

RULES FOR REASONING FROM KNOWLEDGE AND LACK OF KNOWLEDGE

The modern logic textbooks typically classify the *argumentum ad ignorantiam* or argument from ignorance as a fallacy, although many of them increasingly hasten to add that in some cases this form of reasoning can be nonfallacious. In the search of 240 logic textbooks and critical thinking manuals reported in (Walton, 1996, p. 25), 55 were found to have something to say about argument from ignorance. Most of them only had a very short section on this type of argument, and dismissed it as fallacious using examples of problematic arguments concerning topics like unidentified flying objects, extrasensory perception, and the existence of God. For example in Copi (1982, p. 101), the fallacy of *argumentum ad ignorantiam* is illustrated by the argument that there must be ghosts because no one has ever been able to prove that there aren't any. A problem with this approach is that outside the teaching in the current logic textbooks, argument from lack of knowledge, which appears to be essentially the same as argument from ignorance, has often been recognized as a legitimate, reasonable, and very common way of drawing conclusions from a knowledge base, for example in computer science.

It is typical of knowledge-based systems of the kind currently used in artificial intelligence that the system searches through a knowledge base drawing inferences based on rules (called forms of argument in logic, like *modus ponens*) and then once the search has been completed, draws a conclusion (Russell and Norvig, 1995, p. 151). But in the vast majority of cases of such searches of the kind found in realistic arguments, for practical reasons, like the costs or time requires for continued searching, a conclusion may have to be drawn even when it is known that the knowledge base is incomplete. The facts are not all in yet. We do not have what is called epistemic closure (Reiter, 1987), meaning that the database can be assumed to be complete, containing all the knowledge that exists in a domain and that might be collected by continued searching. The tool used in computing systems to ensure closure of a system to justify drawing a conclusion from the search is called the closed world assumption, also known as negation by failure. This is an assumption that allows the user of the system to draw the inference that any proposition not specifically known to be true can be taken to the false: "for example, if the system uses a database of facts, and a particular fact is not included in the database, then that fact is assumed to be false" (Coppin, 2004, p. 480).¹ This kind of reasoning looks very much like it fits the description of argument from ignorance, or lack of knowledge, but it is not regarded as fallacious.

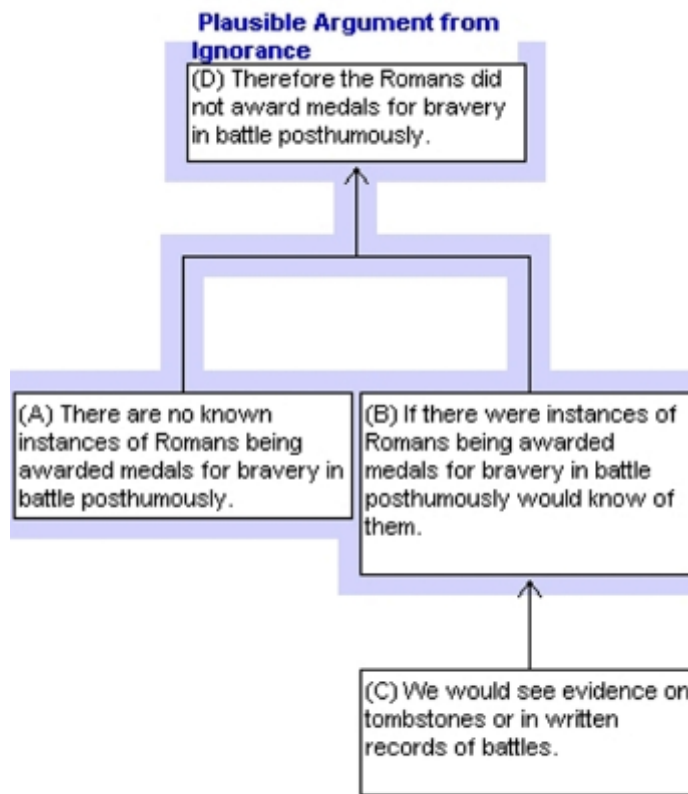
In this paper, I begin by examining some common cases of the *argumentum ad ignorantiam*. I then proceed to try to explicate how lack of knowledge reasoning could be defined by examining current standards for defining knowledge in logic, and by formulating four axioms of knowledge commonly found in current systems of epistemic logic that fit the current view. Next, I go on to examine some realistic cases of lack of knowledge reasoning that are hard or even impossible to evaluate fairly and accurately by these currently accepted axioms defining knowledge. These cases are characterized by reasoning based on lack of knowledge, and they tell us a lot about how we need to rethink what knowledge is and rational standards for evaluating knowledge-based arguments.

¹ The most famous example of a system that uses the closed world assumption is PROLOG, an artificial intelligence system that has negation as failure.

1. *Argumentum Ad Ignorantiam*

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, p. 138). The structure of the argument can be visualized as an argument diagram.²

Figure 1: Argument Diagram for the Roman Medals Example



The two premises A and B function together in a linked argument that supports conclusion D. C offers evidence to support the conditional statement B. The argumentation scheme by which A and B are linked to support conclusion D is displayed on the diagram – plausible argument from ignorance. This scheme will be presented in section 3 below. It is a presumptive scheme, in this instance, that allows a conclusion to be drawn tentatively on the basis of what is known and what is not. The *Araucaria* visualization tool for constructing an argument diagram of an example of argumentation

² The argument diagram above was constructed using the diagramming tool *Araucaria* (Reed and Rowe, 2002). It can be obtained free at this site: <http://www.computing.dundee.ac.uk/staff/creed/araucaria/>

found in a text of discourse in a given case enables the scheme to be applied to the argument as presented in the case. This application reveals the structure of the argument, displaying how the components of this scheme match the propositions given as premises and conclusions in the argument in the text. Such an application has many uses, both in computing and in the educational task of teaching critical thinking to students, like the identification of missing premises (Rowe, Macagno, Reed and Walton, 2006). In section 3 below, it will be shown how *Araucaria* matches the scheme to the given argument.

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. If new evidence comes in, the argument may default. That is the reason the argumentation scheme selected to represent the structure of the argument shown in figure 1 is that of plausible argument from ignorance, as opposed to deductive argument from ignorance. The latter is (at least arguably³) a deductively valid form of argument.

As the examples above suggest, plausible argument from ignorance can sometimes be a reasonable kind of argument even if it is defeasible, and thus can default as more information comes to be known in a case. Thus even though it is not, in many instances, a conclusive argument, it can still be a basis for drawing a conclusion by rational argumentation. It can be used to shift a weight of presumption toward one side in a balance of considerations on a controversial issue or continuing search for knowledge. Its acceptance should be based on burden of proof, and the strength of the argument depends on depth of search. For example, suppose that even after a thorough security investigation, no evidence has been found that Ed is a spy. One might reasonably draw the conclusion by *argumentum ad ignorantiam* that he is not a spy, at least as far as the evidence so far indicates. The strength of the argument depends on how thorough the security search into Ed’s activities was. Or to cite a similar, and also familiar kind of case, suppose a selected sample of rats have been injected with a certain drug, but no harmful consequences were observed when they were tested. Such negative evidence presumably has some strength as evidence supporting the conclusion that the drug is safe (Witte, Kerwin and Witte, 1991). Still, negative evidence of this sort has been generally regarded in scientific research as a less significant experimental finding than positive evidence, even though some see this nonequivalence as a bias that is rationally unjustified (Sterling, Rosenbaum and Weinkam, 1995). Thus although the question of how negative evidence should be weighed is controversial, it is generally accepted that it should have some value as evidence in many instances in scientific research.

Argument from ignorance is a form of reasoning that has gained some recognition and acceptance in computing, as a form of knowledge-based reasoning. Reiter (1987, p. 150) offered the example of an airline flight schedule showing a list of flight numbers paired with of cities connected by each flight. It would not be helpful to list all the pairs not connected by a flight. As Reiter put it (p. 150), “We certainly would not want to include in this data base all flights and the city pairs they do *not* connect, which clearly

³ See section 7.

would be an overwhelming amount of information.” Instead the person at the airport consulting the list can use argument from ignorance to draw the conclusion that if two cities are not stated as connected in the list given, there is no flight connecting them. Here the function of the argument from ignorance is to enable the user to draw conclusions from information not presented in a knowledge base. If a proposition is not explicitly included in the knowledge base, the user can draw the conclusion that it is false. As noted at the beginning of this paper, the principle used to enable a system to draw such a conclusion is called the closed world assumption in computing (Coppin, 2004, p. 80).

Another simple example of a comparable sort (Walton, 1996, p. 251) shows how argument from ignorance should be judged as strong or weak depending on evidence. Suppose Wilma can't find her pen, and she asks Bruce, "Is it in the desk?" Bruce replies, "I don't think so, because I have a pretty good idea what's in the desk", and then Wilma continues, "Yes, but did you search the desk?" Bruce replies "Yes". Wilma then asks him, "How thoroughly did you search through the desk?" When Bruce first replied that he didn't think the pen was in the desk, he was drawing his conclusion on the basis of an argument from ignorance. His conclusion was that, as far as he knew, the pen was not in the desk. But this was quite a weak argument, because as Wilma's next question showed, he hadn't looked in the desk, or at least not recently or thoroughly. She doubted the strength of his argument from ignorance by questioning what can be called the depth-of-search premise. If Bruce were to then search through the whole desk very thoroughly, and not find the pen, his argument from ignorance would now be much stronger.

The question is then raised: if the argument from ignorance can be correct in some instances as represented by the forms above, what kind of correct reasoning does it represent? It seems to be an epistemic form of reasoning based on assumptions about what is not and not known in a case. Its strength or weakness as an argument in a given case depends on how thorough the search for knowledge was in that case. As the search continues and more evidence comes in, the argument from ignorance can become stronger. If there was very little evidence, because not much of a search had yet been undertaken, the argument could be very weak. But still, it might not be entirely worthless. For one thing, it might be a basis for deciding whether to continue searching.

Another question that needs to be discussed is that of whether the argument from ignorance should be defined as fallacious, or should be seen as a kind of argument that can sometimes be reasonable and sometimes fallacious. This form of argument has often been recognized in many disciplines outside logic. For example in computing it is called a lack of knowledge inference, defined as an inference where an agent fails to find a query item of information that a certain proposition is true, stored in its knowledge base, and concludes that this proposition must be false (Collins, Warnock, Aiello and Miller, 1975). In history this type of reasoning is called the *ex silentio* argument - see the Roman medals example below. In science this form of inference is commonly called reasoning from negative evidence, as noted above. Computing systems based on what is called autoepistemic reasoning (Konolige, 1988) have been designed to enable an agent to draw inferences from its own lack of knowledge about a proposition. In all these instances of the form of negative reasoning involved is seen as reasonable, not fallacious. Whereas in logic what appears to be the same kind of reasoning, judging from the examples in the logic textbooks, has traditionally been classified as the *ad ignorantiam* fallacy.

Should the expression ‘argument from ignorance’ be reserved for the fallacious instances of arguments from lack of evidence, or should this expression be regarded as equivalent to arguments from lack of evidence, also called arguments from lack of knowledge, and by many other more positive-sounding names? The traditional expression in logic, ‘argument from ignorance’ has a negative sound to it. We generally judge ignorance to be a bad thing. Therefore it may be suggested that the expression ‘argument from ignorance’ should be reserved for the fallacious type of inference from lack of evidence. But perhaps ignorance is not completely a bad thing, in all instances. For example, Socrates described his awareness of his lack of knowledge as a kind of wisdom. For these reasons, I would say that it is better to use the expression ‘argument from ignorance’ in such a way that it can represent the reasonable kinds of inferences based on lack of knowledge as well as the fallacious ones. Thus the job for the study of informal fallacies is that of devising criteria that will enable us to judge in a given case whether a given argument from ignorance is fallacious or reasonable.

An objection is that it is very hard to see that there can be nonfallacious arguments from ignorance, given that the examples of lack of evidence arguments cited above are partly based on positive evidence as well as negative evidence. In the Roman medals example, the arguers agree that there are no records of awarding medals posthumously. The premise is that there are no historical records of awarding medals posthumously. But this can be described as positive knowledge, for it cannot be characterized exclusively as not knowing or not being able to prove that a proposition is the case. The same applies to the Bin Laden example. We know that if Bin Laden was dead, there would be much chatter picked up on radio transmissions. Since we now observe that there is hardly any chatter of this kind, we can conclude that probably Bin Laden is not dead. It can be argued that this case should not be classified as an argument from ignorance, because it does not conform to the scheme that a particular proposition is not known to be true, therefore it is false. We might conclude from these observations that it is impossible that there could ever be nonfallacious arguments from ignorance.

The truth is, however, that such arguments never take place in a vacuum. They always take place during the context of an investigation in which there is a database that has already been collected. In other words there is already a body of evidence comprised of what might be described as positive knowledge. Then the particular proposition that is said to be unproved, or not known to be true, is utilized as the specific premise to draw the conclusion that this particular proposition must be false. It does not follow in such cases that the argument is exclusively based on ignorance, or lack of knowledge. Such examples are based on a combination of ignorance with what is taken to be knowledge, or at least a body of evidence collected in an investigation at some particular point as the investigation proceeds. Thus when we turn to analyzing the form of the argument from ignorance in the next section, the argumentation schemes characterizing it are defined partly in terms of ignorance, but also partly in terms of knowledge that is assumed to have been collected as an investigation proceeds. As it proceeds, assumptions need to be made about both what is known and what is not known at a particular point.

2. Four Axioms of Knowledge as a Modality for Rational Agents

In order to approach the problem of how to identify analyze and evaluate arguments based on ignorance, or lack of knowledge, we need to first try to get some grasp of what knowledge is taken to be according to the current conception of it in epistemology. According to the current idealized conception of knowledge, based on abstract systems of modal logic representing truth in possible worlds, knowledge is taken to be represented as a body of true propositions that scientific research aspires to discovering and proving. According to this conception, knowledge consists of a set of propositions that are consistent with each other, that are true, and that will never be shown to be false as scientific research proceeds in the future. According to this model, knowledge is separate from the agents or investigators who search after it, for example in scientific research. Some have suggested that this idealistic model of knowledge can be contrasted with a more pragmatic model in which knowledge is defined in relation to an investigation in which propositions originally taken as known to be true can be shown later in the investigation to turn out to be false (Walton, 2005). But before examining what the pragmatic alternative is, let us carefully defined some of the key rationality assumptions in the current idealistic model of knowledge.

According to the standard way of representing knowledge in analytical philosophy, the expression ‘ a knows that A ’, for agent a , and proposition A , can be treated as a modality, a way a proposition can be true or false. Once it is treated in this way, it is also tempting to makes several idealized assumptions about it that correspond to familiar axioms of modal logic. One of these assumptions is that knowledge can only consist of true propositions. This assumption is formally represented in epistemic modal logic by the following axiom: if A is known to be true then A is true.⁴ The variable i stands in for agents a, b, c, \dots who know or don’t know propositions. Thus the expression $K_i A$ reads ‘agent i knows proposition A ’. The assumption that knowledge can consist only of true propositions could be called the veridicality axiom⁵, formalized as follows.

Veridicality of Knowledge Axiom: $K_i A \supset A$

The veridicality axiom represents a Platonic view of knowledge according to which genuine knowledge is only of the true propositions. It is a rationality assumption, but one that is an idealization as a representation of knowledge in the context of knowledge-based systems in computing. There a knowledge base is thought to consist of what are taken now to be true propositions even though at some later point in the collection of data, some of them might be shown to be false or dubious. What is referred to here is of course the well-known property of defeasibility in AI. In this sense of ‘knowledge’ what is taken now provisionally to be given the status of knowledge can be defeated as new data comes in, changing what is to be considered knowledge at the later point.⁶

In addition to the veridicality axiom, there are three additional assumptions widely accepted, although controversial, that characterize knowledge as a modality. The second

⁴ Capital letters A, B, \dots , refer to statements (propositions), entities that are true or false.

⁵ Girle (2003, p. 110) calls this axiom the veridicality principle.

⁶ There is a large philosophical issue here of how ‘knowledge’ should be defined. Many epistemologists would say that this defeasible sense of the term is merely true belief. The term ‘knowledge’ as used in computing, and science generally, often has an honorific sense, referring to what is accepted by the scientific community at any given time.

one is the assumption that an agent knows all the logical consequences of any statement that it knows. This assumption is formally expressed in modal logic as follows.

Deductive Closure of Knowledge Axiom: $K_i(A \supset B) \supset (K_i A \supset K_i B)$

The deductive closure axiom is a strong rationality assumption that, in effect, makes a rational agent logically omniscient about its own knowledge base. This axiom has often been doubted, and it is often thought to represent only a highly idealized notion of knowledge of a kind that would be appropriate for a highly rational agent. Girle (2003, p. 110) calls this axiom “distribution”, attributing it to Hintikka. It would appear to be too strong to represent reasoning of everyday human agents (Meyer and van der Hoek, 1995). Indeed, if knowledge implies belief, a problem is that this axiom quickly produces a belief set that is so large as to be cognitively unmanageable.

A third epistemic rationality assumption is that every proposition known by a rational agent is logically consistent with every other proposition it knows.

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

The consistency axiom says that a rational agent never knows a set of propositions that is inconsistent. Girle (2003, p. 111) calls it the consistency principle, putting it in the following equivalent form: if i knows that A then i does not know that not A . This modal idealization would also seem to be unrealistic as applied to the knowledge base of a normal human agent. For this kind of knowledge base is large, and made up of a lot of apparently unconnected propositions. Thus it is likely to contain hidden inconsistencies. We often take sets of propositions as knowledge that later turn out to be inconsistent. It may even be too strong to represent agents engaged in scientific discovery, because newly discovered knowledge may contradict old knowledge.⁷

The fourth rationality assumption is the iteration axiom.

The Iteration Axiom: $K_i A \supset K_i K_i A$

This fourth axiom says 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 iterations. Thus whether the iteration axiom applies to the knowledge base of a normal agent is doubtful.

The idealized modal conception of knowledge represented by these four axioms has provided philosophers with a powerful set of tools for analyzing the notion of knowledge. But even early on, there were doubts about the gap between the idealization and a realistic view of everyday statements about knowledge, knowing and not knowing. Hintikka (1962) at first hypothesized that epistemic modal logics represented by these four axioms can be taken to model the everyday and scientific concepts of knowledge. Later however, he narrowed his position from a more broadly descriptive and normative one to a stance only claiming an idealized model of knowledge-based reasoning of a rational agent (Girle, 2003, p. 121). Recent work in AI on rational agents, and in

⁷ Once again, the term ‘knowledge’ often seems to have an honorific meaning in science. See note 3.

particular on how a rational agent uses the knowledge it has or does not have under conditions of uncertainty and lack of knowledge have led to a move away from these rationality assumptions and toward a defeasible conception of knowledge. According to this conception, an agent may know something now, but later on, as it learns more, it may find out that this proposition is false, or can be replaced by a better one that is more justified by the new data. The rational agent overcomes the inconsistency by accepting the new proposition and retracting the old one from its knowledge base. This model represents a dynamic view of knowledge in which the knowledge base grows and adapts to inconsistencies, as well as to the discovery of new knowledge that replaces the old.

Recent developments in artificial intelligence (Meyer and van der Hoek) have suggested that the modal conception of knowledge may not be all that helpful. In Russell and Norvig (1995 p. 151) a knowledge base is described as “a set of representation of facts about the world”, and each individual representation is a sentence. The sentences are expressed in what is called a knowledge representation language (p. 152). Characteristics of a knowledge base are that there must be (1) a way to add new sentences to it, (2) a way to query what is known, and (3) a way to delete sentences from it that found to be false or inaccurate (p. 152). This means that a knowledge base must have a question-asking system, that can ask questions to the knowledge base. Once an answer is given in the form of a sentence, it can then be added to the knowledge base, or if a sentence is found to be false, or not known to be true, it can be deleted from the knowledge base. The conception of knowledge embodied in a knowledge base of the kind used in computing is very different from the conception of knowledge embodied in the idealized modal conception of knowledge represented by the four axioms above.

Recent research in artificial intelligence (Collins, Warnock, Aiello and Miller, 1975; Reiter, 1980; Konolige, 1988; Branting, 2000; Coppin, 2004) has suggested, however, that such arguments from lack of knowledge are extremely common, both in everyday reasoning and in computing, and that they are quite reasonable in many instances, even though they are defeasible. Such findings suggest that the four axioms of idealized epistemic logic above are not generally characteristic of how everyday conversational participants reason from their knowledge and lack of knowledge. Two problems posed by the research on the argument from ignorance will be shown below to be especially crucial in this regard. One is the problem of the criterion for determining whether a given knowledge base can correctly be said to be complete, meaning that all the true propositions in it can be classified as known. The other problem is one that is also vitally important in analyzing the argument from ignorance – that of determining in a given case the difference between what an agent knows and what it does not know. It is argued below that the best solutions to these two problems need to be based on a non-idealistic view of knowledge as defeasible, meaning that a proposition now known may later be refuted (defeated as knowledge).

3. Three Models of the Form of Argument from Ignorance

If the argument from ignorance can be reasonable in some instances, even though it is a defeasible argument, it must have some general form as an argument. Its form can be represented in three models, a formal model, an epistemic model and a dialectical model. The formal model represents the argument from ignorance using a modal knowledge

operator K . ' KA ' stands for the expression 'proposition A is known to be true'.⁸ In this model, the argument from ignorance takes two forms.

Positive Logical Form of Argument from Ignorance

$\sim KA$

If A then KA

Therefore $\sim A$

Negative Logical Form of Argument from Ignorance

$\sim K\sim A$

If $\sim A$ then $K\sim A$

Therefore A

The negative logical form is clearly an instance of *modus tollens* (MT): If A then B ; not- B ; therefore not- A . The positive logical form is also an instance of MT provided KA is equivalent to $\sim\sim KA$, by the rule of double negation.

If epistemic closure can be taken to have been established in a case, the argument from ignorance has a deductively valid form of argument, namely that of *modus tollens*: The reason is that the conditional premise can be assumed to hold absolutely as true, because the knowledge base has been searched completely. However in realistic instances, lack of knowledge of all the circumstances may apply to the case. This means that the conditional premise is subject to default as further knowledge comes in, and hence the argument from ignorance is best treated as not being deductively valid.

The problem with applying the logical model to common cases of the argument from ignorance concerns the conditional premise. In deductive logic, this conditional is taken in an absolute way, meaning that it only comes out false if the antecedent is true and the consequent false. The problem is that conditionals of the type found in common cases of arguments from ignorance are not absolutistic in this way. They can hold relative to what is known, even if epistemic closure does not obtain in relation to the whole knowledge base. Such a conditional is open-ended in the search process, and can hold relative only to what is normally or usually expected to be known in a given case. This looser type of conditional is represented in the second (conditional) premise of the epistemic model. This model takes into account another factor. In knowledge-based systems, like expert systems, a domain of knowledge D (field of scientific expertise) may be cited.

Argumentation Scheme for Plausible Argument from Ignorance (Walton, 1992, p. 386)

It has not been established that all the true propositions in D are contained in K .

⁸ By not putting in the subscript to the K -operator as done in the set of axioms above, we here ascend to a higher level of abstraction where some set of rational agents is presumed to be constant.

A is a special type of proposition such that if A were true, A would normally or usually be expected to be in K .

A is in D .

A is not in K .

For all A in D , A is either true or false.

Therefore, it is plausible to presume that A is false (subject to further investigations in D).

The epistemic model can hold an argument from ignorance to be acceptable as a reasonable inference to draw subject to further investigations in a knowledge base.

In the pen case, the conditional premise represents a depth-of-search assumption. Wilma asking Bruce if the pen were in the desk, he would know it. In other words, she asks him whether he has searched the desk thoroughly enough to rule out the situation where the pen is in the desk but he does not know that it is. So conceived, the argument from ignorance has an argumentation scheme that shows it is based on two premises about what can be presumed to be known or not known in a case.

The third model is simpler than the first two. The basic argumentations scheme is based not just on what is known or not known to be true, but also on what would be known if it were true. This model fits the basic argumentation scheme below (Walton, 1996, p. 254).

Basic Argumentation Scheme for Argument from Ignorance

If A were true, A would be known (proved, presumed) to be true.

A is not known (proved, presumed) to be true.

Therefore, A is (presumably) false.

The basic argumentation scheme displayed above has explanatory power because of its simplicity. Even though a knowledge base is incomplete, or at least not known to be complete, it can still enable a conclusion to be drawn by rational argumentation on a basis of this kind of argument from ignorance. In such an instance, the argumentation scheme becomes a defeasible form of argument, holding only tentatively as the search through the knowledge base proceeds further. Hence such a defeasible argument from ignorance is only accepted provisionally subject to the asking of critical questions as the search proceeds.

As promised in the analysis of the argument from ignorance in the Roman Medals case visualized in figure, I will now show how *Araucaria* enables the scheme to be applied to the argument as presented in the text of this case. The propositions given as premises and conclusions in figure 1 are listed below in the key list.

Key List for the Roman Medals Example

(A) There are no known instances of Romans being awarded medals for bravery in battle posthumously.

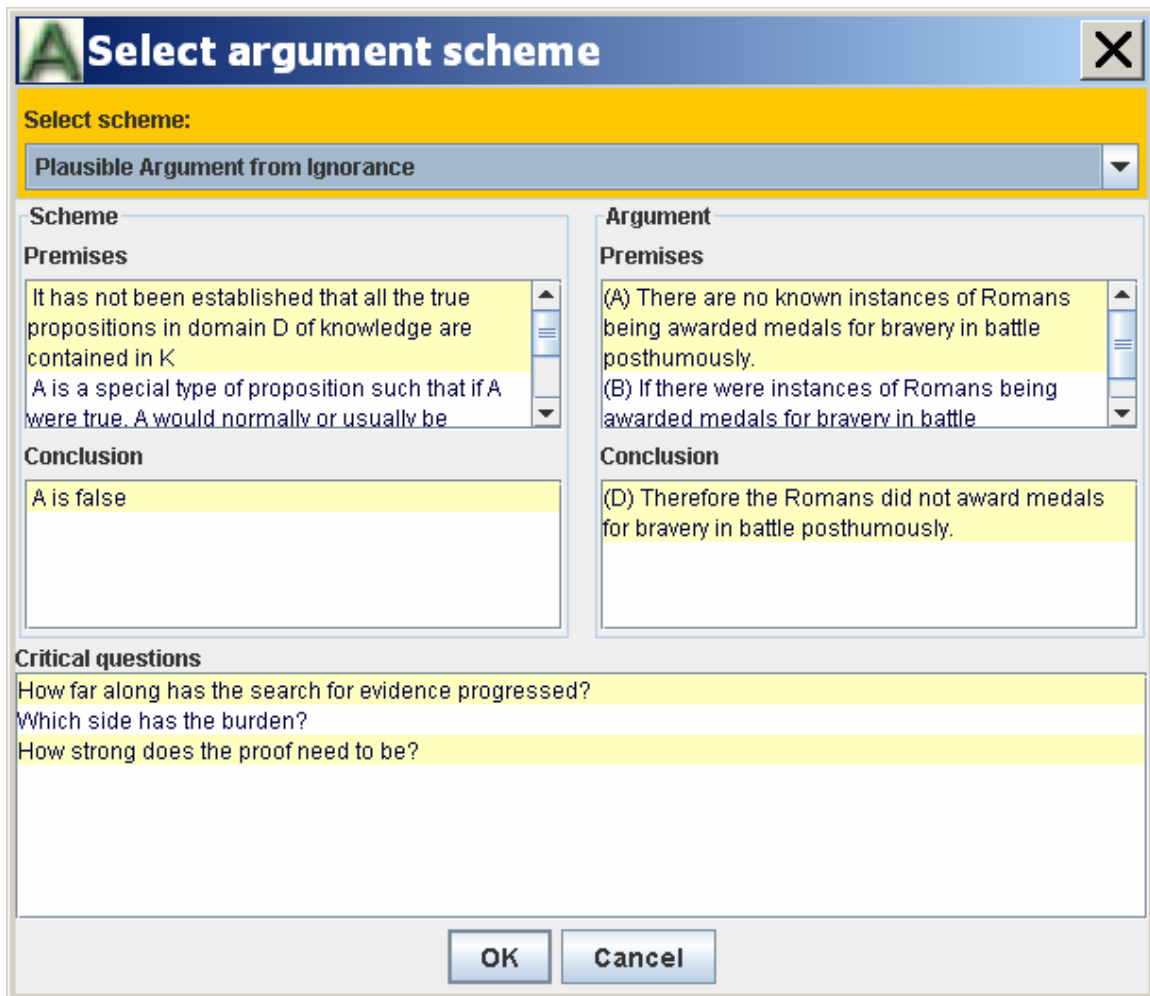
(B) If there were instances of Romans being awarded medals for bravery in battle posthumously would know of them.

(C) We would see evidence on tombstones or in written records of battles.

(D) Therefore the Romans did not award medals for bravery in battle posthumously.

This application reveals the structure of the argument, displaying how the components of this scheme match the propositions given in the example argument.

Figure 2: Screen Shot of the Menu for Argument Scheme Selection in *Araucaria*



The scheme for plausible argument from ignorance is displayed on the left. The individual propositions in the given argument matching the premises and the conclusion of the scheme are displayed on the right. In both instances, the list is incomplete. The list on the right, when complete, matches the key list for the Roman medals example just above.

The menu shown in figure 2 also presents (in abbreviated form) the critical questions matching the given argument in the Roman medals example. The three appropriate critical questions matching the argumentation scheme for the plausible argument from ignorance are more fully stated below.

Critical Questions Matching the Dialectical Scheme for Argument from Ignorance

CQ1: How far along the search for evidence has progressed?

CQ2: Which side has this burden in the dialogue as a whole? In other words, it asks what is the ultimate *probandum* and who is supposed to prove it?

CQ3: How strong does the proof need to be in order for this party to be successful in fulfilling the burden?

CQ1 concerns depth-of-search in the knowledge base. As the examples above have also illustrated, the argument from ignorance typically arises during a process in which knowledge is being collected but at a stage in which the search is not yet complete. To evaluate such a case, CQ1 is obviously very important. At any rate, one can see the value of applying the argumentation scheme to a given case has carried out above. The critical questions prompt the user to seek out the weak points in the argument that need further support by collecting more evidence. The critical questions matching the scheme are the devices primarily used to evaluate the argument from ignorance.

4. Burden of Proof

CQ2 and CQ3 concern burden of proof in cases where the argument from ignorance is a defeasible form of argument put forward by a proponent. It holds tentatively in a dialogue, subject to the asking of critical questions by the respondent. CQ2 and CQ3 questions about the burden. Disputes can arise of the following form.

Proponent: I assert A.

Respondent: Prove it.

Proponent: You disprove it!

Which side has this burden in the dialogue as a whole? CQ2 asks what is the ultimate *probandum* and who is supposed to prove it? CQ3 asks how strong the proof needs to be in order for this party to be successful in fulfilling the burden. As the examples above illustrate, the argument from ignorance often arises at a mid-point in a dialogue during the argumentation stage. To evaluate such a case, two steps have to be taken. First, we have to go back to the starting point of the dialogue where the ultimate *probandum* was

set. Second we have to look at the collection of evidence in the dialogue so far, in relation to that starting point.

Sometimes this back-and-forth process leads to a kind of situation called the *ad ignorantiam* “tug of war” (Walton, 1996, p. 118), where the following case was cited. In a debate in the Canadian House of Commons, the issue was Opposition 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, p. 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” (p. 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.

Disputes about burden of proof were expressed in a dialogue format above. More generally, in dialogue theory, such attempts to shift the burden of proof to the other side can lead to an infinite regress, illustrated by the following sequence of dialogue.

W1: Why *A*?

B2: Why not-*A*?

W3: Why not-not-*A*?

Etc.

As noted above, it may not be possible to decide which side should have the burden of proof in some cases, and the decision may need to be made by a third party. To prevent an infinite regress from blocking progress of the dialogue, the discussion needs to move to a meta-dialogue. Such an interval tries to resolve the issue of which side should have the burden of proof. The embedded meta-dialogue considers evidence from the dialogue up to that point, and the type of dialogue generally. The first task is to determine the ultimate *probandum* of each party, something that should have been determined at the opening stage of the dialogue. The question concerns the role each party took on in the beginning. But as shown by some of the examples considered above, such a problem may not be so easy to solve in some cases. The proposition in dispute as to which side has the burden of proof or disproof may represent an issue that has arisen during the middle of the argumentation stage. It may be unclear how it is related to the ultimate *probandum* of the one side or the other. And thus there may be no easy way to determine who should have to prove or disprove it before anyone should have to accept it. These kinds of cases are problematic, but they do not necessarily represent instance of the fallacious argument from ignorance. They are merely disputes about burden of proof.

The solution to such problems of burden shifting is that the dialogue needs to be supplemented by a metadialogue in which a third party determines which side has the burden of proof. In a trial the judge has the function of making such determinations

(Farley and Freeman, 1995; Leenes, 2001). It may not always be possible to decide which party has the burden. In some cases, *A* may not be relevant to either ultimate *probandum*. But if *A* is relevant to the ultimate *probandum* of one side, then that side will have the burden of proving *A*. The above considerations suggest that the argument from ignorance is often a reasonable form of argument, even though it is often a relatively weak one that is not conclusive by itself. How strong it is depends on how far a search for evidence has progressed. Whether the argument should be accepted or not depends on the burden of proof set for it. If such arguments can be reasonable, why have they for so long been classified as fallacious in logic? This is a question of some significance.

5. Fallacious Arguments from Ignorance

The classic cases often cited as instances of the *argumentum ad ignorantiam* fallacy are witchcraft trials, like the Salem witchcraft trials of 1692, and the McCarthy tribunals of the 1950's. The Salem witchcraft trials were legal proceedings, but the accusation of being in league with the devil was difficult or impossible for the accused party to disprove. Even having a strange appearance or having an "aura" visible only to the witness or the accusers could be positive evidence of guilt. The McCarthy tribunals were not legal trials, but as televised proceedings with examination of witnesses they looked a lot like trials to viewers. The charge of being a "communist sympathizer", like the charge of being in league with the devil, could be based on the flimsiest kind of evidence, but was very difficult to refute. As is often said, proving a negative claim is usually much harder than proving a positive one. In some instances, very flimsy allegations were made. In the case cited by (Copi, 1982, p. 112) as an example of the *ad ignorantiam* fallacy, Senator McCarthy argued as follows: "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." The problem with this kind of argument is the reversal of the burden of proof. The accused was unfairly put in the position of having to try to fulfill a heavy, or even impossible, burden of proof. He or she had to prove her innocence, by using negative arguments like, "I never joined the communist party", or "I never distributed leaflets". As has often been said, it is very difficult to prove a negative assertion, because such a negative claim often involves a kind of generality. To prove you never did something is harder than proving some positive claim that you did something at a specific time and place. Rightly, in our Anglo-American legal system the burden of proof in a criminal case is on the prosecution. The defense merely has to show that the prosecution's argument is not strong enough to fulfill its burden or proving "beyond reasonable doubt". In the McCarthy type of case, there is an attempt to reverse of this burden of proof. But even over and above that, there is a kind of closure to evidence, indicated by the insistence that even no evidence is enough to prove the charge. It is as though the knowledge base is closed off by one side who then insists that the dialogue is over and that its conclusion must be drawn.

But the fallacy in such a case is not just based on a reversal of burden of proof, although that is part of the problem. The other part of it is that the one side is using an aggressive tactic of interrogation that attempts to force premature closure of the investigation. 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. The problem is that

if the accused party tries to ask the appropriate critical questions, he would be discredited as a Communist who cannot be trusted to tell the truth, etc. The fallacy is one of failure of due process in which the capability for critical questioning is shut down. There is lack of knowledge, but the lack of knowledge does not prove what the conclusion it is supposed to, because not enough knowledge has been collected by continuing the investigation far enough to put knowledge and lack of knowledge together for a proper proof.

To sum up, the fallacious kind of *argumentum ad ignorantiam* arises in this kind of case where several tactics are combined. First, not enough evidence has been collected to prove the conclusion by continuing the investigation dialogue by asking further questions and collecting sufficient evidence. Second, the accuser twists the burden of proof or disproof around the other way so that the accused party is forced to try to take on the monumental task of trying to prove he is innocent. Third, the accuser tries to achieve premature closure of the dialogue by shutting down the accused party's ability to raise critical questions that will be taken seriously by the tribunal. There may be very little evidence of guilt, and hence the argument is from lack of knowledge. But also, there is not enough knowledge to fit with the lack of knowledge in the body of evidence needed to support the conclusion. The arguer tries to push the evidence ahead as being conclusive, and as obviating the need for further investigation or examination. The lack of knowledge is combined with premature epistemic closure. The closed world assumption is invoked when it is not justified, and then used to declare that the conclusion can be immediately drawn without further investigation. It is perhaps for this reason that van Eemeren and Grootendorst (2004, p, 182) observe that the *argumentum ad ignorantiam* is frequently found in conjunction with a false dilemma.

6. Epistemic Closure

Epistemic closure (Reiter, 1987) is defined as the condition whereby a collection of new information in a search is closed off, meaning that the evidence is assumed to be complete. In typical legal or political argumentation, where epistemic closure generally cannot be assumed, evaluating an argument from ignorance becomes an issue of burden of proof. 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. But weakness of the argument, by itself, should not be enough for us to categorize it as fallacious. Fallacies are based on argumentation tactics showing a pattern of deception.

If the strength of an argument from ignorance depends partly on closure, what is the criterion for proper closure of a collection of evidence? Let's start with a simple very general rule. The general rule for evaluating epistemic closure is the following criterion.

General Closure Rule

An investigation, or indeed any type of dialogue containing rational argumentation and the collection of evidence, should be closed at the closing stage.

This rule is reminiscent of Yogi Berra's classic statement, "It's not over until it's over". The closure rule rests on the assumption that any argument, like the argument from ignorance, needs to be evaluated not only in light of its logical form (semantic form), but also in relation to a context of dialogue representing a pragmatic framework of use. The first problem in evaluating any contested case, for example on an issue of burden of proof, is to determine the type of dialogue. Persuasion dialogue can be represented by the critical discussion model of rational argumentation presented by van Eemeren and Grootendorst (1984; 1987; 1992). The communal or collective goal of the critical discussion is for both parties to resolve a conflict of opinions by rational argumentation. Each party also has an individual goal of trying to advocate its own viewpoint to the other and prove its ultimate thesis (conclusion). The one viewpoint is opposed to the other, and this opposition is the conflict of opinions to be resolved. The conflict is first identified at a confrontation stage, and the dialogue proceeds through three other stages, an opening stage, an argumentation stage and a closing stage. There are ten rules for the critical discussion given in (van Eemeren and Grootendorst, 1987, pp. 284-293). These rules make the participants defend their claims with rational argumentation and forbid them from attacking each other unfairly. However, it is not just the persuasion type of dialogue that has these four stages. Any investigation in which evidence is being collected can be viewed as a sequence of rational argumentation with four stages.

The classification of basic types of dialogue below (Walton, 1998) postulates six types of dialogue, as represented in table 1. These are not the only kinds of dialogue that represent pragmatic frameworks for rational argumentation. They only represent the most common types that have proved to be especially important in the study of fallacies.

BASIC TYPES OF DIALOGUE

TYPE OF DIALOGUE	INITIAL SITUATION	PARTICIPANT'S GOAL	GOAL OF DIALOGUE
Persuasion	Conflict of Opinions	Persuade Other Party	Resolve or Clarify Issue
Inquiry	Need to Have Proof	Find and Verify Evidence	Prove (Disprove) Hypothesis
Negotiation	Conflict of Interests	Get What You Most Want	Reasonable Settlement that Both Can Live With
Information-Seeking	Need Information	Acquire or Give Information	Exchange Information
Deliberation	Dilemma or Practical Choice	Co-ordinate Goals and Actions	Decide Best Available Course of Action
Eristic	Personal Conflict	Verbally Hit Out at Opponent	Reveal Deeper Basis of Conflict

Table 1

Each type of dialogue has its goal and its four stages. In each instance then, epistemic closure should only be declared at the closing stage, when all of the evidence has been collected, or all the arguments on both sides have been put forward and considered.

This much said however, it has to be admitted that many cases, especially the problematic ones considered above where disputes about burden of proof arise, cannot be solved by simple application of the general closure rule. The reason is that, in many cases of everyday, legal and scientific argumentation, the dialogue is still in the argumentation stage, and yet some decision has to be made about whether to provisionally accept a given argument, like an argument from ignorance. Many ordinary cases of argumentation, especially in deliberation dialogue, occur in conditions of uncertainty and lack of knowledge (Bench-Capon, 1998). In deliberation in everyday life, public as well as private, we rarely know all the evidence in the case, and yet a decision has to be made before incurring the costs and time needed to collect more evidence. We often have to make a decision about whether to keep collecting knowledge and risk delay, or make a decision now, even under conditions of lack of knowledge. In such cases, for practical reasons, the rational course of action may be to invoke epistemic closure and draw a conclusion from the knowledge currently obtained. The bottom line is that in many cases, it may be fair and rational to provisionally draw a conclusion from a knowledge base by invoking epistemic closure, even though it is clear that the knowledge base is incomplete. To say it is incomplete means that it does not contain all the true propositions in a given domain of knowledge. This means, in other words, that more knowledge is out there that could be collected and added to the knowledge base.

For these reasons, the argument from ignorance is best treated as a defeasible form of argument in many instances. If the knowledge base is fully closed, meaning that the dialogue has reached the closing stage, and all the evidence has been collected and assessed, the argument from ignorance can even be treated as deductively valid. As shown above, it is an argument of the *modus tollens* form. But as many of the examples considered above indicate, the closure stage has not been reached yet, and thus it may be rational to provisionally draw a conclusion based on the evidence collected so far, as long as we regard the argument from ignorance as defeasible. This means it should be seen as open to critical questioning. But now the problem is that such cases always need to be decided on a balance of considerations. The argumentation scheme and its matching set of critical questions are the tools needed for ruling on each case. And the argument must be evaluated not just on these localized matters, but also on the basis of the type of dialogue and the global burden of proof for that type of dialogue. Thus the second tool is the typology of dialogues. However, the disputed cases of burden of proof shifting above have shown that even these tools, by themselves, may not be enough to properly evaluate an argument from ignorance in a problematic case. In some cases, the dialogue may have to move to a metadialogue interval in which a third party examines the arguments on both sides and makes some determination of which side should have the burden at that point.

How is the fallacy of *argumentum ad ignorantiam* explainable in relation to epistemic closure? 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 on proper conditions of closure, and not mere speculation or supposition. In some cases, as shown above, the fallacy is also one of failure of due process in which there has been an illicit shift in the burden of proof. The classic case often cited is the McCarthy tribunal. The defendant in such a case must prove he is innocent, but this is impossible because the tribunal is not open to such an outcome. The rule violated in such a case is that both sides in a dispute must have commitment to due process. This means they must take turns questioning and advancing their arguments, and they must abide by the rules for each stage of the dialogue they are in. One must not try to shut the other down during the argumentation stage by invoking epistemic closure of a kind that can only be appropriate for the closing stage.

7. New Defeasible Knowledge Rules for Rational Agents

In this concluding section, five defeasible rules for evaluating knowledge-based arguments are formulated that apply to both arguments from knowledge and arguments from ignorance. I argue that they should replace the currently accepted axioms taken to define knowledge and lack of knowledge and provide reasoning standards for inferences drawn from them. This approach provides further evidence of recent work (Walton, 1996, 2005) showing that the argument from ignorance is not always fallacious, but is often a reasonable but defeasible form of argument.

Knowledge-based reasoning of the same kind represented by the argument from ignorance clearly implies some sort of principles like the veridicality of knowledge axiom and the other three axioms, but not expressed in the absolute, Platonic format these axioms are formulated in. The Platonic format expresses the idea that knowledge is of the fixed and unchanging truths. This format is a high idealization that is not commensurate with the way knowledge bases really work in computing – say, in expert systems technology. To represent a more realistic format, knowledge needs to be based on evidence but open to discovery of new knowledge as evidence continues to be collected. The notions of verifiability and falsifiability of a hypothesis are important here. Something that is knowledge should be open to refutation as new evidence comes in. If this evidence shows it is false, or even that it is open to enough doubt, it should be given up as knowledge. This defeasible format requires that in some instances, a proposition can properly be said to be knowledge, on the basis of the evidence at that time, but then at a later time, it can be shown by new evidence that it needs now to be retracted as no longer being accepted as knowledge. On this new model, knowledge is seen as a species of commitment in the sense of Walton and Krabbe (1995). To say a proposition is known to be true means that it is currently accepted as true, based on the evidence now available, and this evidence gives good reasons to accept it as true.⁹ However, to say a proposition *A* is known to be true, in this sense, is compatible with its being discovered to be false at some point in the future as the collection of evidence proceeds.

⁹ Once again the question is raised of whether knowledge can be abstracted from the agents that are held to possess it. The thesis above relativizes knowledge to a community of knowers. But who are they? Are they individual agents, the community of scientists, or the general population? I shall make no comment on this philosophical question.

This sense of the term ‘knowledge’ implies that knowledge represents acceptance in a field, or domain of knowledge. But it is more than just acceptance, or even justified belief. Knowledge represents the set of propositions that are generally accepted in a field, based on the methods and evidence in that field as a domain of knowledge. But the veridicality axiom does not hold. What does hold is the following defeasible veridicality rule.

Defeasible Veridicality Rule

If A is known to be true then A is generally accepted as true in a domain of knowledge based on the body of evidence in that domain, even though A may later be found to be false as new evidence comes in.

Such a set of propositions is presumed to be consistent, but in reality it may not be. Yet if an inconsistency is found in it, it needs to be removed, or at least the set of propositions constituting knowledge needs to be altered so that the inconsistency in that set no longer exists. Thus the consistency of knowledge holds, but only in a modified form.

Defeasible Consistency of Knowledge Rule

For any two propositions A and B , if A and B are known to be true, then it is assumed that A and B are consistent with each other, but if they are found not to be consistent with each other, one or the other must be rejected as being knowledge.

This rule allows that a rational agent might be properly said to know a set of propositions that contains an inconsistency within it. The possibility exists. However, it is assumed that such a set is consistent. Thus if it found to be inconsistent, the inconsistency must be dealt with and somehow resolved and removed.

What does this imply about the deductive closure of knowledge axiom? Rational agents do not need to be logically omniscient. The set of logical consequences of any given set of propositions in a knowledge base is infinitely large. It is not necessary to assume that such a set of propositions known to be true is closed under deductive logical consequence. Only the following version of this axiom is necessary

Defeasible Closure of Knowledge Rule

If A is known to be true and B is a logical consequence of A then B also has to be accepted as known to be true once it has been proved that it follows from A .

Thus the closure axiom is related to the consistency axiom, and to the following defeasible rule of refutation.

Defeasible Rule of Refutation

If A is accepted as known to be true but is shown by a valid chain of reasoning to imply a proposition B that is known to be false, that is inconsistent with A or any other

proposition that is known to be true, or that has evidence against it, A must be given up as known to be true.

The defeasible rule of refutation defines the general conditions for retraction of a proposition from a knowledge base.

Finally, the iteration axiom has to be considered in its defeasible form. If a rational agent knows that A is true, must it be said that it knows that it knows that it knows that A is true? To examine this question we have to ask what it means to say that agent knows that it knows that a proposition A is true, under the defeasible conception of knowledge. It means that there is a basic claim and a secondary claim. The basic claim is that there is good evidence for A and A is generally accepted as true based on this evidence. The secondary claim is that the agent who knows A has good evidence that the basic claim is true and it is generally accepted that the basic claim is true. This axiom does not seem to be necessary to the defeasible view of knowledge, as far as one can tell. And again there is the problem of its recursive expansion. For example if an agent knows that it knows A to be true, it follows from the iteration axiom that it knows that it knows that it knows A to be true. It is hard to make sense of such iterations, and they do not seem to be necessary to the defeasible conception of knowledge. Thus it seems best to at least provisionally conclude that the iteration axiom fails.

Finally, the fifth rule is that corresponding to the argument from ignorance.

Defeasible Rule for Argument from Ignorance

On the suppositions that (a) if A were to be true it would be known to be true, and (b) A is not known to be true, it follows that (c) A can be assumed to be false.

The defeasible character of this fifth rule is evident from the way it has been phrased above. It rests on the assumption that a knowledge base can be closed provisionally in order to draw a conclusion by defeasible reasoning. For example, suppose that in deliberation we can no longer continue to afford searching for new evidence, and have to make a decision about whether to accept A as true or operate on the assumption that A is false. We know that the knowledge base is fairly extensive, even if we are not in a position to say it is complete. Thus we can say with some confidence that if A were true, it would in all likelihood be in the knowledge base. For example, suppose that we know that the expert system knows a lot about coffee production in South America, and thus we know that if Guyana were a major coffee producer, the system would know that. But there is no proposition in the knowledge base saying that Guyana is a major coffee producer. We could be justified on the basis of the defeasible axiom from ignorance in drawing the conclusion that the proposition ‘Guyana is a major coffee producer’ is false.¹⁰ Of course, we do not know that for sure, in the Platonic sense that it definitely has to be true beyond all doubt or reservations. Still, it’s a pretty good provisional conclusion to draw from what we know and what we don’t know.

8. Conclusions and a Further Problem to be Solved

¹⁰ This example is similar to one given by Collins, Aiello and Miller (1975, p. 398).

In conclusion, some consideration should be made on how to formalize the five defeasible rules for knowledge-based reasoning above. Using the Philonian conditional of classical deductive logic is not adequate to model the conditionals expressed informally in the five rules. The truth-functional Philonian conditional is only adequate in cases where the knowledge base has been closed and will not be re-opened to admit more evidence as the argumentation continues. The Philonian conditional, in the author's view, represents the kind of case where the conditional has the following form: if A , and all other evidence is held constant and fixed, and not subject to further collection of knowledge then B . The defeasible conditional represents this kind of case: if A , then generally but subject to exceptions and possible reversal in the future as new evidence come in, then it is reasonable to tentatively accept B . On this view, admittedly a controversial one, there exists a defeasible *modus ponens* as well as a deductively valid one. A proposal formalizing a kind of conditional that could represent both types of reasoning in multi-agent systems has been put forward in (Reed and Walton, 2004). But there is a further problem to be solved. Is deductive logic ever applicable to cases of arguments based on knowledge and lack of knowledge, or are such cases irreducibly defeasible? This issue remains highly controversial.

The issue is whether we can ever invoke the closed world assumption, implying that there is no longer any lack of knowledge in a given investigation, and then proceed by deductive reasoning, or whether there is always lack of knowledge in any real case of argumentation. Artificial intelligence is full of example cases of artificial problems that can be used as exercises for students to illustrate solving the problem by searching through a database and examining all the possible solutions. The most famous example is the blocks world problem, where the world is circumscribed artificially. In blocks world, there is only a flat surface with several piles of blocks, and the only action allowed is to take a block from the top of one pile and either move it to the top of another pile, or move it on to the flat surface (Coppin, 2004, 428-430). In this example, there are only a small, finite number of solutions to the problem, representing different ways the blocks can be moved to achieve the desired outcome. But is an artificial situation like that in the blocks world problem ever representative of a new real situation, where unanticipated disturbances are possible? McCarthy (1980) raised this question by imagining someone attempting to solve such a problem by asking questions like, 'What would happen if the table fell down?' He invented a form of non monotonic reasoning called circumscription to deal with cases of this sort. It works in a way similar to the closed world assumption by allowing us to assume that no facts are true other than those stated in the expression. As applied to the blocks world problem, for example, it would allow us to conclude that any facts not explicitly stated in the problem are not true.

But the problem is that the closed world assumption, or a circumscription device like that introduced by McCarthy (1980), can never be unconditionally applied to real world situations, excluding what are called acts of God in law. For example, blocks world may be struck by lightning. We can artificially exclude such unanticipated possibilities in an artificial problem, but as applied to examples of reasoning in the real world, like the examples of argument from ignorance we studied above, this kind of exclusion is not realistically warranted. Indeed, as we saw in the examples of argument from ignorance, the fallacy in such cases is precisely a closed-minded approach that invokes epistemic closure to seal off any further investigation or collecting of data.

Thus the problem remains of whether we can rightly say that in such cases, we know all there is to know, and there is no longer any lack of knowledge. Theoretically, a database can be closed off by invoking the closed world assumption, or some other device like circumscription, and then we can evaluate the reasoning on that basis. In such cases where the knowledge base has been finally closed off, because a dialogue has reached the closure stage, and the closed world assumption has been invoked, the knowledge-based argument could be evaluated as deductively valid. In such cases the four axioms in section 2, the veridicality axiom, the deductive closure axiom, the consistency of knowledge axiom and the iteration axiom, can properly be applied as criteria for judging the validity or invalidity of the argument.

But the problem remains that these axioms do not work when applied to defeasible knowledge-based arguments, like the argument from ignorance, that are put forward as means for drawing a provisional conclusion during the argumentation stage. To supplant these four axioms we need defeasible rules for knowledge-based argumentation. These defeasible rules were formulated in the previous section using the basic notions of evidence and general acceptance. Basically a proposition is said to be known to be true in this defeasible sense if it is generally accepted as true in a domain of knowledge (field) and if the evidence in the field is strong enough to support its acceptance. How strong the evidence needs to be in a given field is a matter of burden of proof we leave open. If this approach is right, and we argue here that it is, the closed world assumption can only be justified on the basis of burden of proof. It can be justified in artificial cases, like the blocks world of them, where a *ceteris paribus* clause is invoked. But are any real cases like that? It is open to argumentation that they are not. It remains reasonable to propose the hypothesis that in real cases, there is always a mixture of knowledge and lack of knowledge. Thus both arguments from knowledge and arguments from lack of knowledge need to be evaluated in the context of an ongoing dialogue or investigation. The basic conclusion of the paper is that knowledge-based arguments like the argument from ignorance need to be evaluated by criteria for epistemic closure that are pragmatic in nature and need to be formulated and applied differently at different stages of an investigation or discussion. The problem remains open of how to decide in a specific case when the dialogue can be closed off and the closed world assumption invoked, so that by fiat, there is no longer any absence of knowledge.

In light of these defeasible versions of the original four axioms for knowledge considered in section 2, what needs to be said then about conditions for closure of a knowledge base? The answer is that a knowledge base can be closed off for pragmatic reasons in order to draw tentative conclusions from what is known and not known at a given point during a process in which evidence is being collected (Branting, 2000). We can say that the discussion is over or that the investigation is concluded at this point, in order to now go ahead and make a decision or draw a conclusion based on what is now known and not known. We can make such a decision for closure rationally even though later, we might re-open the discussion or investigation and start to collect more evidence. Thus tentatively we can draw conclusions from the evidence that has been collected to this point, based on rules of defeasible reasoning like the five rules formulated above.

References

Trevor Bench-Capon, 'Specification and Implementation of Toulmin Dialogue Game', *Jurix 1998: The Eleventh Conference*, ed. J. Hage et al., Nijmegen, Gerard Noodt Instituut, 1998, 5-20.

L. Karl Branting, *Reasoning with Rules and Precedents: A Computational Model of Legal Analysis*, Dordrecht, Kluwer, 2000.

Allan Collins, Eleanor H. Warnock, Nelleke Aiello and Mark L. Miller, 'Reasoning from Incomplete Knowledge', *Representation and Understanding*, ed. Daniel G. Bobrow and Allan Collins, New York, Academic Press, 1975, 383-415.

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

Ben Coppin, *Artificial Intelligence Illuminated*, Sudbury, Massachusetts, Jones and Bartlett Publishers, 2004.

Arthur M. Farley and Kathleen Freeman, 'Burden of Proof in Legal Argumentation', *International Conference on Artificial Intelligence and Law*, College Park, Maryland, 1995, 156-164.

Rod Girle, *Possible Worlds*, Chesham, England, Acumen Publishing Limited, 2003.

Jaakko Hintikka, *Knowledge and Belief: An Introduction to the Logic of the Two Notions*, New York, Cornell University Press, 1962.

Kurt Konlge, 'On the Relation between Default and Autoepistemic Logic', *Artificial Intelligence*, 35, 1988, 343-382.

Ronald E. Leenes, 'Burden of Proof in Dialogue Games and Dutch Civil Procedure', *Proceedings of the 8th International Conference on Artificial Intelligence and Law*, St. Louis, Missouri, ACM Press, 2001, 109-118.

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

John McCarthy, 'Circumscription: A Form of Non-Monotonic Reasoning', *Computation and Intelligence: Collected Readings*, ed. George F. Luger, Cambridge, Mass., MIT Press, 1980.

John-Jules Meyer and Wiebe van der Hoek, *Epistemic Logic for Computer Science and Artificial Intelligence*, Cambridge Tracts in Theoretical Computer Science 41, Cambridge University Press, 1995.

Chris Reed and Glenn Rowe, 'Araucaria: Software for Puzzles in Argument Diagramming and XML', *Technical Report*, Department of Applied Computing, University of Dundee, 2002.

Chris Reed and Doug Walton, 'Towards a Formal and Implemented Model of Argumentation Schemes in Agent Communication', in I. Rahwan, P. Moraitis, and C. Reed, (eds) *Proceedings of ArgMAS 2004*, Springer Verlag, 2004.

Raymond Reiter, 'A Logic for Default Reasoning', *Artificial Intelligence*, 13, 1980, 81-132.

Raymond Reiter, 'Nonmonotonic Reasoning', *Annual Review of Computer Science*, 2, 1987, 147-186.

Glenn Rowe, Fabrizio Macagno, Chris Reed and Douglas Walton, 'Araucaria as a Tool for Diagramming Arguments in Teaching and Studying Philosophy', *Teaching Philosophy*, 29, 2006, 111-124.

Stuart J. Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, Saddle River, New Jersey, Prentice Hall, 1995.

T. D. Sterling, W. L. Rosenbaum and J. J. Weinkam, 'Publication Decisions Revisited: The Effect of the Outcome of Statistical Tests on the Decision to Publish and Vice Versa', *The American Statistician*, 49, 1995, 108-112.

Frans H. van Eemeren and Rob Grootendorst, *Speech Acts in Communicative Discussions*, Dordrecht, Foris, 1984.

Frans H. van Eemeren and Rob Grootendorst, 'Fallacies in Pragma-Dialectical Perspective', *Argumentation*, 1, 1987, 283-301.

Frans H. van Eemeren and Rob Grootendorst, *Argumentation, Communication and Fallacies*, Hillsdale, N. J., Erlbaum, 1992.

Frans H. van Eemeren and Rob Grootendorst, *A Systematic Theory of Argumentation*, Cambridge, Cambridge University Press, 2004.

Douglas Walton, 'Nonfallacious Arguments from Ignorance', *American Philosophical Quarterly*, Vol. 29, 1992.

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

Douglas Walton, *Appeal to Expert Opinion*, University Park, Penn State Press, 1997.

Douglas N. Walton and Erik C. W. Krabbe, *Commitment in Dialogue*, Albany, New York, State University of New York Press, 1995.

Douglas Walton, 'Pragmatic and Idealized Models of Knowledge and Ignorance', *American Philosophical Quarterly*, 42, 2005, 59-69.

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