

Critical Study

Errors of Reasoning, Naturalizing the Logic of Inference

John Woods

London: College Publications, 2013. Pp. xvii, 1-572. Paperback ISBN 978-1-84890-114-8.

Amazon.ca CDN\$32.54, Amazon UK £16, Amazon.com US\$25

Critical Study by **JAMES B. FREEMAN**

Department of Philosophy
Hunter College of The City University of New York
695 Park Avenue
New York, NY 10065
U.S.A.
jfreeman@HUNTER.CUNY.EDU

Keywords: causal response description of knowing, command and control model of knowing, concept-list misalignment thesis, EAUI-error, attractive, universal, incorrigible, externalism, internalism, traditional concept of fallacy, traditional list of fallacies

John Woods's stated aim in this monograph is to investigate premise-conclusion reasoning "with a special emphasis of factors in play when it falls into error, when bad reasoning is mistaken for good" (1).¹ Woods further states that he believes that logic has not done well in investigating these topics. To do better, logic needs some reconstruction to become "empirically sensitive," taking account of empirical information, and also "epistemologically aware" taking account of the "cognitive natures" of "real life" reasoners (2). Such a study will involve giving higher priority to studying errors of reasoning. It should not be content to understand error solely as failure of deductive validity or inductive strength. Woods indicates that much of the book will be devoted to answering what he phrases as Ham-

¹ Citations from Woods will simply present the page number or numbers within parentheses.

blin's Question: Why is research into fallacies in such a sorry state? Can one possible explanation be that the items in the traditional list are not really fallacies and are actually "virtuous ways if reasoning"? (7) Woods has given us a big book (over 500 pages) and a prolix one, looking at many side issues along the way. My purpose in this critical study is to present the core of Woods's position and, as they arise, to examine certain points critically.

Woods sees a significant reason why logic is in a sorry state with respect to fallacy theory as acceptance of the deductive validity/inductive probability standards for determining connection adequacy. As with a number of other authors, for example Walton (1992), L. J. Cohen (1977), Rescher (1976), and Pollock (1970), he holds that empirical observation of reasoning reveals that humans reason in a third way. Woods contends the proper study is the province of an empirically sensitive naturalized logic. Unlike formal logic, naturalized logic will include an account of reasoners, beings (including artificial devices) that do things (15). Premise-conclusion reasoning of such reasoners involves logical consequence. To a given type of premise-conclusion reasoning, the question arises of a corresponding consequence relation: are there such relations specific to that type of reasoning and what are their characteristic properties? This, says Woods, is the logician's central task.

We may distinguish between consequence having and consequence drawing—something going on in a reasoner's mind (24). Woods claims that it is part of logic's role to indicate when it is proper to draw a consequence, which is conditional upon "one's interests and one's resources" (26). Besides the errors of drawing a consequence that is not a consequence, or drawing a consequence which somehow goes against one's interests, there is the error of not seeing that a proposition is a consequence (26). Given a statement one wants to defend as a conclusion, there is the converse problem of finding a set of premises "properly supported" (26) which has that statement as a consequence. Woods claims that both issues "are legitimate components of a duly psychologized logic of erroneous reasoning" (28).

What conditions, then, will make for a good or adequate theory of reasoning? What are the evaluative consequences of specifying these conditions? If human reasoning falls short of the requirements of this theory of good reasoning, is it irrational? The *irrationality thesis* regards not living up to the ideal requirements of a theory as irrational. The *approximate rationality thesis* grants that although human reasoning does not conform to all the rules of logic, it does instantiate heuristics that result in thinking which approximates what the theory ideally requires.

Both theses put the blame for the disconnect on the human side. Woods asserts, albeit with hedging, that “the irrationality thesis is dismissible out of hand” (44). The approximate rationality thesis is more plausible. Humans engage in many activities, which they perform more or less well. Why should reasoning be excluded? There are heuristic procedures that, if followed, yield approximately correct results, results which may be the best we can ask of human reasoners. But, Woods asks, why should one regard such reasoning as *inferior*? Where is the argument that this reasoning *falls short* of that required by the theoretic ideal? Woods dismisses the view that the theoretical ideal is better than practical heuristic simply because of its “mathematical virtuosity” (38). Why, he asks, can we not hold that heuristics are the right standard and a theory’s ideal requirements are a deviation? Why does ideality confer normativity? To counter the viewpoint that it does, Woods proposes the “convergence of the normal and the normative, NN”: barring reasons to the contrary, “how we *do* reason from premises to conclusion is how we *should* reason” (52). Why should we not go with the conclusion of someone experienced in a subject matter, even if that person’s overall reasoning falls short of some formal ideal?

Woods holds that normatively good reasoning is “accurate” and “apt.” By accurate, Woods means that the reasoning is rightly done and produces the right answer. By apt, he means that it is reasonable to reason that way, given the circumstances (54). The NN convergence thesis can now be clarified this way: *Ceteris paribus*, typically how one reasons in real life is accurate, apt, or both. An instance of human reasoning is to be presumed correct until shown not to be—innocent until proven guilty. An inquiry into an empirical theory of error in reasoning can be guided properly by this presumption.

Woods’s comments call for critical reflection. First, if there is an error at some point in someone’s reasoning, does that show the person irrational or the reasoning irrational? One’s answer in part depends on one’s theory of rationality of which there is a plurality. Woods completely agrees with this point. Is the appraisal of only *approximate* rationality fair? One can further press the question of what standards are being required of the reasoning by approximate rationality. This point is totally conventional. If premises gave inductively strong support to a conclusion but did not deductively imply it, is the reasoning irrational? Different reasoning calls for different standards. If the rationality thesis, and even the approximate rationality thesis, is surreptitiously imposing the wrong standards of reasoning, the thesis itself is irrational! Woods holds that there are third way standards beyond deduction and induction, and identifying them

constitutes “no more important and widely open problem for the logic of error” (67).

I find a far more serious problem with the NN thesis itself. There are certain ways of forming basic (uninferred) beliefs which some epistemologists have standardly regarded as presumptively reliable—perception, introspection, memory, reason employed to recognize elementary analytic conceptual relations. Unless there is evidence that these mechanisms are not functioning properly or that the environment is somehow deceptive, there is a presumption for the resulting beliefs. But a belief arrived at through inference is not a basic but an inferred belief. Is it true that any mode of reasoning from premise to conclusion is presumptively reliable? I think not. As L. J. Cohen points out in (1992, 131), in some cases of inference, there is a place for voluntary intervention between the forming of a belief with a particular perceptual content and the acceptance of that belief as a premise for what we might further believe or do. One can withhold acceptance in the absence of sufficient evidence. But on the other hand, there might be no intervention. Is there a presumption that there is voluntary intervention in general? Consider an inference to explanation. If evidence suggests a certain explanation, even if one has a strong inclination to believe it, is there a presumption for that belief without further testing? Is there a presumption that in forming an explanatory belief, one’s immediately entertained belief is the best explanation? To be fair to Woods, he qualifies the NN convergence thesis with “at a first pass” (52). But unless Woods can show that the extension of this qualification is not sufficiently widespread to rebut the NN convergence thesis, he has not shown his point, that we should grant presumption to the NN convergence thesis.

Indeed, some of Woods’s later remarks further reinforce my point, e.g., when he discusses data-bending. We do not have data that a theory can simply proceed to explain or to take as confirmatory evidence. Rather, candidate data need “pretheoretical massaging” (71). Suppose, Woods proposes, that one is constructing a theory of a type of behavior, K, open to normative assessment. To be adequate, such a theory must take account of actual instances of K-behavior and to “adequately preserve” what is pretheoretically known of K-behavior” (71). The massaging of data may involve some reconceptualization of what is an instance of K-behavior. But such reconceptualization opens the door to error. Data can be mismassaged, up to the point of just making data up. The framework of an inquiry K determines how the data of K are conceptualized. Further, “Method influences the conceptualization of data. How the data are conceptualized influences, in turn, what counts as theoretically adequate responses to them” (75). In light of these admis-

sions, how does the presumption for the thesis of NN convergence fare?

Woods turns in Chapters 3 and 4 to the issue of knowing. He presents theses that will be important for much of the subsequent discussion. He takes certain claims about knowledge to be givens.

The cognitive abundance thesis: Human beings have knowledge, lots of it (86).

The error abundance thesis: Human beings make errors, lots of them (86).

This assertion also allows for two further assertions:

The enough already thesis: Human beings are right enough about enough of the right things enough of the time to survive and prosper (88).

The knowing-well-being proportionality thesis: To a significant degree the scale of human flourishing matches the scale of cognitive well-doing (89).

Further observation of human cognitive beings supports:

The cognidiversity thesis: Cognidiversity is a fit and necessary adjustment to nature's diversity (90). That is, humans not only know a lot, there are lots of things to know about.

Further reflection on or observation of human knowing shows that human knowledge is a causal response to the environment:

A causal response description of knowing (CR-characterization): A subject knows that α provided that α is true, he believes that α , his belief was produced by belief-forming devices in good working order and functioning herein the way they are meant to, operating on good information and in the absence of environmental distortion or interference (93).

I cannot fail to note the resemblance of this characterization to Plantinga's characterization of knowledge in (1993): Knowledge is warranted true belief, where a belief has warrant just in case it is produced through the operation of a properly-functioning belief-generating mechanism, operating in a non-distorting environment for which it was designed to operate, oriented to the truth (in contrast to other aims such as reassurance), with an objectively high probability of reliability.

Woods indicates another characterization of knowledge, what he calls the command and control model (CC):

A command and control model of knowing: Knowledge is a case making achievement, in which knowing that α depends on the knower's constructing a successful argument for it, or at least having the argument ready to hand and within his timely reach (97).

This is a version of internalism. As Woods points out, the CC conception is very much in line with the justified true belief definition of knowledge. As Plantinga points out in (1993), it is connected with a deontological view of knowledge. To be justified in holding a belief, one must have satisfied one's epistemic duties. What does this mean?

We have explored these issues in some detail in Chapter Four of (2005). I want to recall certain points from this discussion, since they are very relevant to points Woods seeks to make here. The classical foundationalists Descartes and Locke held a deontological or internalist view of justification. It is possible that one may conscientiously believe that doing X is his duty, when in fact it is not. His duty is to perform some incompatible action Y. He follows his conscience and does X. Has he acted wrongly? Although objectively wrong, since his duty lay in doing Y and not X, it is not subjectively wrong, since he sincerely believed that doing X was required. We may thus distinguish between subjective and objective duty. Descartes and Locke both believed that one could identify a class of cases where one could simply "see" one's objective epistemic duty, where one's objective and subjective duty coincide. For Descartes, we can just recognize when certain ideas are sufficiently clear and distinct. For Locke, certain propositions are immediately evident. The question of epistemic duty does not arise here, since our acceptance is involuntary. We may also recognize immediately that these propositions support further propositions. Our objective duty is to accept only those propositions properly supported by the body of immediately supported propositions. But since we can recognize immediately evident propositions and the support structure between propositions, we can recognize that we are within our epistemic rights to accept the supported propositions. Here again, our objective and subjective duties coincide, since for Locke our epistemic duty requires our accepting only what is properly supported by the body of immediate evidence.

Descartes and Locke then are internalists. We have internal access to whether the conditions of objective epistemic duty are satisfied. At least implicit in this characterization is that when we are justified in holding a belief, we are not only aware

of the evidence on which it is based, but we are aware it justifies those beliefs. This is the meta-awareness requirement. For Descartes, we are justified then not just when we are aware of clear and distinct ideas but that they are clear and distinct. For Locke, epistemic justification involves not only having a good reason but being aware that it is a good reason. As we shall see shortly, this meta-awareness requirement leads to the downfall of internalism.

Woods has problems with the CC model. Its requirement of making justification a necessary condition for knowledge apparently rules out as knowledge many instances we regard as knowledge. “There [are] simply too many cases of human knowledge unattended by any discoverable presence of justification” (98). I would add that characterizing justification as having an *argument* for the claim to be justified is in effect to impose the meta-awareness requirement. Seeing why it is problematic is straightforward. The requirement means that one’s perceptual awareness of a tree in full leaf outside one’s office window is insufficient justification for one’s believing that there is a tree in full leaf outside one’s office window. Clearly one could have the perceptual awareness without having any justificatory argument. Indeed, what would such a justificatory argument be? How would one argue cogently for the veridicality of sense perception? What non-question-begging premises could one take as starting points? An argument, being a finite structure, must have starting points—basic premises. But these basic premises must be justified, i.e., argued for. That means we need further argumentation. We are clearly into an infinite regress, whose upshot is that we never have justified belief. In fact, the situation is even worse. Not only would the need for premises generate an infinite regress, the principles of inference or the claims that the premises do in fact justify the move from premises to conclusion need justification, i.e., justification through argument, also. The regress explodes in at least two directions. Indeed, it is easy to see that the CC model should lead directly into complete skepticism. Woods calls such a position big box skepticism.

Woods recognizes that making being able to produce a justification a necessary condition for knowledge would have disastrous consequences for the approach he is developing. There would be few instances of knowledge, since there are few instances of justification in this sense. Hence, the justification requirement “costs us the cognitive abundance thesis, ... the cognitiversity thesis, ... the enough already thesis. It violates the condition that a theory of knowledge not be too hard on knowledge” (103). For such reasons, Woods wants to reject justification as a requirement for knowledge.

However, Woods allows that the CR-model is also open specifically to the Gettier problem. He constructs his Gettier counterexample this way. Suppose Sarah has driven a Buick for many years and Harry, through his acquaintance with Sarah believes this statement. His belief then arises as a causal response to what he has encountered in his acquaintance with Sarah. That is, the belief “was produced by belief-forming devices in good working order and functioning herein the way they are meant to, operating on good information and in the absence of environmental distraction or interference” (93). Harry also knows that a Buick is an American car. So Harry believes that Sarah drives an American car. But Sarah has just started to drive a Cadillac, and Harry does not know this. So it is still true that Sarah drives an American car. But under these circumstances, is it right to say that Harry *knows* that she does? Woods responds this way: The new information that Sarah drives a Cadillac functions as a defeater to the inferential link from Harry’s premises to his conclusion. Adapting Toulmin’s notion of warrant, we can represent the warrant of the inference this way:

From	It has been the case that $(\exists y)(Dxy \ \& \ By)$ and $(\forall y)(By \supset Ay)$
To infer	It is now the case that $(\exists y)(Dxy \ \& \ Ay)$

The information about Sarah’s new car is consistent with both the premises and conclusion, but—to use Pollock’s term in (1995, pp. 85-86)—undercuts the inference in this case.

Woods proposes refining the notion of “good information” in this characterization of CR-knowledge. “Information is good for belief just when it is accurate, current and complete” (114). Woods sees immediately that this requirement of completeness apparently leads to serious consequences for his project. Having complete information is rare. Hence occasions on which we have good information are rare. The knowledge abundance thesis would be imperiled. Woods replies by specifying the meaning of complete information, “all information causally necessary for the formation of that belief yet consistent with its truth at the time in question” (115). So a belief may be well-formed—the premises of an inference may be knowledge—and the principle of an inference adequate (the inference is “well reasoned” in Woods’s phrase), yet the argument fails to transmit knowledge from the premises to the conclusion,

Woods thus distinguishes between knowledge and well-produced belief and well-reasoned inference. If all the other CR conditions are satisfied except for information being complete, we may still have well-produced belief which is nonetheless not knowledge. Requiring good information to be complete would

mean that good information is a rare commodity. Hence the Cognitive Abundance thesis—that human beings have lots of knowledge—becomes tenuous. Woods admits this. But he seems to treat this dilemma quite lightly. He says that well-produced belief and well-reasoned inference lack knowledge-generating transmitting guarantees ... is nothing short of what fallibilism requires” (115). But he also says “Cognitive Abundance tells us ... not to make too much of Error Abundance” (115). I for one am confused at this point. Is Woods holding that we still have plenty of knowledge? Then what of the abundance of incomplete information? Is he saying that well-produced beliefs are enough and we need not aim for knowledge? If an argument has acceptable premises and instances an adequate inference rule, but one for which defeating information which is unknown to us exists, our belief that the conclusion holds is well-produced, we have a *prima facie* case for the conclusion. Is Woods saying that is enough or does he mean to hold that there are sufficiently many *ultima facie* cases to preserve the Cognitive Abundance thesis? I do not know.

There is a further, and I believe much more serious, problem for the CR-model. As presented, it is an externalist account of knowing. Consider the characterization again:

A subject knows that α provided that α is true, he believes that α , his belief was produced by belief-forming devices in good working order and functioning herein the way they are meant to, operating on good information in the absence of environmental distraction or interference.
(93)

The objection is that externalism allows someone to have justification without being aware of the evidence for that justified belief. Although Woods wants to reject the concept of justification, it is clear that the concept he wants to reject is a normative, internalist concept in line with the CC-model. But, as we have just seen, Woods allows that beliefs that may not constitute knowledge may be well-produced. So let us substitute being well-produced for being justified. Now suppose a subject believes α , although α is not true, but otherwise satisfies the CR-condition for knowledge. Although the belief is not knowledge, is it well-produced? Must the subject be aware of the good information for his or her belief-forming devices to be able to operate on that information?

Since Woods several times gives an example involving space aliens, he should not object to the following example. Suppose space aliens from a society technologically significantly advanced beyond that of earth visit this planet. They kidnap a

human being and implant a cognitive device in his brain. The device responds to radio signals by producing in the subject of this experiment an unshakably strong belief that there is a space alien in the vicinity. The space aliens have administered anesthesia while operating on the subject, which has produced complete amnesia of the whole incident. So the subject has no recollection of his encounter with the space aliens and no awareness that they have implanted the device in his brain. The space aliens retreat to their space ship, but through their sophisticated technology continue to observe the subject. They send signals to the device. Each time the subject exclaims that there is a space alien nearby. Is his belief well-produced? Consider: He believes the proposition. His belief was produced by a device in good working order, functioning in the way intended. It responds to the signal the aliens are sending. The signal is information. The implanted device responds only to signals sent from the space ship and thus is accurate and correct. There are no disturbing factors in the environment. Is the belief then well-produced according to the CR-model (taken as presenting conditions for well-produced belief rather than knowledge)? But the subject is totally unaware of any evidence for his belief.

Now take this scenario a step further. The only beings who have access to the transmitters capable of sending a signal to the implanted device are the space aliens themselves. Further suppose the range of the transmitter is limited, so that the implanted device will receive a signal only if the transmitter and thus at least one space alien is in the vicinity. Hence the subject's unshakably strong belief is true. So according to the CR-model, our subject knows that there is a space alien in the vicinity although he is aware of no evidence for his belief. Is this right? Is his belief knowledge? I find that it strains the concept of knowledge to admit that the subject's belief constitutes knowledge. Awareness of some evidence for a belief is a necessary condition for knowledge and indeed for a well-produced belief.

This does not mean, however, that to have knowledge, one must satisfy the internalism of the CC-model. Nor does it mean that one must embrace the CC-model's concept of justification as the only proper way to characterize the concept of justification, especially when trying to characterize knowledge as justified true belief. Several philosophers have proposed nondeontological concepts of justification, Alston in (1985) prominent among them. A belief's being epistemically good for a subject *S* who holds that belief means that *S* has "*adequate* grounds for believing that *p*, where *adequate* grounds are those sufficiently indicative to the truth of *p*" (1985, p. 71, italics in original). A ground for a statement *p* is some indication of the truth of *p*. A

ground could also be a perceptual experience. When I perceive a tree in full green leaf outside my office window, I have a ground for believing that there is a tree in full green leaf outside my office window. Notice that I *have* that evidence or internal access to it. My having the perceptual experience of seeing the tree is my evidence for the tree.

When are grounds adequate grounds? Alston wants to explicate the notion through objective probability. If the ground is *adequate*, the objective probability that the statement it is grounding is true is high² (Alston 1988, p. 269). But while q may be an adequate ground for p , q together with other information r of which S is aware may constitute a very inadequate ground for p . The information r serves to rebut the move from p to q . For q to constitute adequate grounds for p from S 's perspective, q must not be overridden. These comments motivate Alston's concept of justification J_{eg} ("e" for evaluative and "g" for grounds). "S is J_{eg} in believing that p iff ... S's belief that p was based on adequate grounds and S lacked sufficient overriding reasons to the contrary" (1985, p. 77). In appealing to objective probability in explicating the concept of adequate grounds, Alston is incorporating the reliability of the belief-generating mechanism into the criterion of ground adequacy for belief. Notice that absence of awareness of rebutting information, an internalist condition, is sufficient for J_{eg} . One does not need to be aware that the mechanism is objectively reliable or have evidence of its objective reliability or a good argument for it. In presenting Alston's explication of J_{eg} justification and adequate ground in (2005, p. 82), we remarked

Making the criterion of adequacy the reliability of the mechanism that generated the belief commits Alston to a form of externalism. But by insisting that to be justified a belief must be based on adequate grounds, Alston is requiring that to have a justified belief, one must have, be aware of, adequate evidence for that belief. This introduces an internalist character into Alston's account of justification, for he points out, "Grounds must be other psychological state(s) of the same subject..., which are internal to the subject in an obvious sense" (1985, p. 78). This leads Alston to characterize his view as an internalist externalism.

We may contrast Alston's concept of J_{eg} justification with the concept of knowledge (or justification) on both the CC and CR-

² This definition thus introduces the thorny notion of probability into the account of justification. We cannot deal with the problematic concept of probability here.

models. Let's review these two models. According to the CC-command and control-model, "knowledge is the end-state of epistemic undertakings whose strategies for execution and whose conditions of success lie within the command and control of the would-be knower, and so are matters of the agent's volition" (97). So it makes perfect sense to talk of epistemic duty in connection with this model. It presents a thoroughgoing internalist conception of knowledge and justification. As such, it is open to the objection that it leads to skepticism. One is aware of being appeared to tree-in-full-green-leafly. But believing that there is a tree in full green leaf outside my office window involves taking a step. It is thus a matter of volition. But why is taking that step justified? To satisfy my epistemic duty, I need to have a justification for that step. But how am I to have that justification? What are the grounds of *its* justification? The step from the grounds of that justification to the justification itself requires justification. But as is obvious, we are in an infinite regress here, with the upshot that the possibility of knowledge (or justified belief) is cut off.

By contrast, the CR-causal response-model of knowledge claims that:

A subject knows that α provided that α is true, he believes that α , his belief was produced by belief-forming devices in good working order and functioning herein the way they are meant to, operating on good information and in the absence of environmental distraction or interference. (93)

As our discussion of the concept of "good information" shows, the concept is completely externalist. Hence, inspection of the CR-model shows it to present a thoroughgoing externalist account of knowledge, liable to the objection that according to the model one could know that p without having any evidence for p . But on Alston's view, to have justification one must have evidence and be aware that one has evidence and not be aware of defeaters. The charge against an externalist reliabilism is deflected. On the other hand, one's being aware of evidence—at least in the perceptual case—is not a matter of one's volition. Furthermore, one is not responsible for showing that the steps from the evidence to the belief are reliable. It is enough that they *be* reliable, not that S be able to give a justification for taking the step.

We have illustrated this contrast between the CC-model, the CR-model, and the J_{eg} concept of justification primarily through perceptual belief. But it applies to other types of beliefs, including inferred beliefs. The concept of J_{eg} has implications

for what Woods has said about inference, which we will develop in due course. Woods indicates that his discussion of these models is part of his overall argument against traditional epistemology's making justification a "wholly general condition on knowledge" (133). It is "the wrong epistemology for a naturalized logic of mistakes of reasoning" (133) But it seems that Woods has overlooked a whole family of concepts of justification in his argument.

Woods regards fallacy tradition vis-a-vis developing a proper conception of knowledge, differently, finding the concept acceptable as a characterization of fallacy, but finding the traditional list of eighteen as enumerating the extension of this concept as extremely wanting. In the traditional concept, fallacies are instances of reasoning which appear good but are not (134). The common notion of fallacies sees them as widespread misconceptions, and from this Woods takes the notion that a fallacy occurs widely. Woods also notes that fallacies are attractive—people are drawn to commit them and, as habits of thought, they are habits difficult to break—they are incorrigible. Thus, Woods proposes that fallacies are characterized by "the adjectives 'error', 'attractive', 'universal', and 'incorrigible'" (135). He adds to this also that fallacies are bad and takes these five concepts as the defining conditions of "fallacy."

Woods next turns to explicating these concepts. To understand error correctly, Woods contends, one must distinguish between consequence-having and consequence drawing. The distinction may be easily illustrated with begging the question. The conclusion repeats a premise, so the conclusion is a consequence of that premise. But when a statement has to be defended, drawing that statement from itself is inapt. The universality condition holds not necessarily when a pattern of reasoning is always, invariably committed, but when it is frequently committed. Woods adds that frequency is determined by occasion. For example, one can commit the *ad hominem* fallacy only when one is reasoning from a premise attributing something bad to a person who has made an assertion. We have an error only when this attribution is inapt. The universality requirement further adds that this error occurs more often in situations that are occasions for committing the fallacy than errors occur in situations just calling for reasoning. *Ad hominem* reasoning may similarly illustrate the attractiveness condition. Woods does not make explicit what attractiveness includes. The attractiveness he is concerned with arises in situations where a pattern of reasoning which is ordinarily correct is inappropriate. To rebut what someone has said, we may say something about that person's lack of reliability. This may be a perfectly appropriate way to reason against that person's assertion. Information about the person's unreliability to

make a certain claim is ordinarily relevant to undercut his reliability and thus undercut the force of his argument. For the negative information to constitute the fallacy of abusive *ad hominem*, it must be irrelevant to the question of his reliability.

We may be able to recognize that a piece of reasoning involves error, but still find the reasoning attractive and be inclined to reason that way. This leads to the incorrigibility condition. The reasoning appears right, is shown wrong, but still appears right and consequently attractive even after its wrongness is detected. Woods takes the fact that a pattern of reasoning which in certain instances is the wrong way to reason but is generally the right way to reason as explaining instances of the incorrigibility condition. The “occasioned re-commission is not typically the wrong way to reason” (145). Woods notes that if his analysis of the traditional fallacy characterization is correct, it creates a high burden of proof for someone who wants to show that a pattern of reasoning, say a traditionally recognized type of fallacy, really is a fallacy. One would have to show in particular that the fallacy is universal and incorrigible as the concepts have been explicated here.

Woods next turns his attention to how the concept of belief contributes to the concept of error. He first distinguishes *s*-beliefs from non-*s*-beliefs. Suppose Harry needs to find out whether α is the case. This determines an agenda for Harry where thinking in a certain way that α or thinking in a certain way that not- α are possible end states. The “in a certain way” “means that the end-state achieved by thinking in that way is a ‘cognitively saturated’ end state” (155). Cognitive saturation means being “in a state in which the cognitive impulse that drove [one’s] agenda has subsided” (155). Such beliefs are *s*-beliefs. A non-*s* belief is an “agenda non-closing belief” (156).

We may also distinguish between first-person and third-person belief ascription. Lou, Harry’s friend, may very well say “Harry believes that α although α is false.” This is perfectly consistent. But *Harry* cannot say “I believe that α but α is false.” This assertion is pragmatically inconsistent. Woods calls *s*-beliefs whose first person ascription will not allow the person ascribing also to assert that the belief is false first-person *s*-beliefs (156). They have this particular feature for an understanding of error: “They make error unrecognizable. That is to say, they make α ’s error unrecognizable in the first-person *s*-case in the here and now” (156).

But the somewhere else and later are different. Critics or one’s own reflection may bring one to question one’s *s*-beliefs, even to see that a belief was erroneous. So, although saying “I believe that α but α is false” is pragmatically incoherent, saying “I believed that α but α is false” is not. As Woods puts it, “Error

detection when it is possible at all is an after-thought” (163). However, Woods wants to make the point that one’s current belief that one’s previous belief was erroneous might be an erroneous belief also, although one could not recognize that in the here and now of having the belief that the previous belief was erroneous. Woods calls this the “no escape thesis.” He coins the notion of an epistemic bubble to express this thesis. One may experience oneself as knowing that α , but this knowing is not actual, yet one cannot distinguish in one’s current state “between his knowing that α and his experiencing himself as knowing that α ” (162). Likewise, when one believes now that what one believed earlier is an error or when one believes that someone else is in error concerning what they believe now, one is still in an epistemic bubble with respect to what one believes now. Only incorrigible states of knowledge escape the end-state bubble.

Could one have avoided an erroneous belief one holds to begin with? Woods approaches this question by defining what he calls an s-belief. Suppose one’s agenda is to know whether α . Suppose one comes to think either that α or that $\sim\alpha$ in such a way that one is satisfied that one’s question has been answered. He is in a cognitively saturated state, “the cognitive impulse that drove his agenda has subsided” (p. 155). But if one thinks that α and one’s thinking is a saturated end state, α could still be false. So one’s thinking that α is an error. Woods argues that we cannot avoid being in such states. “Nothing in the empirical record leads the slightest support to the suggestion that preventing ourselves from being in s-states is something that lies within our capacity to bring off on a scale that would qualify as a general policy” (171-72). Woods has argued that states where we are in error do not have markers that would readily flag that they are not erroneous. In line with the skeptics of antiquity, he rejects the Stoic idea of cataleptic beliefs.³ So if part of being human is to yearn to come to states of satisfaction as to whether α or $\sim\alpha$, our having such epistemic yearnings makes unavoidable our being in error on occasion.

Yet Woods believes this is not cause for epistemological skepticism or more generally a pessimism about our belief-generating capacities. He supports an epistemological progressivism towards error correction by citing principles he has already put forward. We do favor newly established beliefs over old. The convergence of the normal and the normative (52) proposes that *ceteris paribus*, how we reason from premise to conclusion is the way we should reason. The Enough Already Thesis (88) claims that humans are right about enough of the right

³ See Stough (1969), especially pp. 38-44.

things enough of the time to thrive. The nature of an error correction process also gives reason for progressivism. If state 2 corrects state 1 and 3 corrects state 2, seldom if ever will state 3 return us to state 1. Through this process, error is successively reduced.

Woods concludes Chapter 5 on error with the question of whether these considerations constitute a solution to the error detection problem. To repeat, for him a proposition is an error if it is incompatible with one or more true propositions and it must be false all along, not just false now although it had been true. He feels that the CR-model gives us sufficient grounds to claim the error-detection problem solved. We can solve the problem if “(a) ... we have error detection devices and (b) by and large our devices work as they should” (p. 183). On the CR-model, we do not have to be able to specify these devices or to show why they are reliable. The how and why of these error detection methods are the tasks of the cognitive scientists. We do not need to have their answer to be justified in identifying a proposition as an error, just as we need not know how perception works to be justified in believing there is a tree in full leaf outside my office window upon being appeared to in the appropriate way.

Does Woods run afoul of the basic objection against pure externalism? From the perspective of Alston’s J_{eg} criterion, we can applaud his insistence that there is no requirement to explain why the error-detection devices work. What then about the need to be aware of evidence? I think that here Woods has covered the case in his characterization of error-detection. “For human agents, error detection ... is rooted in their possession of phenomenal states the recognition of whose own erroneousness is not *concurrently* possible for them” (163), but such recognition is actual in a later state. Awareness of these phenomenal states constitutes the needed awareness of evidence. Further, to detect error, Woods also insists that where believing proposition E would be an error, to recognize E’s erroneousness, “there must be one or more true propositions with which E is *seen to be* incompatible” (181, italics added). So one must be aware of evidence of error for a proper detection of error. The objection against pure externalism does not hold.

One can still have reservations here on two grounds. First, just what are the human error detection devices and their presumptive reliability? We can readily enumerate human belief-generating mechanisms—sense perception, *a priori* intuition, taking the word of others. An epistemological tradition vouches for a presumption for the beliefs these mechanisms generate.⁴ Woods owes us an account for these error-detection devices.

⁴ See Rescher (1977), p. 37.

Secondly, Woods owes us an account of how one recognizes the truth of the propositions seen incompatible with E and an account of how the person detecting the error comes to believe their incompatibility with E. Incompatibility is a logical relation which is recognized *a priori*. Recognizing incompatibility, then, is not problematic. But must the propositions seen incompatible with E be seen true or is seeing that they are justified beliefs enough? Woods is in a dilemma here. If the incompatible propositions are true but not seen true, then how does their recognized incompatibility constitute a sign of E's being erroneous? If one simply recognizes that they are justified, that there is a presumption in their favor, then they may still not be true. At best we may have a presumption of erroneousness, but is that sufficient to solve the error detection problem? Woods owes us answers here.

Woods titles Chapter 6 "Economizing." He points out that cognitive procedures that lead to error need not be an overall bad thing. They give us occasion "to learn from experience. They are fruitful contexts for trial and error" (p. 185) Woods's next point is that due to the convergence of the normative and the normal, errors made in drawing a conclusion are far less frequent than errors in accepting the premises one draws a conclusion from. Human beings are limited in terms of cognitive resources available to them, in particular the amount of information which is available to them, the time to process, i.e., draw inferences from that information, and their capacity to remember it all. Having made this observation, Woods asserts what I take to be an empirical claim that individual cognitive agents proportion their cognitive targets to their cognitive resources (194). These resources ordinarily will not allow the cognitive agents to satisfy the standards of a deductive argument nor of a confirmation argument. But as Woods puts it, the default requirements for such targets to be acceptable neither call for deductive validity nor inductive strength. Invalidating inductive weaknesses are not reasons for rejecting that the reasoning is good (193). Rather, "something is an error of reasoning only if it violates a rule of right reasoning that is *contextually* in force" (193, italics in original).

Woods goes further here, making a statement about the connection between rules and the rightness of the reasoning they mandate. "What," he asks, "is a correct rule of reasoning? It is a rule made correct by the rightness of the reasoning it mandates. So it is not the rule that makes Harry's reasoning right. It is the reasoning that makes the rule correct" (p. 195). There may be a lot to ponder here. But surely Woods' characterization is unacceptably vague. What is "right reasoning"? Woods next considers the charge that humans' projecting from samples to a popu-

lation frequently involves the error of hasty conclusion. He claims this charge itself is a hastily drawn conclusion. Woods questions that, except in certain contexts, we are not generalizing to universal but generic generalizations. To be justified in inferring a generic generalization, one would still need to generalize from a sufficiently representative sample, “a sample of appropriate size and randomness” (212). But, Woods points out, generalizations from small samples need not be hasty generalizations, if such generalizing involves pattern recognition and that “pattern recognition is tied to our ability to recognize natural kinds” (213). Of course, one’s generalizing can still be in error. One may recognize a pattern which a particular case does not exhibit. But, for Woods, this is not a bad thing.

Against the charge that causal reasoning frequently involves the errors of confusing correlation with cause and temporal succession with causal succession, Woods replies by first reminding us that given the traditional definition of fallacy, to claim that human reasoning instances the false cause fallacy or the *post hoc* fallacy, those patterns of reasoning must occur with a significant frequency. Correlation and constant temporal succession suggest there may be a causal connection. However we shall be subject to the charge of fallacy only if, in such situations, we judge that there is a causal connection “at a rate that out-distances the frequency of error-making in general” (217). However, as a matter of fact, we do quite well in *correctly* recognizing causal connections and discriminating them from correlations and mere temporal successions. Hence, “It is hard to see how it could be typical of us to mismanage these discriminations with a frequency that outruns the general rate of error making” (218).

When properly understood, humans are not fallaciously projecting universal generalizations but rather generic generalizations. What then of reasoning from these generic generalizations? Some have argued that such reasoning constitutes a third way of reasoning, beyond deduction and induction. Woods addresses this question in Chapter 7. If reasoning from premises to conclusion is done correctly, a certain relation holds between the set of premises and the conclusion. In third way reasoning, “The underlying R-relation is subject to rupture upon the addition of true premises consistent with the old ones and also with the conclusion” (221). Woods’s using the word “premises” in characterizing statements introducing this new information is unfortunate. Such statements do not function as premises but as defeaters, to use Pollock’s term. By claiming these statements are consistent with the conclusion, Woods seems to be restricting the defeaters just to Pollock’s undercutting defeaters. Ordinarily, the testimony of a witness positively supports the conclusion of a case

an advocate seeks to build. But evidence that the witness is somehow unreliable ruptures that relation, even though that evidence is consistent with both the premises and the conclusion. The fact that R has been ruptured yet both premises and conclusion remain true means, for Woods, that the rupture is “alethically benign” (221). One wishes Woods would have defined this term. In the case of an undercutting defeater, one could say that the general relation between premises and conclusion is left untouched by this undercutting defeater, although in this particular case it is undercut. R’s are not-Q, even though some can be P, falsifies that P’s are Q. How then would a generic statement of the form “P’s are Q” be falsified? Woods has a clear answer here. Suppose R is a species of P where the extension of R is a proper subset of the extension of P. P’s in general are Q but R’s in general are not. That R’s in general are not-Q falsifies that P’s are Q. It takes not one example but a regularity between a definable subset to falsify the generic relation. The important thing as Woods sees it is that R be a species. If a P accidentally is not Q, this fact does not overturn P’s are Q. One bird with a broken wing does not falsify that birds fly. But a whole species of birds which do not fly does. So where α is not Q because α is a member of R, to reason

P’s are Q
 α is a P
 Therefore α is a Q

does not infer a false conclusion from true premises. Rather the first premise is not true. Just what is the significance of R’s being a *subspecies* of P? Woods does not say. One could argue that if the extension of R is a species, to predicate R of α is to indicate that R expresses to some extent the *essence* of what it is to be an R, that R is not an accidental property. This is a plausible interpretation which deserves developing. In particular it requires meeting or at least considering Quine’s objections to essences.

Woods indicates that his goal in Chapter Eight also is basically negative. A number of logicians have proposed logics of defeasible inference or non-monotonic logics. Woods finds that these logics have little contribution to make to properly arriving at the tasks of drawing consequences nonmonotonically from sets of premises. He wants to maintain that when someone, reasoning non-monotonically, draws a conclusion from premises, he is not discerning a consequence relation between the premises and the conclusion drawn from them. Rather, when one draws α from Θ , one’s belief-producing devices produce the belief that α and also produce “some disposition to recognize that it was

the information in Θ that got [one] to see that α or made [one] think that α " (292). This position affirms Woods' externalism, which we have considered before. He proceeds to argue for his view by arguing against the consequence relation view. I find his argument obscure, and since it is not part of his positive program, I shall not consider it here. I also believe the view can be counterexampled. Suppose an Alpha-Centaurean scientist unobserved arrives on earth to study humans' belief-generating capacities experimentally. He kidnaps Harry and implants a belief-generating device in his brain. Anesthesia destroys any memory of the kidnapping or the operation. The device is designed to produce a belief in Harry that there is a tiger in the vicinity when he sees a perfectly round sphere and hears the blast of a horn and that these experiences constitute sufficient evidence for the tiger. The scientist knows that Harry takes walks in a park near his home. He goes to the park and puts a large sphere in front of a bush. When Harry walks by, he cannot miss it or miss forming the startling perceptual belief that there is a perfect sphere in front of him. The scientist then hides behind the bush, horn at the ready. Harry comes by and sees the sphere. No doubt, he really notices it. The scientist blows the horn, BOOP! Harry screams "tiger" and runs away as fast as he can. All belief-generating mechanisms are functioning as they should. There is no environmental distortion. Has not Harry *inferred* that there is a tiger from the two perceptual beliefs that there is a perfect sphere in front of him and a horn has blasted by way of what the scientist did to his brain to now reason according to this inference rule?

In Chapter 9, Woods turns to the question of error in connection with taking the word of others. Here, not reasoning but misinformation is the cause of error. Woods sees us as vulnerable to false statements of fact and distinctly vulnerable. Believing what we are told and telling others have an adaptive advantage and much of what we believe is believed on the basis of the word of others.

Does *knowing* a proposition through the word of others presuppose justification? If one claims to *know* that α on being told that α , one cannot say both that one knows that α but that one has doubts about the reliability of the source which told him. But how did one come to know that the source was reliable? Woods believes that the CR-model can help us here. When someone tells us that α , our belief-forming device or mechanism would not produce in us the belief that α unless it concurrently produced the belief that the one who told us that α was reliable. This is not to say that we are justified in believing that the source is reliable. We are *caused* to have certain beliefs, in particular that the source is reliable by our properly-functioning be-

belief-generating mechanism. “These causal facts are facts about a certain way in which our knowledge is produced. It is a way that makes no claim on justification” (303). We can see that it makes no claim on justification when we remember that according to the CR-model, knowledge requires that the belief-generating mechanism be in good working order, functioning as designed, operating on good information, and with environmentally distorting factors absent. None of these factors presuppose that the subject is aware of that factor. By the same token, I cannot but see in this argument, in Bertrand Russell’s words, theft over honest toil. We have a belief that a certain source is reliable. But that belief was produced by a mechanism of whose reliability *in this case* we know nothing. In particular, we may form the belief that α is reliable with no awareness that our belief generating mechanism was operating on good information. Can it be said then that we *know* the source to be reliable in this case?

Woods continues by asking whether the causal regularity of our belief-forming system has normative significance. He refers to the convergence of the normative and the normal principle, which we discussed earlier, as reason for affirming *normative* significance. Recall that the principle asserts “As a first pass, and when there aren’t particular reasons to the contrary, how we *do* reason from premises to conclusion *is* typically how we *should* reason” (52). But given that our belief in the conclusion is *caused*, is the agent reasoning here? I think not. If belief in the conclusion is somehow *caused*, then one may have the belief without any awareness that one’s premises were relevant to the conclusion in a way that *conveyed* belief from the premises to the conclusion. One could believe those premises, and then become aware that one believed the conclusion, without in any way seeing that belief in the premises led to belief in the conclusion by seeing that if one believed the premises, one should believe the conclusion also, at least *ceteris paribus*. Woods acknowledges the revolutionary implications of his position here. “It renders nugatory any critical examination of the argument’s force. It makes the goodness of the argument moot” (305). By contrast, do not our comments immediately above show Woods’s position radically defective?

Woods’ externalism indeed leads him into a rather extreme position when he discusses mathematical knowledge. You prove a theorem α . You come to recognize some of the premises of the proof by being taught them, where the knowledge of that premise has been handed down through a long chain of teachers. Woods comments, “You couldn’t possibly have attained your knowledge that α without the distal tellings that your proof *relied on*” (306, italics added). (X tells Y that α distally when X does not tell Y directly but through a chain of intermediate inter-

locutors.) But just what does “relied on” mean here? I think Woods means causal dependence of one teacher teaching the next in the chain. This is correct as far as the context of discovery is concerned. But can we be satisfied that the requirements of the context of justification have been satisfied? For someone to come to recognize α as a mathematical theorem, would not one have to “see” for oneself that α was true *a priori* immediately or that α followed from certain axioms which he saw to be true *a priori* where the steps of inference each could be seen valid *a priori*? I still do not see how implication can be explicated externalistically as Woods would have us believe.

Obviously, although we do accept much of what we are told, we should not accept all of it. As Woods points out, in a situation when we are told something, there may be a trigger which indicates that before acceptance there should be a due diligence search. Woods sees such situations as rare. He does enumerate some triggers for the due diligence exercise. That some interlocutor’s word is an evaluation, not a description, triggers ordinarily that it should not merely be accepted but only when properly defended. A perceptual report which involves misperception and is immediately corrected, an interpretation, e.g., “Mother Theresa had a generous disposition”, recognition of some unreliability about the subject matter, are all triggers that one should not simply accept what one has been told. However, Woods greater concern in this discussion is not with these triggers but with the fact that in mechanisms that generate beliefs which may become premises of our reasoning, there may be many conditions which may produce error but which are not accompanied by triggers. Woods explains that when one acquires a belief either through perception, say-so, or inference, one is experiencing belief change. With perception or say-so, there may be a long chain to the eventual production of one’s belief. By contrast, an inferential chain may be short. Woods proposes talking about the number of steps leading to a belief in a particular case as “the *surface* of a medium of belief-change” (330, italics in original). “The larger the surface size of a medium of belief change, the greater the likelihood of error” (331). Woods proposes this hypothesis as intuitive and worthy of empirical test.

From discussing errors which may arise when being told something, Woods turns to not being told in Chapters 10 and 11. In Chapter 10, he identifies what he calls virtual beliefs, which are not genuine beliefs at all. Humans learn in a rough and ready way how the world works, at least so far as it affects our getting on in the world. But in many cases, we are not aware consciously that we have these “beliefs,” which makes them virtual. They constitute a storehouse of expectations for a given person, who

may learn something occurrently, but quickly store it as a virtual belief. We become conscious of our virtual beliefs when the world contravenes our expectations. They pose a special problem for knowledge, since they may be used in reasoning yet “when errors are committed, it only stands to reason that they will go unspotted” (358), until we come to a surprising situation. Yet one cannot do without virtual beliefs.

Chapter 11 is devoted to abduction as a context in which to explore how our ignorance of something we desire to know but do not contribute to good reasoning (364). Woods rejects Peirce’s view that a hypothesis conjectured on the basis of its explaining some surprising fact constitutes neither knowledge nor justified belief and should not be believed until properly tested. Rather Woods holds that on many cases those who have arrived at the hypothesis by abduction and have come to believe the hypothesis on the basis of this abduction can actually have new knowledge. He illustrates his point with this example: Coming home at a time you expect no one else there, you find the back door ajar and a window broken. Although several hypotheses explain this fact, you believe that a burglar has entered your home at some time earlier in the day. Suppose this hypothesis is correct—there *is* a burglar in the home. Your belief in your hypothesis is true. Woods adds that if properly produced, that belief constitutes knowledge. “Your hypothesis-forming devices and your belief-forming devices have fallen into an accord in which each partner was in good working order, and operating properly on good information” (377). The question is not whether your believing is justified but whether your belief-forming mechanisms have been properly stimulated. On the CR-model, you have knowledge. On the CC-model, however, you are still in a state of ignorance. Woods concludes that deciding whether or not one has knowledge in this situation depends on settling the rivalry between the CC and CR models (378).

Here I must protest. The belief-generating mechanism is *not* functioning properly in Woods’ scenario, certainly if one thinks of acceptance as the last step in belief-formation. In a perceptual situation, unless one is aware of some rebutting condition, one correctly accepts what one’s perceptual belief-generating mechanism produces. But this is not true for hypotheses, for here in general when a hypothesis is formed, one is aware that there may be alternative hypotheses. As Cohen points out, there is a place for the will to intervene. In such cases, the process or procedure for accepting what one’s hypothesis-generating mechanism has produced is not immediate. As Cohen says,

What is the crucial difference between the kind of mechanism that generates presumptively acceptable beliefs and the kind that generates beliefs which are not presumptively acceptable?... Wherever there is standardly some opportunity for the intrusion of a voluntary element into the mechanism, the kind of belief generated is not presumptively acceptable, because a mistake may be made in the discharge of this voluntary element (1992, pp. 130-31).

To believe with justification that a hypothesis gives us the best explanation, one needs justification for that hypothesis. So, unless Woods can show that inference to a hypothesis is always inference to the best hypothesis, his claim that properly-functioning belief-generating mechanisms which generate hypotheses generate beliefs which constitute knowledge is just plain wrong.

Woods sees abductive reasoning as a way of showing that not all reasoning according to the patterns of affirming the consequent and denying the antecedent is fallacious. Instead of the bald

1. If α , then β
2. β
3. So α

which violates the validity standard, one is actually reasoning

1. If α were the case, β would be a matter of course
2. β
3. So, defeasibly, it is reasonable to conjecture that α .

One reasoning this way is most likely not trying to meet the validity standard. If so, one should not judge the reasoning according to that standard, but according to the standard for which it is intended. "Since interpretation precedes assessment, the default position with respect to reasonings in this gross form is to find a reading of them under which they come out all right. One such is an abductive reasoning" (383).

In Chapter 12 entitled "Asking," Woods presents an extended discussion of begging the question. His concept of this form, though, seems quite idiosyncratic given the standard textbook characterization. We do not have it that an argument begs the question if the conclusion appears as a premise—albeit typically otherwise worded—or is presupposed for some premise to be meaningful. For Woods, one interlocutor begs the question against another if in arguing against his thesis τ she presents a

premise α , negatively relevant to τ , which he does not concede nor can one ascribe to him. Woods asks whether begging the question so characterized is an error of reasoning. He concludes not. The interlocutor has not made a mistake in inferring not- τ from some set of premises containing α . Rather the interlocutor has mistakenly attributed α to her opposing interlocutor. If she is assuming that by putting forward this premise, she is putting forward a statement that he accepts, she has made a mistaken assumption about her interlocutor or has made an error in selecting α as a premise for her argument.

Woods terms begging the question in its familiar textbook sense “babbling.” He argues that instances do not satisfy the criteria for being a fallacy, either in Aristotle’s sense or the traditional sense. The Aristotelian concept of fallacy applies just to syllogisms and an argument consisting simply of a conclusion and its repetition as a premise is no syllogism. On the traditional conception, the challenger will obviously not accept as a premise a conclusion she has challenged. The ascription of babbling does just that, and thus makes a premise attribution error. On the traditional conception, a fallacy has to appear cogent, even though it is not. Making an obvious premise attribution error does not appear cogent and so babbling is not a fallacy, although it is an error. I must protest here also. Even if the Aristotelian concept of fallacy applies only to syllogisms, and the traditional conception requires that the argument appear cogent, an argument that repeats the conclusion as a premise or presupposes that conclusion fails to give a challenger a cogent reason, presuming, as we may, that she does not accept that conclusion. Otherwise, why argue? If an argument fails to present a reason for the conclusion which a challenger should find acceptable, why is it not a fallacy?

In Chapter 13, “Getting Personal,” Woods turns to *ad hominem* arguments. Are they fallacies? He concludes that they also are not. As Woods sees it, the proponent of an *ad hominem* puts forward a retort to an opponent’s claim. The proponent then “concludes from this that the adequacy of her opponent’s case should be called into doubt” (463, italics in original). It may also be that “She further concludes from *this* that there is reason to think that her interlocutor’s position is false” (463, italics in original). The word “concludes” is italicized because Woods wants to make this point: If “concludes” means that the proponent regards her inference as deductively valid, the opponent is clearly in error. Why is this inferring not a fallacy? Woods cites two reasons based on his conception of fallacy. First, there is not “a jot of evidence in the empirical record” that such inferences are made with deductive intent” (463). Woods declines to judge whether “concludes” is possibly used with non-deductive intent

because although “called into doubt” and “reason to think that” suggest a non-deductive type of concluding, he believes that a non-deductive logic has yet to be worked out “definitively, [so] that citing it at this stage is more promissory than helpful” (464). Secondly, the available empirical data do not support that such inferences are made with sufficient frequency.

Woods entitles Chapter 14 “Taking” in the sense of taking for granted or given. What is taken for granted Woods calls data. In developing a theory of a given subject matter, e.g., the behavior of organisms of a certain species, one is constrained by the data concerning their behavior. Woods asks two questions concerning theories—how are they related to their data and, given our understanding of this relation, what may we say of “how psychologists and applied mathematicians deal with what they take to be errors of statistical and probabilistic reasoning” (469). He turns to model based theories, where modeling involves idealization and abstraction. This feature immediately raises two problems. Idealization involves what is false to the data, while abstraction overlooks (“suppresses” is Woods’s term) certain features of the data. So how do model based theories enhance our understanding of the data?

To answer this question, Woods point out that the data are not raw. Their very representation processes the data in some “nontrivial ways [which] are distortive” (471). Representing objects as colored or things using concepts that are precisely tailored to process the data in distorting ways seems contrary to respecting the data. But such distortion is not peculiar to empirical science, but a feature of cognitive encounters with empirical reality. Yet science may distort further. Representational and linguistic distortion is part of our perceptual and cognitive mechanisms, part of how we generate beliefs about the world. Data bearing such distortions are products of our perceptual or cognitive mechanisms. These distortions are by no means willed. Further distortion is willful on the part of the theorist. This voluntary distortion attempts to prepare the data “for theoretical engagement” (471). From these considerations, Woods concludes that there are two types of models—theoretical models (T-models) and data models (D-models). But Woods regards the distortion in both these types of models as benign. D-model distortion is benign first because we may regard such models as explications of the “raw” data. They are designed to “prime” the data for the application of further theoretical models. To help us understand this, Woods point out that in economics there is no empirical evidence that utilities are infinitely divisible. But neo-classical economists stipulate that they are. Why? “It enables economic theory to appropriate the fire power of the calculus” (474). But instead of the conceptual clarification of explication,

the further distortion of the T-model constitutes “a reconceptualization of the raw” (475), a new concept altogether of what the theory is about. Two distinct concepts are involved and ought not to be confused.

Woods turns in the last three sections of this chapter to three specific fallacies involving reasoning with probability which humans have been charged with—the conjunction, gambler’s, and base rate fallacies. Those claiming that the conjunction fallacy occurs with significant frequency cite examples like the following:

FACTS: Linda is 31, single, outspoken, extremely intelligent. In college, she was a philosophy major, active in anti-discrimination, social justice, and anti-war movements.

QUESTION: Given this information, which of the following two statements has the higher probability?

1. Linda is a bank teller.
2. Linda is a bank teller active in the feminist movement (See 478.)

The correct answer is (1), but a large majority choose (2). Does this fact require us to concede that people have a problem with the probabilistic reasoning involved here? Woods points out that a necessary condition for this interpretation is assuming that people are reasoning according to the classical probability calculus. But is this assumption correct, even likely? There is at least one alternate explication. The word “probability” can be (mis)interpreted as “plausibility.” The facts given all support the second conjunct of (2) and not the common conjunct of (1) and (2). Woods believes that plausibility theory, although not nearly as developed as probability theory, gives the more satisfactory answer here. So the subjects of the experiment are not engaged in fallacious probabilistic reasoning but are reasoning plausibilistically.

The gambler’s fallacy is easily illustrated. Woods asks us to suppose we have the very unusual run of 16 straight tails in flipping a coin. What is the probability that the next flip is a head? The correct answer is $\frac{1}{2}$. If someone argues that because a run of 17 straight tails is extremely low, $1/131,072 \approx 0.0000065$, the probability that the next flip will show heads is greater than .5, the person has committed the gambler’s fallacy. Woods again argues that we should not regard this pattern of reasoning as a fallacy. Unfortunately I find his reasoning virtually opaque here.

One can also easily illustrate the base-rate fallacy. Here again is Woods’ example: Suppose D names some disease. Suppose the rate of D in the homosexual community is three times

greater than in the non-homosexual community. Suppose Leslie has D. What is the probability that Leslie is a homosexual? Asking just what is that probability given just the facts presented is to pose a question impossible to answer. But one can argue that we are not being asked to make a determination just on this information alone but on the base rate of homosexuality in the population. Suppose it is 10 per cent and we add that information to the previous two facts. Now can we calculate the precise probability that Leslie is a homosexual? Woods reports that experimental evidence shows that people tend to give a wrong answer here—to overestimate the probability of Leslie's being a homosexual by overlooking the base rate of homosexuality in the population. He counters that the appellation is wrong. People are not making a mistake in reasoning. Rather they do not know how to compute the probability, more specifically they don't know how to apply Bayes' theorem, and thus guess incorrectly. But, as Woods stresses, guessing wrongly is not a fallacy. In most cases, when asked to do something one does not know how to do, one admits that incapacity outright. One does not guess. Since it is part of the definition of fallacy that committing that fallacy occurs frequently, at least in a given population, "in its guessing-form, the base rate fallacy is not a fallacy" (491).

In the final chapter, Chapter 15, Woods turns his attention to the concept of fallacy overall. He believes that he has now said enough, throughout this essay, to substantiate his claim that the traditional concept of fallacy and the traditional list of fallacies do not, by and large, match up. He has argued that fourteen out of the eighteen patterns on the traditional list are not fallacies in the traditional sense and postulates that the remaining four could be handled in the same way. He asks whether there are fallacies not on the traditional list, patterns that *do* nonetheless satisfy the traditional definition, or whether the concept of fallacy is empty. He thinks the concept not empty. He has defended the Error Abundance thesis, "Human beings make errors, lots of them" (88). Now unless these errors are so diverse that sorting them into classes makes little sense, "There are errors that it is characteristic of us all to commit in numbers that fulfill the frequency condition of Error Abundance" (513). Woods puts forward a candidate—Powers' Paradox (513):⁵

1. If $\alpha \supset \beta$ then if α then β .
2. If, if α then β , then $\alpha \supset \beta$.
3. Therefore, $\alpha \supset \beta$ if and only if if α then β .

⁵ Woods cites a private conversation with Lawrence Powers as acquainting him with this argument.

If one accepts this argument, then one must concede that material implication and implication coincide. (Woods' point would be easier to see if he had written " α implies β " rather than "if α then β ".) But this is obviously wrong. It is not the case that every false statement implies any statement whatsoever or that a true statement is implied by any statement. The problem is with premise (1).

Does Powers' Paradox satisfy the EAUI criteria—error, attractive, universal, incorrigible, which Woods believes an argument must satisfy to be a fallacy? Woods claims without presenting evidence that the pattern satisfies A and, when dealing with those already familiar with the material conditional, "such as logic students in large culturally diverse universities" (514), the pattern satisfies U with respect to this population. Woods's experience teaching logic students in various universities gives him the empirical evidence for this second claim. This experience also gives him evidence for the I criterion. Students can be instructed in the error of equivocation involved in this argument pattern yet still be taken in by it when presented with this pattern again—the mistake is incorrigible. There you have an empirical argument that Powers' Paradox satisfies the EAUI criteria and so the concept of fallacy does not have an empty extension. Logicians (and cognitive scientists) may be able to discover more.

Having made this point, Woods now turns to answering three questions. Hamblin asks why logic has not developed an appropriately deep theory of the fallacies. His diagnosis is that logic is dialectical and the traditional fallacies are dialectical errors. Since modern logic does not conceive of logic as dialectical, it is not in a position to develop such a theory. By contrast, Woods says that *his* answer is the concept-list misalignment thesis. He has argued across the board that fourteen at least if the traditional list of fallacies, together with certain others, fail to satisfy that EAUI criteria. Trying to understand why these patterns are fallacies, as traditionally understood is wrongheaded. Woods regards this argument sufficient to answer Hamblin's Question.

Finocchiaro has asked why "have mainstream logicians been unable to see this failure of fit between concept and its purported extension?" (516) Woods responds by pointing to the Fallacies Dilemma. A fallacy instances "bad reasoning whose badness is inapparent" (516). So to recognize that an argument is a fallacy one must see it "as bad and yet not see it as bad" (516). If an argument satisfies the EAUI criteria, it will be an error that is due to incorrigibility, which one will frequently continue to make, even after the error is pointed out. This "postdiagnostic recidivism" makes the error that much harder to recognize. I am not moved by Woods's answer here. Recognizing the

error in a particular argument may be difficult, but what if one presents a logical analogy—an argument with the same principle of inference but with premises obviously acceptable and conclusion obviously unacceptable? The Fallacies Dilemma argument seems to have forgotten about employing analogies to recognize fallacious arguments. Will such recognition help with the recidivism rate? Here is an open empirical question.

In addition to the error detection problem, Woods points out that mainstream logicians basically recognize two standards for connection adequacy in arguments—deductive validity and inductive strength. But Woods claims that most human reasoning is third-way reasoning, How does this contribute to answering Finocchiaro's Question? If the mainstream logicians are looking to diagnose (connection adequacy) fallaciousness as either deductive invalidity or failure of inductive strength and arguments in the extension if the traditional eighteen are neither, again it appears wrongheaded to apply the mainstream criteria to understand their fallaciousness. Likewise, why should seeking to apply these criteria lead to recognizing that these arguments satisfy the EAUI criteria for fallaciousness or that these criteria do not explain the fallaciousness of these arguments? Finally, Woods turns to his own question: "What explains this lingering attachment to the traditional list?" (519). Woods responds by pointing out again that genuine fallacies are hard to recognize. If so, then there is a barrier against recognizing them as the data on which a theory of fallacies needs to be built. Stick then with the instances of the traditional list as examples of fallacies. Fallacy theorists are not heeding the proper data for fallacy theory.

In the very last section of the very last chapter of this book, Woods takes stock. His naturalized logic project is akin to naturalized epistemology, but is only in its infancy. He has certain strong convictions, both negative and positive. His negative convictions are expressed in his argument that the traditional list of fallacies and the traditional characterization of fallacies are misaligned, that validity and inductive strength are the wrong standards to apply in assessing fallacies, that arguments should not be modeled as premise-conclusion reasoning, and that ideal modelling instances "normative presumptuousness" (520). On the positive side, he is convinced that "We owe our cognitive well-being to the smooth functioning of our cognitive devices" (520), which function smoothly by default, that we can choose premises correctly, distinguish relevant from irrelevant, plausible from implausible, and projectable from non-projectable. But he has no understanding of what smooth-functioning involves or how we are able to make these discernments. Woods leaves these as open questions.

To my mind, there are rather bigger questions for Woods's project. His positive convictions express his externalism. However, as we have pointed out earlier, externalism as an epistemological theory has distinct problems. We can construct scenarios where one's cognitive faculties are functioning properly but one is aware of no evidence for certain specific beliefs they generate. Are such beliefs then acceptable as premises for what to believe or do, to recall Ennis's memorable characterization of critical thinking? There is an even further problem. Unlike an externalist such as Plantinga who identifies our cognitive faculties before discussing their reliability, Woods alleges that there are cognitive faculties that function properly without, in a number of instances, identifying these faculties. How then can we accept the claim that they are reliable by default, if we do not know what is being alleged to be reliable? Of course, externalism is attractive for someone advocating an empirical theory. Our cognitive mechanisms and their reliability should be open, at least to a certain extent, to empirical investigation. But this empirical concern should not lead us to overlook the problems with externalism or to view, as we have already noted, externalism and internalism as the only two alternatives. Alston's externalist internalism as a theory of justification constitutes a third alternative, one that, we have argued, avoids the objection against pure externalism or pure internalism.

I leave Woods essay with an even bigger question. What is the value of this empirical investigation for the central questions of argument evaluation: Are the premises of an argument acceptable? Are they so connected to the conclusion of the argument that their acceptability is transferred to the conclusion? In light of the premises and the warrant of the argument—warrant in Toulmin's sense—may we take the conclusion itself as a further premise for what to believe or do? How does empirical information about the frequency of a certain pattern of reasoning, or the difficulty of abandoning it contribute to our seeing that the reasoning fails to render a conclusion acceptable in light of the premises? For pedagogical purposes, information about the frequency of a pattern of fallacious reasoning may indicate the desirability of teaching recognizing that pattern of reasoning. Students may then be forearmed against this pitfall. But does it tell us anything about *why* this "popular" pitfall is a pitfall? Woods has made a case that arguments instancing patterns traditionally identified as fallacious do not satisfy the traditional definition of a fallacy. But if arguments instancing these patterns, at least in some cases, fail on the premise acceptability or connection adequacy grounds, what is the *logical* interest of this further empirical fact? For a convincing account of the importance of his project, Woods needs to answer these questions.

References

- Alston, William. (1985) Concepts of epistemic justification. *Monist* 68 pp. 67-89.
- Alston, William. (1988) An internalist externalism. *Synthese* 74 pp. 265-83.
- Cohen. L. Jonathan. (1977) *The probable and the provable*. Oxford: The Clarendon Press.
- Cohen. L. Jonathan. (1992) *An essay on belief and acceptance*. Oxford: The Clarendon Press.
- Freeman, James B. (2005) *Acceptable premises: An epistemic approach to an informal logic problem*. Cambridge: Cambridge University Press
- Plantinga, Alvin. (1993) *Warrant and proper function*. New York: Oxford University Press.
- Pollock, John L. (1970) The structure of epistemic justification. *American Philosophical Quarterly*, 7 pp. 62-78.
- Rescher, Nicholas. (1976) *Plausible reasoning: An introduction to the theory and practice of plausibilistic inference*. Assen/Amsterdam: Van Gorcum.
- Stough, Charlotte L. (1969) *Greek skepticism: A study in epistemology*. Berkeley and Los Angeles: University of California Press.
- Walton, Douglas. (1992) *Plausible argument in everyday conversation*. Albany: State University of New York Press.
- Woods, John. (2013) *Errors of reasoning: Naturalizing the logic of inference*. London: College Publications.