also rational for Sam to accept $\sim\text{IMP}$, despite its extreme improbability, because $\sim\text{IMP}$ is derivable from a consistent subset of Sam's acceptance system. Intuitively, any theory of rational acceptance which entails that it is rational to accept both p and $\sim p$ is *ipso facto* inadequate. One way for Lehrer to circumvent this objection is by placing an additional, probability restriction on deductive closure as follows:

(RDC₂) It is rational for a person S to infer that p on the basis of a minimally inconsistent set S_1 iff (i) S can validly deduce p from a consistent subset of S_1 and (ii) p is no less probable than n , where n specifies some fixed probability below which no rationally acceptable proposition can fall.

Of course, it is notoriously difficult to specify n , because many extremely improbable scientific propositions, e.g. propositions concerning the nature of black holes, are nevertheless rationally acceptable. Thus, in the end (RDC₂) may prove chimerical as a solution to the present objection.

Objection 3: Despite the apparent incongruity with his other views, Lehrer has rightly observed that:

A person who has discovered that he accepts an inconsistent set of statements may *reasonably* continue to do so when, for lack of time or insight, he sees no satisfactory way of extricating himself²⁷ (emphasis added).

27. Keith Lehrer, "Reason and Consistency", *op. cit.*, pp. 160.

An example will illustrate the point. Phil, a philosophy graduate student in a seminar on free will, has come to realize that he accepts the following minimally inconsistent propositions:

- (D) Determinism is true
- (F) Human beings are free agents.
- (I) Freedom and determinism are incompatible.

He accepts that D, because he accepts that he has never witnessed an uncaused event and does not believe there are any. He accepts that F, because he accepts that he and his fellow human beings could have done otherwise. And he accepts that I, because he thinks that all compatibilistic arguments equivocate with respect to "freedom" and because he cannot imagine how a causally determined agent could do otherwise. An astute metaphysician might be able to show Phil the error of his ways, but he has not met such a metaphysician, and nothing he has read so far has been able to show him the way out of his metaphysical fly bottle. The crucial thing to note is that F, D, and I already belong to Phil's acceptance system, and it would be irrational for him to randomly reject one of these propositions, because, given his predicament, he might very well end up rejecting one of the true propositions. Hence, as Lehrer admits, it is rational for Phil to continue to accept F, D, and I, which is equivalent to saying that Phil is personally justified in continuing to accept F, D, and I. Things get worse for Lehrer's coherence theory once we suppose that Phil also accepts the following propositions:

- (R) An agent is morally responsible for an action if and only if she performed that action freely.
- (K) Lee Harvey Oswald killed John F. Kennedy.

From consistent subsets of his acceptance system, Phil can validly deduce explicitly contradictory propositions as follows: From F, R, and K, he can validly infer that Oswald was morally responsible for killing Kennedy, and from D, I, R, and K, he can validly deduce that Oswald was not morally responsible for killing Kennedy. Since, given his predicament, it is rational for him to continue to accept D, F, I, R, and K, if rational acceptance is closed under deduction,

then it is rational for Phil to accept the two contradictory propositions just mentioned. Moreover, neither (RDC₁) nor (RDC₂) is of any use in blocking this objection. Thus, in order to avoid this objection, Lehrer would have to deny deductive closure outright. Unfortunately, Lehrer has argued elsewhere that:

...the set of accepted statements must be deductively closed. The reason for this is that it would be irrational for a person to refuse to accept the deductive consequences of the statements she accepts.²⁸

Let us conclude by returning to our earlier question: Can a coherence theory, such as Lehrer's, which sanctions the acceptance of minimally inconsistent sets be sustained? The answer is that such a coherence theory can only be sustained by denying that rational acceptance is closed under deduction, a move Lehrer seems unwilling to make. Of course, to avoid denying deductive closure, he might insist that consistency is necessary for coherence and argue that his coherence theory does not sanction the acceptance of minimally inconsistent sets. Unfortunately, such a move is problematic in two ways. First, if Lehrer makes consistency necessary for coherence, then he will fall prey to the first horn of the dilemma with which we began. Second, since intuitively we are indeed justified in accepting IMP, if Lehrer's coherence theory entails that we are not so-justified, it will yield the wrong result. In short, Lehrer must either hold that consistency is necessary for coherence and personal justification or deny that rational acceptance is closed under deduction, neither of which seems acceptable.

28. Keith Lehrer, "Induction, Rational Acceptance, and Minimally Inconsistent Sets", *op. cit.*, pp. 113.