American Philosophical Quarterly
Volume 42, Number 1, January 2005

PRAGMATIC AND IDEALIZED MODELS OF KNOWLEDGE AND IGNORANCE

Douglas Walton

It is a pervasive assumption in recent analytical philosophy that knowledge can be defined as a modality representing a rational agent's true and consistent beliefs. Such views are based on rationality assumptions. One is that knowledge can only consist of true propositions. This way of speaking is sharply at odds with the way we speak about knowledge, for example, in computing, where a so-called knowledge base can be a database, that is, a set of data that has been collected and is thought to consist of true propositions, even though, realistically speaking, many of them might later be shown to be false or untenable. In this study, opposed to the truth-implying idealized sense of "knowledge" there is postulated a pragmatic sense of the term in which knowledge is sometime defeated, or shown to be false as new data comes in.[1]

The idealized conception of knowledge has worked its way, through the developments of epistemic modal logic based on the abstract notion of possible worlds, into analytical philosophy. Possible worlds, introduced as a technical device in the work of Kripke (1963) on modal logic, dominated the vocabulary and methods of analytical philosophy from that point on. This device purported to provide philosophers with a powerful explanatory apparatus for explaining the notion of knowledge. Hintikka (1962) started by arguing that his epistemic modal logics represented what is meant by the concept of knowledge, but later shifted from a descriptive to a normative approach in which he argued that epistemic logic represents an idealized model of how a rational agent knows (Girle 2003: 121). Girle's comment (121) is whether such normative models are real enough or too ideal, suggesting that, through such idealization, philosophers may have been led rather uncritically down a path that has confused more than intelligently guided the study of everyday arguments, based on the notion of knowledge that philosophical analysis should to be able to analyze and understand as part of their methodology.

This paper puts forward a contrast between this idealized model of knowledge and a pragmatic model. On the latter approach, the term "knowledge" is defined in terms of a search through a database that may be more or less complete. As this dynamic process continues, more and more propositions are collected, and may be verified and falsified. Two factors are shown to be very important in judging knowledge in this framework. One is the question of whether the database is complete, in the sense that all the true propositions are known. The other is the contrast between what is known and what is not known. On this model, knowledge is defeasible, meaning that

a proposition now known may later be refuted (defeated as knowledge). It is argued that the pragmatic model enables us to make better sense of a very common form of reasoning known as the lack-of-knowledge inference, traditionally classified as a fallacy in logic, the *argumentum ad ignorantiam*.

1. The Idealized Model of Knowledge

There is a sense of the term "knowledge" that has been widely accepted in analytical philosophy in recent years, in which knowledge is defined as a species of true belief. This meaning represents a rather strict sense of the term, implying that only true propositions can be known. It follows that if a proposition was thought to be true at one time, but then was found to be false at a later time, it cannot properly be said to have been known at the earlier time. This view is, of course, a very old one. Plato liked it, because he felt that knowledge can only be of the true and unchanging reality (the eternal forms). Yet both in Plato and in recent analytical philosophy, the difficulties of the view that knowledge is justified true belief have been apparent (especially as revealed by the Gettier problem). The view that knowledge can only be of true propositions is modeled by an axiom widely accepted in formal systems of epistemic modal logic: if A is known to be true then A is true.[2] Where i represents an agent who can know propositions, and the expression K_iA reads "agent i knows proposition A," this axiom, which could be called the veridicality axiom,[3] can be formally expressed as follows.

Veridicality Axiom: $K_iA \supset A$

The veridicality axiom is a rationality assumption, and a fairly strong one. It represents what might be called a Platonic view of knowledge in which real knowledge is only of the fixed and unalterable truths. Another axiom commonly accepted in epistemic modal logics is this one: if A is a tautology then i knows that A. This axiom can also be called a rationality assumption, meaning that it could represent an artificial agent, a machine that never makes logical errors in what it knows. But doubts have been expressed whether such an axiom could represent any real agent; "Ideal agents are unreal agents, and maybe they are too unreal to be considered seriously" (Girle 2003: 108). Still, artificial agents of the kind used in AI need to carry out instructions based on logical reasoning.

Another rationality assumption characteristic of some epistemic logics is that an agent knows all the logical consequences of any statement that it knows. This assumption can be formally expressed as the following axiom of epistemic logic.

The Deductive Closure Axiom:
$K_i(A \supset B) \supset (K_iA \supset K_iB)$

Here "if A then B" is represented by the material conditional, defined as only being false if A is true and B is false. Thus the deductive closure axiom is an extremely strong rationality assumption, making the rational agent logically omniscient. The deductive closure axiom states that a rational agent knows all the logical consequences (closed under the material conditional) of all the statements it knows. Girle (2003: 110) calls this axiom "distribution," attributing it to Hintikka. This rationality assumption would be far too strong to represent the reasoning of an ordinary human agent who would have a lot of beliefs. Calculating and comparing all the logical consequences of all these beliefs would be a recursive process that would involve large computations.

There are two other rationality assumptions made in some modal epistemic logics that are also quite strong. One says that every proposition known by a rational agent is logically consistent with every other proposition it knows.

The Consistency Axiom: $\neg K_i(A \wedge \neg A)$

The consistency axiom says essentially that a rational agent never knows a contradiction. It

would appear to be equivalent to the formula Girle (2003: 111) calls the consistency principle: if i knows that A then i does not know that not A. Once again, this rationality assumption would be hard to guarantee for a normal human agent, who might be said to know a lot of propositions. In practical terms, once such a knowledge base gets quite large, it is quite likely to contain hidden inconsistencies.

Additional axioms that represent rationality assumptions about knowers and what they know concern iterated modalities like the iteration axiom.

The Iteration Axiom: $K_iA \supset K_iK_iA$

The iteration axiom rules that if an agent knows A to be true then it knows that it knows that A is true. Such iterations can be expanded to any number of agents, and are therefore highly perplexing. Whether this kind of axiom applies to a normal agent, once again, appears to be highly dubious.

The four rationality assumptions above represent notions of knowledge that may not apply to ordinary human beings, or to artificial agents of the kind found on the internet, even though they may apply to idealized reasoning of some kinds studied in science, like mathematical reasoning in Euclidean geometry, of the kind based on deductive inferences from axioms to theorems. They can be said to represent variations based on an idealized conception of knowledge. The veridicality axiom is a basic epistemic rationality assumption, excluding all false propositions from knowledge. The other axioms can be added on, one at a time, representing stricter and stricter rationality assumptions on the concept of knowledge so defined. When you put them all together, you get a very strict and idealized conception of knowledge. An agent that has knowledge in this sense is such an idealized knower that one might well question whether any real agents of this sort exist. Still, an agent to which all four axioms apply can represent an ideal standard of rationality.

2. The Inquiry as a Model of Proof

Corresponding to the idealized model of knowledge, there is dynamic model of the process by which knowledge is acquired and verified, based on arguments from premise to conclusions, called the inquiry. The inquiry represents an epistemological view called foundationalism. On this view, knowledge is built up by inferences only from premises that are established and verified as axioms that are not subject to doubt—on concrete foundations, so to speak. Only from such solidly established premises can a conclusion be drawn that represents knowledge. In this model of a knowledge collection and verification procedure, only a proposition that has been proved can be called knowledge, and proving it means proving it beyond doubt. Thus once established, a proposition that has been proved in this manner should never need to be retracted. Of course, retraction is common in the kinds of searches and proofs we carry out in everyday reasoning. And it is not unknown even in scientific reasoning. This observation suggests that the inquiry may not be the only model of searching for data to prove something.

There can be various reasons for searching for information. One may be to collect information that is needed to make intelligent decisions when deliberating on what to do. For example, all the information telling someone at an airport monitor which flights leave to which destinations at what times is presented so that such a person can make informed decisions about which gate to go to in order to board a flight, and so forth. The information is collected and displayed to help such deliberations. In other cases, information (data) may be collected in order to prove or disprove a hypothesis in a scientific investigation. The purpose of an inquiry is to prove or disprove a hypothesis by collecting relevant evidence. In general, the goal of an inquiry is to prove that a particular proposition is definitely true

or false, based on premises that are known to be true. In some cases, what may be proved is that the proposition in question cannot be proved, when all the relevant evidence that might prove it has been collected. The aim of the inquiry is to establish knowledge by drawing conclusions only from premises that are known to be true or false. The aim is to base the inquiry on firm foundations, so there will not be a need in the future to have to retract conclusions that were thought to be true but turned out later to be false. The goal of the inquiry is to have argumentation that is *cumulative,* meaning that once a particular proposition is accepted as true at any particular point in the inquiry it will remain and hold at all successive points as the inquiry proceeds. Cumulativeness is an ideal of the inquiry as a normative model of how the collection of knowledge should proceed. In real cases of that we might call inquiries, retractions occur. However, from the point of view of inquiry as a model of argumentation, a retraction is regarded as evidence of a "mistake." For example, when retractions of published findings of scientific research do occur, they are generally an occasion of discomfort, indicating retraction in an inquiry is something to be avoided if possible (Broad and Wade 1982: 181–192).

Aristotle wrote (*Posterior Analytics*, 71 b 26) that in a demonstration (*apodeixis*), the premises must be "primary and indemonstrable," so that the premises are better known than the conclusion and prior to it (71 b 29). The model of a demonstration is Euclidean geometry. The premises are axioms laid down as starting points and all theorems are derived from the axioms, by deductively valid rules of inferences that are numbered. As each new theorem is derived from axioms or previously proved theorems it is assigned a higher number. Since each theorem follows deductively from fixed axioms or previously proved theorems, there is never any need to retract any conclusion drawn in the sequence of proofs. This model represents the inquiry very well.

Does the establishment of scientific knowledge take the form of an inquiry of the kind described above? Controversies on the issue have raged in recent years in the so-called "culture wars." It seems very doubtful, according to the research in sociology on how scientific argumentation actually takes place, that it takes the form of the inquiry. The scientists form groups, follow paradigms, and defend their favorite models against those of their opponents. Scientific opinions are often divided. New theories or findings are often ridiculed at first, then later settle into orthodoxy. None of this seems much like the model of an inquiry in which results are firmly fixed and never later retracted. On the other hand, scientists often invoke the model of inquiry as an ideal of scientific research presented to their students or to the general public. Whether this ideal represents what really happens certainly has been contested. But it needs to be recognized that there are stages of scientific inquiry. At an earlier discovery stage, a hypothesis may be a creative guess pointing the way to experimentation and testing. At a later stage, a hypothesis may be very well tested, and be accepted by the majority of scientists in a field, even appearing in all the elementary textbooks. At the early stage, a problem for investigation may be identified, even before the stage of collecting data proceeds. There is a later stage of testing, and a still later stage of drawing conclusions from the prior process. Still, if you look at this broad process as a whole, especially as it reaches the later stages, it may be possible to discern a cumulative buildup of evidence as scientific knowledge is consolidated.

The inquiry is a highly idealized model of knowledge, especially in regard to the property of cumulativeness. Does it make sense to talk of defeasible knowledge of the kind that is possible only in a non-cumulative framework? The suggestion may make no sense to

many philosophers, especially those of a Platonic bent. If a proposition was refuted, even if it was accepted as knowledge at any earlier time, surely it makes no sense to say that it is knowledge, or that it was known to be true at that earlier time. On the other hand, we often take things to be "knowledge" that are later on rejected, or shown to be false or untenable as scientifically accepted opinions. Can any sense be made of such a suggestion?

3. A Pragmatic Model of Knowledge

It can be argued that agents of a more realistic kind often have to function under conditions of uncertainty and lack of knowledge. Even in science, old theories and hypotheses give way to new ones as scientific knowledge advances. What we said we knew at one time falls away as an old theory or hypothesis is no longer accepted. Does that mean we did not really know it, at the earlier time? Certainly it does under the idealized standard of epistemic rationality. On that standard, unless the proposition in question is true, it cannot be knowledge. Thus we are led to ask whether there could be an alternative conception of knowledge according to which we could say that propositions currently accepted in scientific theories, or as scientific findings supported by the evidence and generally accepted by scientists in a domain of knowledge, are known to be true. Here knowledge is defeasible, meaning that future evidence that is discovered could turn out to defeat the hypothesis or theory. Such a practical conception of knowledge, opposed to the idealized one outlined above, can be modeled as follows.

In the pragmatic model of epistemic rationality, there is a group of agents engaged in a search for the truth of a matter being investigated. They are collecting data and attempting to verify or falsify hypotheses by testing them against the data collected. But the data they are collecting is incomplete. They have collected a set of data to this point in the search, and they have falsified some hypotheses, based on these data, and verified others. But as the search continues, some of the hypotheses verified at this point will later on turn out to be falsified. Even so, they can describe the propositions they accept at this point in the search, including data statements as well as hypotheses, as representing the knowledge that they currently possess. The pragmatic model of epistemic rationality is dynamic. None of the propositions called knowledge at this point are known to be false. But they might turn out to be false. If shown to be false, they are no longer knowledge. But still, a proposition rightly described as known could turn out to be false. Also, knowledge is not necessarily consistent. A new hypothesis could be inconsistent with an old one, or the set of data could justify either one of two hypotheses, one of which is inconsistent with the other. On this model, something can be correctly described as knowledge at one time in the search, but then at a later time, it is no longer knowledge.

The pragmatic model sees knowledge as a matter of evidence based on practical acceptance rather than on abstract and idealized notions of truth and absolute proof. It sees knowledge as based on acceptance during a dynamic process of collecting evidence and testing it against appearances (data). It is a skeptical view that traces back to Carneades' criterion of evidence. A proposition is evident that appears to be true, but is more evident if it tested against appearances and passes the test. It is still more evident if it is consistent with other propositions that are accepted.[4] It can be argued that this acceptance-based view of evidence was the forerunner of the view of knowledge advocated by the American pragmatists (Doty 1986). Evidence is what is tentatively accepted, based on what appears to be true, at a given point during the collection of evidence in a particular case. According to the pragmatic view, evidence is basically

something that seems to be true in the form of a plausible argument that gives a good but defeasible reason to support a conclusion that is doubtful. This view is an alternative to the justified belief model of evidence. On the pragmatic view, knowledge is no longer defined as justified true belief, or as any kind of belief. It is based on argument, based in turn on rational acceptance. To say that a proposition is knowledge, on this view, is to say that it is not only based on evidence, but also to say that it can be proved, meaning that it is established as accepted in a domain of knowledge, up to the standards of burden proof appropriate for and accepted in that domain.

The pragmatic account of knowledge starts with a knowledge system, meaning a working system with an agent having a database. It is assumed that we have a basic idea of the principles of how it works and we can offer a description of its operation, its methods, and its standards for acceptance. It has a methodology for collecting data, for testing the data, and then for accepting or rejecting hypotheses based on such testing procedures and on what is currently accepted or not in the field as knowledge.

The basic idea behind the pragmatic model is the notion of the search. A search is a collection of data that meets some criterion or standard of what is to be collected in order to answer a question. For example, I ask whether a certain file labeled "Search" is in the folder "My Documents." I scan through "My Documents" visually. I see no file in that folder labeled "Search," and so I conclude that there is no "Search" file in "My Documents." This is a negative finding, but it could justify a definite conclusion that there is no "Search" file in "My Documents" if my search has been thorough enough. This negative finding could be significant, because it could lead me to search elsewhere for this file. A search of this kind is conducted by an agent. An agent is an entity that has some information or knowledge and can use it to take action to carry out a goal. Moreover, an agent can know what some of the consequences of its actions are, and so can use this feedback to guide its actions further along the path to its goal. The notion of a search can be a social one, in many instances. A group of agents can be engaged in the search, and some of them can know things that others do not know.

On the pragmatic model, a search can be judged to have been carried out to various depths, and there are practical criteria for how deep the search has been. Take the example of a security search. A government employee may have to undergo a background search for criminal convictions and so forth, as a condition of employment. But someone who applies for a job at the FBI may have to undergo a much more thorough search. Such a search may be very deep, in the case of a suspected spy for example, but even an extremely deep search may not be successful in finding what is there, but cannot be found. For example, a spy may be a "mole," but no search may uncover him, if all the incriminating data has been successfully concealed. These observations raise the question of what the standards are for judging the success or failure of a search for knowledge according to the practical model of epistemic rationality. The first thing to note is that a knowledge base can be judged complete.

The meaning of the term "complete" is that all the knowledge in the domain of knowledge has been collected in the search. The condition that a database is complete in this sense is called the closed world assumption in computing (Reiter 1980: 69), meaning that all the knowledge one can collect in the database is included in it. Reiter (1980: 69) offered the following example of a question-answering system to show how the closed world assumption is applied to a database in artificial intelligence: Consider a database representing an airline flight schedule and the query "Does Air Canada flight 113 connect Vancouver

with New York?" The question-answering system searches through its database in an attempt to prove the proposition CONNECT (AC113, VAN, NY). If the proof succeeds, the system answers "yes." Otherwise it answers "no." The problem is whether the database is complete. Does it contain all the possible flights connecting Vancouver with New York? In a typical airport database, all available flights to or from that airport would be listed, and so the closed world assumption is met.

This meaning of "complete" is pragmatic. It means that all the knowledge available at this point has been collected in the database. It does not mean that this set of propositions is fixed for all time. For example, new flights may be added at an airport, or old ones may be cancelled. Still, if all the flights currently available are listed, the database is complete, in the pragmatic sense of the term cited above. In practical matters of security searches and airport flight listings, the closed world assumption can be met, even if there is no guarantee that the veridicality axiom is met. The consistency axiom would be assumed, since a correct airport listing of flights, for example, would be presumed to be consistent.

On the pragmatic model, completeness is determined by setting a burden of proof appropriate for the kind of investigation, search, or inquiry that is underway. Burden of proof has two factors, comprising how high a level of strength of argument is needed for proof, and which party or parties to the inquiry has to furnish the proof. The inquiry is an idealized investigation procedure, as it never admits of retractions, as shown above. Thus to represent a pragmatic notion of knowledge that will fit various frameworks like that of a forensic investigation into an air disaster, or a security investigation for clearance to access sensitive documents, models of search with conditions for success weaker than that of the investigation need to be studied. For example, a search may be ruled complete if it has run through designated procedures that are thought to have exhausted a database thoroughly enough to move ahead with a decision for action, subject to retraction of the finding if new data comes to be known.

4. *Ad Ignorantiam* Arguments

The *argumentum ad ignorantiam* is an argument of the form: proposition *A* is not known (proved, established) to be true (false), therefore *A* is false (true). Essentially the same type of argumentation is called "negative evidence" in science, and the *ex silentio* argument in history. A good example of the latter is the historical conclusion that the Romans did not award medals posthumously, inferred from the lack of historical evidence that a posthumous award ever took place (Maxfield 1981: 138). Arguments from ignorance are often used as means of making a tentative conjecture under conditions of uncertainty and lack of so-called hard evidence. For example, in an article on pursuing the elusive Osama Bin Laden (*Newsweek*, August 19, 2002, 35–41), the following *ad ignorantiam* argument was used: "One apparent sign Bin Laden is not dead is the relative lack of background chatter picked up on radio and other electronic transmissions." This *ad ignorantiam* argument, like the one about Roman medals, is a reasonable one, even though it is only a less-than-conclusive conjecture based on lack of evidence.

In principle, the argument from ignorance is a reasonable kind of default argument that can rightly be used to shift a weight of presumption in a balance of considerations argument, based on burden of proof, but it depends on a premise indicating depth of search. For example, if no evidence has been found that Ed is a spy, one could reasonably conclude by *argumentum ad ignorantiam* that he is not a spy, but only on the assumption that a thorough security search into Ed's activities has been carried out. Similarly, if a large number of rats have been given a certain drug,

and did not experience any harmful effects of a kind that were anticipated, this negative outcome may have some worth as evidence for the conclusion that ingestion of this drug does not produce this outcome. Negative evidence is generally regarded as less significant than positive evidence, in scientific research. However, some researchers see this priority given to positive over negative evidence as a kind of bias inherent in current methods of scientific research (Witte, Kerwin, and Witte 1991). It has been recognized as a bias that affects statistical findings on the outcomes of scientific research, and that ought to be corrected, or brought into less of an imbalance. Although negative evidence is typically of a tentative sort, depending on how thorough a search was, it can be highly significant in reaching a rational decision on taking action under conditions of uncertainty.

The argument from ignorance is quite familiar in computer science, where it takes the form of negative reasoning from a knowledge base. Reiter (1987: 150) comments on the example of an airline fight schedule presenting a list of flight numbers and pairs of cities they connect. He adds, "We certainly would not want to include in this data base all flights and the city pairs they do *not* connect, which clearly would be an overwhelming amount of information." But by negative reasoning from this knowledge base, a viewer of the flight schedule listed on an airport monitor can infer that, if two cities are not stated as connected in the list given, there is no flight connecting them. This kind of negative reasoning from a given knowledge base is highly familiar in computer science, and is a reasonable species of *argumentum ad ignorantiam*. Such arguments are frequently associated with default inferences and nonmonotonic reasoning, of a kind where lack of knowledge warrants the drawing of a presumptive inference on how best to proceed in a situation of uncertainty. More and more, it is being recognized in computer science, that this kind of reasoning is commonplace and useful, and is generally reasonable, as opposed to being fallacious.

A very simple example adapted from Walton (1996: 251) can be used to show to the reader how common the argument from ignorance is, in everyday reasoning, and how it can function as a reasonable kind of argumentation. Suppose Wilma cannot find her pen, and she asks Bruce, "Is it in the desk?" Bruce replies, "I do not think so, because I have a pretty good idea what's in the desk," and then Wilma continues, "Yes, but have you looked in the desk?" In this small dialogue is contained the basic idea of how the argument from ignorance works. When Bruce replied that he did not think the pen was in the desk, he was drawing an inference on the basis of the argument from ignorance. He was indicating that, as far as he knew, the pen was not in the desk, and then from this premise of lack of positive knowledge, he was drawing the conclusion that (probably, or plausibly), the pen was not in the desk. Wilma's reply questioned how thorough Bruce's search was, perhaps suggesting that he should actually go and look through the desk, to make sure the pen was not there. She was questioning the strength of his argument from ignorance by questioning the depth-of-search premise that is always an additional (in this case nonexplicit) premise in the argument from ignorance.

In general, the argument from ignorance has two premises, as displayed in the form of argumentation scheme below (Walton 1996: 254).

Argumentation Scheme for Argument from Ignorance

Lack of Knowledge Premise: Proposition *A* is not known to be true (false).

Conditional Premise: If *A* were true (false), then *A* would be known to be true (false).

Conclusion: Therefore *A* is false (true).

In the pen case, the depth-of-search premise equates to the conditional premise above. Wilma was asking Bruce, in effect, whether if

the pen were in the desk, he would know that. So generally, we can see how the argument from ignorance is based on a characteristic argumentation scheme that shows its structure as a kind of reasoning that can be correct or reasonable, as used in many cases in everyday reasoning.

Here we have to go into the question of when we are actually justified in determining that a knowledge base is closed. In terms of the argumentation scheme for the argument from ignorance, this amounts to a determination of when the conditional premise is acceptable. On the pragmatic theory, this question is a matter of burden of proof, determined by the accepted standards and methods of the domain of knowledge. Whether this proposition can be taken as known or proved is to be determined by different standards in different domains. For example, proof in a court of law is decided by different standards and methods than proof in a domain of science. In the case of the airline database we can be justified in saying that the database is complete if we know that the system kept up to date on a constant updating of flights.

The question is then raised: if the argument from ignorance is correct in some form, what kind of correct reasoning does it represent? It seems to be an epistemic form of reasoning based on assumptions about what is known and not known in a case. But the kinds of cases of everyday use of such arguments examined above do not seem to meet any of the idealized assumptions about knowledge expressed by the epistemic axioms of modal set out above. Nor do they fit the framework of the inquiry as a model of establishing knowledge. Instead, they represent everyday examples of reasoning in which an agent searches for something, finds it or not, and then draws an appropriate conclusion. Models of search weaker than that of the inquiry are therefore necessary to fill out the pragmatic model of knowledge fitting common cases of defeasible knowledge-based reasoning.

5. Fallacious Arguments from Ignorance

Like a lot of arguments traditionally classified as fallacies, the argument from ignorance is frequently a weak, presumptive sort of argument that is inconclusive, but shifts a weight of presumption to one side or the other in a dialogue. As such, it is often used as a way of shifting the burden of proof in an argument. Sometimes this back-and-forth process leads to a kind of situation called the *ad ignorantiam* "tug-of-war" (Walton 1996: 118), where the following case was cited. In a debate in the Canadian House of Commons, the issue was Opposition's concern that the embargo on the export of Canadian uranium "for non-peaceful purposes" was not being respected. An opposition minister demanded that the Secretary of State for External Affairs prove that the treaty was being respected, after he had claimed that, as far as he knew, on the information that was available, that it was being respected. The opposition minister asked, "What is your proof?"(Walton 1996: 119). The Secretary of State replied, "I have looked for any weakness in the treaty, and I have found none." He told the Opposition not to be so secretive, "Come forward with your allegations so that we can find out whether they are true or false"(118). The reply was, "Do a proper investigation." In this case, each side tried to shift the burden of proof back to the other side, in a typical *ad ignorantiam* tug-of-war. The problem, in such a case, is to determine on which side the burden of proof should rightly lie in the debate. In cases where it has not been decided, an *ad ignorantiam* argument can go back in forth in this fashion through many moves.

The fallacious kind of *argumentum ad ignorantiam* arises where one party in a dialogue tries to use this kind of argumentation as a tactic to force the burden of proof or disproof around the other way. By such a tactic, the respondent's ability to prove anything,

or even to raise appropriate critical questions in the dialogue, is blocked or impeded. Such cases may occur even in cases where the burden of proof has been clearly set on the proponent's side. At the opening stage of the dialogue, the proponent may try to avoid fulfilling the burden by trying to shift it to the other side, making it appear as though the other side must prove its claim (which may not even be possible), or lose the debate. The most visible kind of case of this sort is the "kangaroo court" or so-called "witch hunt" type of tribunal, which has the trappings of a fair criminal trial, but where the defendant is presumed to be guilty. The classic cases cited here are the witchcraft trials, especially the Salem witchcraft trials of 1692, and the McCarthy tribunals of the 1950s. The Salem witchcraft trials were actually legal proceedings, but the so-called "spectral evidence" claimed by witnesses to prove that someone was a witch (in league with the devil), was difficult to disprove, because it was supposedly only visible to the person testifying. The McCarthy tribunals were televised proceedings that were made to look to viewers a lot like trials, but they were not legal trials. McCarthy posed as having evidence that the defendant was a "communist sympathizer," but the allegations were more often based on innuendo. In one case, cited by Copi (1982: 112), McCarthy stated, "I do not have much information on this except the general statement of the agency that there is nothing in the files to disprove his Communist connections." One problem here for the person accused of being a Communist sympathizer was that it was extremely difficult to prove such a negative thesis—that he is not a Communist sympathizer—because of the vagueness of the charge, and the smear on his trustworthiness created merely by the accusation itself. But the main problem was the reversal of the burden of proof. The accused was put in the position of having to try to fulfill this heavy, or even impossible, burden of proof.

How is the fallacy in such a case a misattribution of commitment? It seems that the fallacy is an epistemic failure in which lack of knowledge is improperly treated as evidence of a kind that should be based on knowledge and not mere speculation or supposition. But the fallacy is also one of failure of due process in which there has been an illicit shift in the burden of proof. The defendant must prove he is innocent, but this is impossible because the tribunal is not open to such an outcome. Both sides in a dispute must have commitment to due process. If one shows evidence of a failure to have this kind of commitment, the other can attack him as an unfair and biased arguer.

6. Conclusions

The idealized model of knowledge cannot deal very effectively with the argument from ignorance, or lack-of-evidence argument as it might less prejudicially be called. It makes this form of argument come out as inherently fallacious rather than as an argument that is reasonable in some cases and fallacious in others. It can model the lack of knowledge inference as a form of rational argument, but it can only deal with cases based on the inquiry model, where retraction is not allowed. As shown by the examples above, many of the cases we have to deal with and evaluate, both in everyday and scientific reasoning, are not of this kind. They are defeasible arguments based on a search through a database that is not complete. They are arguments that are intelligent guesses, used to guide a search forward as more evidence is being collected. The evaluation of epistemic arguments in such cases is best seen as pragmatic, based on an allocation of burden of proof that can move the search process forward in a constructive direction. Thus, by such a search, a conclusion is reached tentatively, even if it may need to be retracted later, as more data is found.

None of this is to say that the idealized model of knowledge is wrong, inherently incorrect, or misguided as a philosophical viewpoint. It is to suggest that when this approach is paired with the pragmatic model, a more useful and balanced epistemology can be produced. The immediate payoff of bringing the pragmatic model alongside the idealized one is that much better sense can be made of the argument from ignorance as a form of argument that can often be reasonable, even if it is defeasible. This move, in turn, offers a better basis for dealing with cases where a decision needs to be arrived at on the question of whether a given argument from ignorance is reasonable or fallacious.

University of Winnipeg

NOTES

1. The author would like to thank the Social Sciences and Humanities Research Council of Canada for a grant, and David Godden for many vitally important comments and corrections.

2. Capital letters *A*, *B*, etc., refer to statements, also called propositions, entities that are true or false.

3. Girle (2003: 110) calls it the veridicality principle.

4. According to Sextus Empiricus, *Outlines of Pyrrhonism* (Mates 1996: 122), something is plausible if it appears to be true. It is even more plausible if it appears to be true and is consistent with other things that appear to be true. And it is even more plausible if it is stable, meaning that it is consistent with other things that appear to be true, and is tested. According to this pragmatic theory, everything we accept as based on evidence is subject to doubt and is plausible only, as opposed to being known (beyond all reasonable doubt) to be true.

REFERENCES

Aristotle. 1984. *Prior Analytics. The Complete Works of Aristotle*, ed. Jonathan Barnes (Princeton, N.J.: Princeton University Press).

Broad, William, and Nicholas Wade. 1982. *Betrayers of the Truth: Fraud and Deceit in the Halls of Science* (New York: Simon & Schuster).

Copi, Irving M. 1982. *Introduction to Logic,* 6th ed. (New York, Macmillan).

Doty, Ralph. 1986. "Carneades, a Forerunner of William James's Pragmatism." *Journal of the History of Ideas*, vol. 47, pp. 133–138.

Girle, Rod. 2003. *Possible Worlds* (Chesham: Acumen Publishing Limited).

Hintikka, Jaakko. 1962. *Knowledge and Belief: An Introduction to the Logic of the Two Notions* (Ithaca, N.Y.: Cornell University Press).

Kripke, Saul. 1963. "Semantical Analysis of Modal Logic I: Normal Propositional Calculi." *Zeitschrift fur Mathematische Logic und Grundlagen der Mathematik*, vol. 9, pp. 67–96.

Mates, Benson, trans. 1996. *The Skeptic Way: Sextus Empiricus's Outlines of Pyrrhonism* (New York: Oxford University Press).

Maxfield, Valerie A. 1981. *The Military Decorations of the Roman Army* (Berkeley: University of California Press).

Reiter, Raymond. 1980. "A Logic for Default Reasoning." *Artificial Intelligence*, vol. 13, pp. 81–132.

Reiter, Raymond. 1987. "Nonmonotonic Reasoning." *Annual Review of Computer Science*, vol. 2, pp. 147–186.

Walton, Douglas. 1996. *Arguments from Ignorance* (University Park: Pennsylvania State University Press).

Witte, Charles L., Ann Kerwin, and Marlys H. Witte. 1991. "On the Importance of Ignorance in Medical Practice and Education." *Interdisciplinary Science Reviews*, vol. 16, pp. 295–298.