Argument and Computation
Vol. 2, Nos. 2–3, June–September 2011, 131–155

Reasoning about knowledge using defeasible logic

Douglas Walton*

Centre for Research in Reasoning, Argumentation, and Rhetoric (CRRAR), University of Windsor, 2500 University Ave. W., Windsor, Ontario, Canada N9B 3Y1

(*Received 9 March 2011; final version received 24 October 2011*)

In this paper, the Carneades argumentation system is extended to represent a procedural view of inquiry in which evidence is marshalled to support or defeat claims to knowledge. The model is a sequence of moves in a collaborative group inquiry in which parties take turns making assertions about what is known or not known, putting forward evidence to support them, and subjecting these moves to criticisms. It is shown how this model of evaluating evidence in an inquiry is based on a defeasible logic using forms of argument that admit exceptions. It is contended that reasoning from absence of knowledge is as important to inquiry as positive reasoning from evidence to knowledge. The philosophical conflict between this view of reasoning about knowledge and the true-belief-plus view is explored by airing objections and replies on both sides.

Keywords: bounded procedural rationality; defeasible logic; inquiry; Carneades Argumentation System; truth and knowledge; evidence; standards of proof; scientific knowledge; falsification; reasoning from absence of knowledge

This paper adopts a viewpoint of bounded procedural rationality to build a model representing the structure of the process whereby reasoning is used to justify the claim that a proposition should have the status of knowledge. Two important elements of the model are the requirements that the process uses evidence both for and against the claim and that it is based on defeasible reasoning. The model is primarily meant to represent scientific knowledge as a body of accepted propositions, but subject to some reservations, it may also be used to represent the reasoning used in knowledge claims in everyday conversational discourse. The model is implemented in the Carneades system, an argument mapping application, with a graphical-user interface, and a software library for supporting argumentation tasks, including argument mapping and visualisation, argument evaluation, applying proof standards that respect the distribution of the burden of proof, and argument construction from defeasible rules. In the model, a proposition can be classified as knowledge if and only if (1) it has been proved in an investigative procedure called an inquiry, (2) to the proof standard appropriate for the inquiry, (3) based on the evidence marshalled during the inquiry, and (4) using the kind of evidence that is admissible in the inquiry.

The paper begins in Section 1 by describing the philosophical conflict between two opposed views of knowledge. The one view, currently the dominant view in epistemology, is shown to be characterised by four defining principles stating that (1) knowledge bases contain only truths, (2) knowledge bases are consistent, (3) knowledge bases are closed under deductive implication, and (4) knowledge bases contain the assumption that if a proposition is known, then it is known that it is known. The other view, a fallibilistic view attributed to Peirce and Popper, holds that a

*E-mail: dwalton@uwindsor.ca

ISSN 1946-2166 print/ISSN 1946-2174 online

http://dx.doi.org/10.1080/19462166.2011.637641
http://www.tandfonline.com

proposition can be established as knowledge during the course of an inquiry, but later rejected as knowledge once new evidence has falsified it. The two views of knowledge are shown to be incompatible. The sceptical roots of this fallibilistic view of knowledge are shown to be traceable back to the ancient scepticism of Arcesilaus and Carneades. In Section 2, an evidence-based model of inquiry that has five defining characteristics is applied to a standard example of reasoning about knowledge using the Carneades argumentation system. The model shown in Section 3 is based on a procedural view of inquiry in which evidence can support or defeat claims to knowledge. The argumentation structure of this procedure is modelled in Section 4 by an extension of the Carneades system as a sequence of moves in a collaborative group inquiry in which parties take turns making assertions and putting forward evidence to support them. It is shown in Section 5 how this model of evaluating evidence in an inquiry is based on a defeasible logic using forms of argumentation that admit exceptions. It is the contention of Sections 6 and 7 that reasoning from absence of knowledge (RAK), or lack of evidence, is as important to inquiry as positive reasoning from evidence to knowledge. The philosophical conflict between the two views of reasoning about knowledge is revisited in Section 8, where the main objections and replies on both sides are discussed. In the last section, Section 9, conclusions of the paper are presented.

1. Two opposed views of reasoning about knowledge

According to the justified true belief analysis, to qualify as knowledge, a proposition[1] *p* must meet the following three requirements: (1) *p* is true, (2) the agent who claims to know that *p* is true believes that *p*, and (3) the agent is justified in believing that *p*. Surprisingly, (1), the veracity condition, has not generated any significant degree of discussion (Steup 2010). Perhaps the most familiar philosophical definition of knowledge is captured by Zagzebski's (1999, p. 93) phrase that knowledge is "true belief plus something else". On this type of account, propositional knowledge is explained as some form of "good true belief" (Zagzebski 1999, p. 99). The justification condition has been explained in a variety of ways, including evidentialist accounts, causal accounts, and reliabalist accounts (to mention a few).

Accounts of knowledge as true-belief-plus have also been adopted by many theorists working within an epistemological approach to argumentation (Lumer 2005a, p. 190; 2005b, p. 215). Indeed, Lumer (2005a, p. 192) has gone so far as to claim that theories which use the term "knowledge" in a sense different than that adopted within normative epistemology – which, Lumer claims, relates knowledge and justified belief to objective truth conditions (e.g. by identifying knowledge with the current stock of expert opinions) – do not count as epistemological approaches (Walton and Godden 2007).

Classical logic has a model-based semantics established on a notion of truth in which truth is a relation between a proposition and an external reality that provides a criterion for telling whether a proposition is true or false. This feature can be seen in classical deductive logic in which the propositions take on the truth values true and false. It has proved to be a useful model to represent mathematical reasoning. Epistemic reasoning has traditionally accepted this model, which is based on four key assumptions about how knowledge relates to truth and logical reasoning. The general assumption has been that knowledge can be modelled in formal systems of modal logic, following four assumptions. The relation $\rightarrow$ represents the material conditional used in classical deductive logic. Rescher (2003, pp. 10–11) takes these four principles, along with some others, to represent defining characteristics of knowledge. He wrote (2003, 10) that although "some writers see the linkage between truth and knowledge as a merely contingent one", this view in not tenable. For the purposes of this paper, these four principles are taken as axioms representing the currently dominant view in epistemology.

(1) veracity: $Kp \rightarrow p$
knowledge bases contain only truths; if p is known then p is true.
(2) consistency: $\sim(Kp \;\&\; K\sim p)$
knowledge bases are consistent; if p is known then it is not the case that not-p is known.
(3) deductive closure: $K(p \rightarrow q) \rightarrow (Kp \rightarrow Kq)$
knowledge bases are closed under deductive implication; knowledge includes all of the logical consequences of any proposition known.
(4) iteration: $Kp \rightarrow KKp$
the contents of knowledge bases are transparent; if p is known, then it is known that p is known.

In his survey of the logic of knowledge, Rescher (2005, p. 4) wrote that the veracity assumption obtains as a general principle for systems of epistemic logic. He formulated this principle as saying "If $K_x p$ then p", where x is an intelligent knower. This 'if–then' is taken to represent a deductive connection. It could be strict implication based on the modal necessity operator, but here we take to represent the material conditional of classical deductive logic.

Cooke (2006, p. 1) defined Peirce's way of viewing fallibilism by linking it to iterated knowledge claims and by contrasting it with the view of knowledge held in traditional epistemology. According to traditional epistemologists, in order to know, an inquiring agent must be in an epistemic position to know that he knows. In contrast, the doctrine of fallibilism does not allow an inquirer to judge that he possesses the truth concerning the matter inquired into. He not only denies the iterated axiom that if A is known to be true then A is known to be true. He even denies that the agent can be in an epistemic position to know that he knows that A is true. Thus, the iteration fails, following the view of the Peircean fallibilist.

Walton (2005) calls any theory of knowledge meeting these four criteria an *idealised* model of knowledge, and argues that any such idealised model is unsuitable for pragmatic purposes of modelling knowledge based on defeasible epistemic reasoning of the kind found in scientific discovery and investigation. Walton argues that in such cases, knowledge bases are not consistent, transparent, or closed under deductive implication. The reason is that scientific knowledge, even if based on confirmation that meets a high standard of proof, needs to always be open to continued testing, and hence needs to be seen as defeasible. Walton (2005) proposed a pragmatic conception of knowledge, which is built upon two common sense platitudes that capture our everyday epistemic situation in the world: (1) a knowledge base can be incomplete, and (2) a knowledge base can be fallible. A knowledge base can be incomplete in the sense that there can be many true claims which are not included in the knowledge base. Further, "[o]n this [pragmatic] model, knowledge is defeasible, meaning that a proposition now known may later be refuted (defeated as knowledge)" (Walton 2005, pp. 59–60), thus allowing for retraction in the process of inquiry, investigation, and discovery.

On the Peircean view, the real aim of an inquiry of the kind that takes place in a finite amount of time and resources for collecting evidence is not that of actually reaching the truth, and knowing that it has been reached, but only that of a firm settling of opinion. Peirce (1984, p. 354) wrote that the "only legitimate aim of reasoning is to ascertain what decision would be agreed upon if the question were sufficiently ventilated". Taking the view that truth is the result of inquiry would "block the path of inquiry because our minds would be closed, and hence, we would never be motivated enough to inquire" (Peirce 1931, p. 6.3). Clearly, acceptance of the veracity axiom is not consistent with Peirce's view of the inquiry. In addition, he warned us not to infer from the premise that we can be substantially certain about many things to the conclusion that we "perfectly know when we know" (Misak 1987, p. 260). It is a corollary that the iteration axiom does not work in the Peircean inquiry either. Maintaining the veracity condition is at odds with the important role that defeasible reasoning plays in the process of evaluating claims to knowledge in an inquiry.

These are strongly opposed views. The veracity principle is so widely accepted in epistemology that epistemologists feel that it would be unthinkable to reject it. Rescher (2003, p. 10) considers that holding that the linkage between knowledge and truth is merely contingent "inflicts violence on the concept of knowledge as it actually operates in discourse". His reason is that the locution "an agent knows that proposition *p* but *p* is not true" is "senseless". He gives two arguments to back up this contention.

The first argument is based on the premise that a person cannot be said to know that something is the case when this person is not prepared to accept it. The conclusion he holds to follow from this premise is that a claim that a person knows that proposition *p* is only tenable when that person holds *p* to be the case. Note, however, that given that premise of this argument is true, it does not follow that proposition *p* has to be true. It only follows that the person holds *p* to be true. In other words, what follows is not truth but acceptance: if a person says that he knows that a proposition is the case, it follows that he must be taken to accept this proposition. What does not follow is that the proposition itself must actually be true.

The second argument is that a person cannot properly be said to know that a proposition is true unless he is prepared to accept the proposition, as a true premise in his thinking and a suitable basis for his actions. But as with the first argument, the conclusion of this argument is that knowledge implies acceptance. The argument does not prove the veracity principle that knowledge implies truth.

Both these arguments can be seen to be open to doubt, based on the defeasible model of reasoning in an inquiry presented in the rest of the paper. It is an acceptance-based model, and does not require any of the four principles of reasoning about knowledge. In this fallibilistic model, because of the principle of falsification of scientific knowledge, disputed cases may need to become the subject of further inquiry. In defeasible epistemic reasoning, there is no guarantee that a conclusion drawn is actually true, with reference to some standard of truth that is a relation between a proposition and an external reality. In defeasible reasoning, the conclusion can be taken to be proved to be true based on the evidence that has been collected so far, but when new evidence is collected that same proposition may be turned out later to be proved to be false, as much as this outcome is meant to be avoided in a careful investigation.

The problem is that we can have lots of evidence in favour of a proposition and no evidence against it, so that on balance, the proposition can rightly be said to be known to be true, but later on, as more evidence comes in, we might find that this proposition is false. It is important for scientific knowledge that it be represented as open to defeat in this kind of situation as new evidence comes in. This defeasibility requirement, however, is inconsistent with the traditional definition of knowledge as justified true belief. This requirement is not consistent with the notion that sometimes propositions rightly accepted as scientific knowledge later on turn out to be disqualified or rejected as knowledge, once new experimental findings come in, or a new theory comes in that offers a better explanation than the previous one.

The historical motivation of the defeasible knowledge approach is to be found in the variant of academic scepticism developed by Arcesilaus, Carneades, Philo of Larissa, and Cicero. As reported by Cicero, the Greek sceptic Arcesilaus adopted the view, appropriated from the Socratic dialogues, that nothing can be apprehended with certainty by the senses or the mind. He concluded that "truth is submerged in the depths" (Thorsrud 2002, p. 6). From these premises, Arcesilaus also drew the conclusion that knowledge is not possible. From his premises and conclusions, we can take it that Arcesilaus assumed, as an implicit premise, that truth is a requirement of genuine knowledge. So reconstructed, his argument runs as follows: we cannot have truth (or be sure we have it); knowledge requires truth; therefore, we cannot have knowledge. His conclusion is the classical one of the sceptic. But one can be a fallibilist without being a sceptic, for there is another route open to the fallibilist. He can concede that while truth cannot be known, at least with certainty beyond all doubt, knowledge of a fallible sort can be obtained through a process of collecting and testing evidence in a systematic inquiry, even though the knowledge obtained as the outcome of this procedure cannot be verified at the end of it (beyond further questioning) as a true proposition.

This view that Cicero attributes to the Academic sceptics suggests an approach to epistemology that does not require the axiom of veracity. According to this epistemology, an open-minded rational agent who is in search of the truth can be motivated and directed by this concern, even though he is sceptical about arriving at the end of the search with knowledge of the truth that cannot later be disputed or doubted. Truth is very important in such an inquiry, because regard for the truth is an ideal that motivates the inquiry. But one can argue that this view is consistent with the view that there are difficulties and human limitations concerning the acquisition of a kind of knowledge that implies the truth.

Peirce held a form of fallibilism that held that all our knowledge is fallible, subject perhaps to only one exception: "No; but there is nothing at all in our knowledge which we have any warrant at all for regarding as absolute in any particular .... If I must make any exception, let it be that the assertion that every assertion but this is fallible, is the only one that is absolutely infallible" (Peirce 1931, p. 2.75). Peirce wrote that many things are "substantially certain" (Peirce 1931, p. 1.152), but that this is different from the kind of absolute certainty that implies truth. On his view truth is an aim of inquiry but it can only be reached during the procedure of an inquiry after an infinite process of evidence and argumentation that would take an infinite time. He concluded that knowledge needs to be seen as an approximation to the truth that has survived the testing inquiry that has examined all plausible views, and selected the one that is most likely to be true. On this view there is an asymmetry, in that while knowledge can be falsified by testing, it can never be verified in the sense that it can be proved to be true beyond doubt.

The distinction between verifiability and falsifiability is central to Popper's philosophy of science and an important aspect of how Popper defined scientific knowledge. On his view, to be correctly considered scientific knowledge, a theory must be falsifiable. However, note that it does not have to be verifiable, at any rate meaning that it has to be true beyond any possibility of doubt. According to Popper's philosophy of critical rationalism, human knowledge is based on reasoning of a kind that is conjectural. He held that scientific hypotheses are falsifiable, but not verifiable, in the sense that positive outcomes in experimental testing can give evidence to support a hypothesis, but do not imply the conclusion that the hypothesis is true, or known to be true. For Popper, the search for truth is one of the strongest motives for scientific discovery. Like Peirce, he held that that a scientific investigation can measurably come closer to truth (verisimilitude) with respect to the amount of truth and falsity it implies. He claimed that verisimilitude is measurable, and contended that it was an aspect of his philosophy of science that scientific knowledge is objective, in the sense that it is (1) based on an evidential procedure that moves toward truth as its goal and (2) is independent of the knowing subject.

The fallibilist approach to reasoning about knowledge is strongly opposed to the generally dominant approach in epistemology that views knowledge as justified true belief. On the fallibilist view, all that is necessary for a proposition to be accepted as knowledge is for the inquiry procedure to prove the proposition, based on the arguments for and against it that are produced and evaluated during the procedure itself.

2. Evidence and knowledge

What is evidence, and how is it related to knowledge? Evidence is made up of a set of evidential data or items of evidence and of inferences reasonably drawn from them. One part of evidence is the perception of the data through the senses, or sensors (in the case of an automated agent or a device with memory, like a black box data recorder). The other is the drawing of conclusions from the perception of such data, based on inference. If I see something that looks like a red car in the parking lot, my perception of the red car is data, and from these data I can draw a reasonable inference to the conclusion that there is a red car in the parking lot. If I say later that there was a red car in the parking lot, this conclusion can also be accepted as evidence, but in such an instance it depends on my memory. If I tell somebody else later that I saw a red car in the parking lot, they can also draw the inference to the conclusion that there was a red car in the parking lot, based on testimonial (witness) evidence.

Knowledge is based on evidence. Indeed, being evidence-based is a requirement of knowledge, in the sense of the term meant here. This meaning of the term comprises evidence in many different disciplines, including ones in science and law, such as archaeology, forensic evidence, medicine, and authentication of works of art and other artefacts.

In general, the relationship between knowledge and evidence has five components. The first component is the ultimate conclusion to be proved, the proposition that is claimed to have the status of knowledge. The second is the body of data that is being put forward as the basis for drawing inferences from the data. The third is the marshalling, or collecting together of the data and conclusion drawn from them that are relevant to support the claim that this proposition can be classified as knowledge. This is the body (or mass) of evidence. The fourth is the chain of reasoning that provides the argumentation, or justification as it might equivalently be called, proving the ultimate proposition from the mass of evidence. The fifth is the standard of proof that this chain of reasoning has to meet in order to prove the ultimate conclusion. Evidence comes to be knowledge through a dialogue procedure called an inquiry which can be modelled as a dialogue system.

A basic proposition is taken in current epistemology to be one like "I see a red patch now" that is not justified as knowledge by some other proposition that functions as evidence supporting it. A basic proposition, in other words, is taken to be knowledge because it is immediately evident, rather than requiring some other evidence to back it up. The possibility of there being basic propositions seems to go against the fallibilistic view that all knowledge is based on evidence collected in an inquiry and supported by reasoning from the evidence that backs up the claim to knowledge. For it would appear that in the case of a basic proposition, something that can be taken as knowledge even if no additional evidential basis supported by rational argumentation is needed to back it up.

A reply to this objection can be given by considering Pollock's (1995) theory of defeasible reasoning. Pollock (1995) drew a distinction between two kinds of arguments that can defeat another argument: rebutting defeaters and undercutting defeaters. A rebutting defeater gives a reason for denying a claim (Pollock 1995, p. 40). Thus, a rebutting defeater attacks the claim or the conclusion the argument it is aimed at. An undercutting defeater attacks the connection between the claim and the reason rather than attacking the claim directly (41). Pollock's classic example (1995, p. 41) clarifies the distinction.

> For instance, suppose *x* looks red to me, but I know that *x* is illuminated by red lights and red lights can make objects look red when they are not. Knowing this defeats the prima facie reason, but it is not a reason for thinking that *x* is *not* red. After all, red objects look red in red light too. This is an *undercutting defeater.* (Pollock's italics)

Suppose I see a red object, a red patch on the wall for example. The fact that what I see is red is immediately evident, and seems to require no argument to support it as knowledge. I know it is red, and it seems that nobody could dispute that, and that therefore I would never have to justify it by citing evidence other than by repeating my claim that I see a red patch.

Pollock, however, showed that my claim to see a red patch in such a case is based on a defeasible argument. The argument runs as follows.

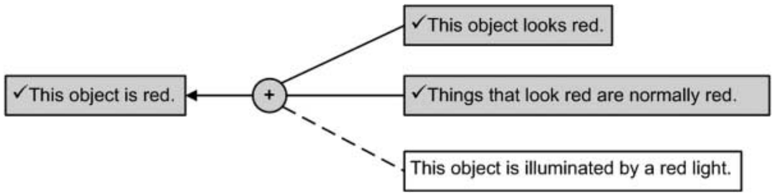

Figure 1. Pollock's red light example in the Carneades model.

Premise A1: This object looks red to me.
Premise A2: When an object looks red, then (normally, but subject to exceptions) it is red.
Conclusion A: This object is red.

The structure of Pollock's argument is displayed in the Carneades argument map (Gordon and Walton 2009) shown in Figure 1. Carneades is a computational model of argumentation consisting of mathematical structures and functions on these structures (Gordon and Walton 2009). Carneades models the structure and applicability of arguments, the acceptability of statements, burdens of proof, and proof standards. Carneades has been implemented using a functional programming language, and has a graphical-user interface: (http://carneades.github.com/).

The conclusion, the statement that the object is red, is shown in the text box at the left. The two text boxes on the right at the top represent premises that support the conclusion. The node in the middle containing the + indicates an argument leading from these two premises to the conclusion. The lines joining these two premises to the conclusion are unbroken, indicating that these two premises support the conclusion. They are premises that behave like assumptions, meaning that they are taken to hold. However, there is a third premise that appears in the text box with no fill. The dotted line joining this premise to the argument indicates that it is the type of premise classified by Carneades as an exception. An exception is a type of premise that is taken not to hold, although it can hold in special evidential circumstances. The motivation of the Carneades system is to explain what happens when a respondent asks a critical question (Walton and Gordon 2005). Carneades approaches this problem by distinguishing two types of premises (in addition to the stated premises of an argument) called assumptions and exceptions. An assumption holds if it is undisputed or accepted, but not if it is rejected. An undisputed ordinary assumption holds if its statement is acceptable, given its proof standard, or if it has been accepted, but not if it has been rejected. Exceptions hold unless the statement of the exception has been proven acceptable.

But why is this argument defeasible on Pollock's theory? The reason is that there is a counter-argument that can undercut the original argument by attacking the connection between the claim and the reason.

This type of evidential situation is shown in Figure 2, where the exception, the statement that the object is illuminated by red light, has now been accepted. If the special circumstances are such that the statement at the left in the unfilled box is known to hold, the argument is defeated.

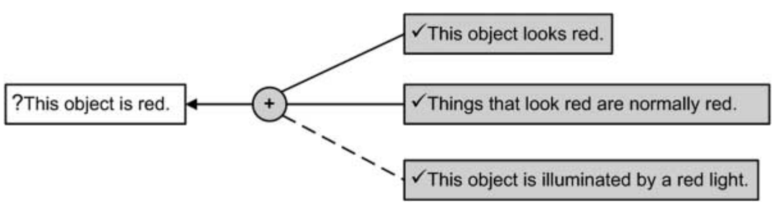

Figure 2. The red light example shown as an argument undercut by an exception.

Once it has been accepted, even though the argument retains the support given to it by the two previous premises, it now defaults once the exception has been activated. Then the argument, indicated by the round circle containing the plus mark, no longer has the support needed to prove the conclusion. The conclusion is now undercut by the exception.

According to Pollock's theory (1995, p. 41), the second argument is an undercutting defeater but not a rebutting defeater of the first one, because the second one is based on a defeasible generalisation, namely the proposition that red objects look red in red light too. The object may still be red, for all we know, despite the second argument above. What is shown is that there is an argument needed to support the claim to knowledge.

The precise nature of this argument has even been defined by Pollock's defeasible perception rule (1995, p. 41): having a percept with content φ is a *prima facie* reason to believe φ. Moreover, the undercutting defeater for any argument based on the perception rule can be formulated as follows (Bex et al. 2003, p. 38): 'the present circumstances are such that having a percept with content φ is not a reliable indicator of φ' undercuts the original argument. This rule has a function similar to that of a critical question matching an argumentation scheme. It leaves open a possibility that can be subject to doubt, and if questioned in response to an argument that fits the scheme, it can shift a burden of proof onto the proponent of the argument to respond to the doubt or give up the argument as a knowledge claim.

What has been shown is that the so-called basic propositions are based on argumentation needed to function as evidence supporting the claim to knowledge made. Such an argument needs to be based on Pollock's defeasible perception rule, the generalisation or warrant on which the inference to the claim depends. What can also be shown is that this form of argument is quite common in everyday conversational argumentation. An interesting legal example (Prakken 2002, p. 858) shows that it is common in legal argumentation as well.

Premise 1: This object looks like an affidavit.
Premise 2: If something looks like an affidavit, then it is an affidavit.
Conclusion: This object is an affidavit.

This argument is defeasible, for on a more detailed reading of the document, it might be found that it is not a real affidavit, but merely a forgery. Under the right conditions, however, it may be justifiable to go ahead on the reasonable assumption that the document is an affidavit without verifying that by having experts check it. It can be tentatively assumed it is an affidavit, because based on appearances and what else we know about the document, we can accept it, given that there is no reason to doubt its genuineness.

In Walton (2006), an argumentation scheme for this kind of defeasible argument has been formulated, called the scheme for argument from appearance.

It appears that this object could be classified under verbal category C. Therefore this object can be classified under verbal category C.

As shown by Pollock, and supported by ancient sceptical objections as well, this form of argument is best seen as defeasible rather than as conclusive. It is best evaluated on a balance of considerations, and should be seen as subject to sceptical doubt. According to the account given in Walton (2005), the scheme has the following matching critical questions attached to it.

CQ_1: Could the appearance of its looking like it could be classified under C be misleading for some reason?
CQ_2: Although it may look like it can be classified under C, could there be grounds for indicating that it might be more justifiable to classify it under another category D?

If either of these questions is asked in a given case, the original claim to knowledge needs to be suspended temporarily until the claimant to knowledge offers a satisfactory answer. Thus, even the most simple and straightforward claim to direct knowledge, like ‘I see a red patch now’, once analysed carefully in the Pollock manner, needs to be seen as a defeasible argument that does at least partly rest on external evidence that may need to be marshalled to support it.

3. The procedural view of inquiry

In this section, a new procedural view of inquiry is outlined that enables a clear formulation of the view that the idea of evidence supporting knowledge being defeasible makes sense. On this model, a proposition is classified as knowledge if it is accepted as supported strongly enough by the evidence to meet an appropriate standard of proof. In the model, a proposition p does not have to be true to be included in knowledge. Even though the external truth condition is dropped, however, there are still links with external reality. The acquisition of knowledge is seen as part of a procedure in which evidence is collected, tested, and measured against standards of proof.

The model thus supports a theory of knowledge that could be classified as pragmatic, in that it varies with the standards of proof appropriate for kinds of inquiry in a field of knowledge and with criteria for it to be considered to be evidence. According to the Carneades model, a group of interacting agents is collecting data as part of a search for the truth of a matter they are collaboratively investigating. As they go along during the search process, they verify or falsify hypotheses by testing them against the data they have collected so far, at the same time as they are engaged in the process of collecting new data. As the search for knowledge continues, some hypotheses become better supported by the evidence, but at the same time, some of the hypotheses previously accepted have to be given up, because they are falsified by the new data that are streaming in. Depending on the type of investigation, for example, it might be a scientific investigation or a legal inquiry, there will be an established proof standard that enables the investigation to determine whether a proposition can be accepted as proved or not (McBurney and Parsons 2001).

Such a pragmatic model of epistemic rationality is procedural, meaning that whether a proposition is accepted as knowledge at any given point depends on the standard of proof and the data that have been collected to that point. A proposition rightly said to be known to be true at a given point in the investigation could later on turn out to be proved to be false. Or at a particular point, the set of data collected at that point could justify two hypotheses, one of which is not consistent with the other. In this model, a particular proposition might rightly be classified as knowledge at one point in the investigation, whereas at a later point, the same proposition might turn out to be no longer classified as knowledge. In general, whether the proposition is rightly said to be knowledge or not depends on its rational acceptance, given the evidence then for it, as balanced against the evidence then against it, at that point in the investigation. In this pragmatic model, knowledge is not defined as justified true belief, or even as any kind of belief. It is based on the evidence collected at a given point in the investigation, on the kinds of arguments that can properly be used to justify a claim in that type of investigation, and on the standard of proof set for knowledge in that type of investigation. On this model, the strict barrier between discovery and verification of knowledge characteristic of older ways of thinking in analytical philosophy is no longer absolute.

Popper’s theory of knowledge is procedural, in that he holds that scientific knowledge advances towards the truth by improving tentative theories through a process of error reduction achieved by criticism and testing. This procedure of conjecture and refutation begins with the formulation of a problem P_1 and proceeds from there to a theory TT that is the conjectured solution to the problem. The next stage is that of error elimination (EE) that consists, in Popper’s words of a “severe critical examination of our conjecture”, or of several competing conjectures, if we have them, and a critical discussion that comparatively evaluates the competing conjectures (Popper 1972, p. 164). Finally, P_2 is the problem situation as it emerges from the process of testing the first hypothesis.

According to Popper (1972, p. 164), the whole procedure takes this form: $P_1 \rightarrow \text{TT} \rightarrow \text{EE} \rightarrow P_2$. This procedure repeats itself through successive refinements of the problem $P_1, P_2, \ldots, P_n$ so that progress can be gained in the movement towards finding the truth of the matter being discussed. On this approach, it is not a requirement for a proposition to be part of scientific knowledge that it be true. It is only required that it be accepted as true based on the evidence provided by testing and criticism, so that the procedure of scientific inquiry of which it is part is moving towards the truth. Popper (1962, p. 312) saw this procedure as a slow, steady, and continuous movement of trial and error that proceeds by successive degrees of improvement. Although he accepted the idea that the task of science is to search for truth, he conceded that on his view of the scientific method we may never get to a hypothesis that is true, or know that it is true when we get it (1962, p. 229). He accepted a variant of an old view that our knowledge is fallible.

This way of viewing the issue turns it into a problem of formulating conditions for the closure of the inquiry procedure in which the collection and processing of evidence should take place. At the opening stage, an appropriate proof standard for that type of inquiry needs to be set, and then a proposition can be said to be proved as acceptable by the investigation when the pro evidence supporting the proposition minus the contra evidence is sufficient to meet that standard. If the evidence supporting the ultimate proposition is strong enough so that the inquiry has reached its standard of proof at that point, then that proposition can be taken as known to be true based on the evidence for and against it that has been assembled and evaluated during the procedure. The degree of corroboration of a scientific theory, according to Popper (1972, 18), is determined by the critical discussion of the theory that has taken place, the degree of testability of the theory, the severity of the tests it has undergone, and the way it has stood up to these tests.

4. The Carneades model of inquiry

A dialogue system is a model of a sequence of exchanges in which two parties (or more) take turns making moves in an orderly manner engaging in speech activities such as questioning, argumentation and explanation in a rule-governed environment, in an orderly way as a transaction between the two participants. Formal dialogues are abstract normative structures that may be used to model argumentation in real dialogues, for example, parliamentary debates or scientific investigations. On the Carneades model, a formal dialogue is defined as an ordered 3-tuple $< O, A, C >$ where O is the opening stage, A is the argumentation stage, and C is the closing stage (Gordon and Walton 2009, p. 5). Dialogue rules (protocols) define what types of moves are allowed and how each type of move can or must be responded to. Commitment rules determine when and how insertions and retractions take place (Walton and Krabbe 1995). Each party has an individual goal and the dialogue itself has a collective goal. The seven basic types of dialogue are persuasion dialogue, discovery, inquiry, negotiation dialogue, information-seeking dialogue, deliberation, and eristic dialogue (Walton and Krabbe 1995). In an inquiry dialogue, the collective goal is to prove a designated statement, or if it cannot be proved by the evidence collected, to prove that it cannot be proved. The best argument standard is typical for deliberation dialogue, whereas a higher standard like reasonable doubt is typical for inquiry dialogue. During the closing stage it is determined, according to the standard of proof set at the opening stage, which party has won or lost the dialogue.

The goal of the inquiry type of dialogue is to prove that a statement designated at the opening stage as the ultimate claim is true or false, or if neither of these findings can be proved, to prove that there is insufficient evidence to prove that it is true or false (Walton 1998, chap. 3). To say a dialogue is *cumulative* means that once a statement has been accepted as true at any point in the argumentation stage of the inquiry, that statement must remain true at every succeeding point in the inquiry through the argumentation stage and to the closing stage. The cumulative argumentation in an inquiry is modelled by the (Kripke 1965) semantics for intuitionistic logic. This model has a tree structure, where the nodes are taken to represent evidential situations at a given point in an investigation in which more evidence comes to be verified. It can never happen, however, that a proposition is falsified and has to be retracted. For this reason, modelling inquiry as strictly cumulative is often called foundationalism, a view not compatible with retraction of an evidence-based commitment when it has been falsified by new evidence.

To build a more realistic model of inquiry, Black and Hunter (2007) built a formal dialogue system for use in medical domains where retraction is necessary because the database is typically incomplete, inconsistent, and has conditions of uncertainty. These conditions obtain in cases where many different health-care professionals cooperate by sharing specialised knowledge in order to provide care for a patient. Black and Hunter (2007, p. 2) explore inquiry dialogues in which strict cumulativeness is not required. This type of system has to have rules for the retraction of commitments (Walton and Krabbe 1995).

The argumentation stage *A* of a dialogue is made up of a sequence of moves, where each move *M* is an ordered pair <SpA, Con>, where *A* is the content of the move and SpA is a speech act representing the type of move whereby *A* was put forward in *D*. For example, there is a speech act for making a claim, or assertion as it can equivalently be called. The speaker can say 'I assert proposition *A*', and the commitment rule for assertions requires that proposition *A* be inserted into the speaker's commitment set. Generally, there is a rule regarding burden of proof such that whenever a speaker puts forward a speech act of this sort, and the hearer challenges it, requiring the speaker to back up his claim with some evidence, the speaker must either provide the appropriate evidence or give up his claim. It is a problem in formal dialogue systems that this rule does not always apply to all assertions. Sometimes a participant may assert statement hypothetically as a hypothesis, even though he cannot presently prove it. However, if the speaker makes a move claiming that he knows proposition *A*, then he not only has to give evidence to back up this claim, but the evidence has to meet a standard of proof that is high enough to sustain a claim to knowledge. In dialogue systems, there is a way to distinguish between knowledge claims, and weaker kinds of claims that require less strong supporting arguments to back them up when they are challenged.

In an adequate model of this sort, knowledge should be defeasible in a way that allows for external evidence from reality to have some role in the procedure, and attainment of truth is not necessary for the procedure to decide whether a proposition is knowledge or not. Instead, what determines whether a proposition is knowledge is the weighing of the evidence for and against it, based on the collection of evidence that has been marshalled during the procedure. For the proposition to be knowledge, the evidence for it has to be stronger than the evidence against it, to a degree stipulated at the opening stage of the inquiry. This degree of strength that is required is called the standard of proof. The standard of proof needs to be high, in order to avoid the practical possibility of a later need to retract the proposition that was claimed to be known to be true.

In this model of defeasible knowledge, the product of the procedure can be called knowledge, and in the Carneades model, the criteria used to judge something as knowledge are linked to the procedure that generates the knowledge as its end state. However, this model does not merely represent a "consensus-based approach", for several reasons. First, the reasoning has to be based on external evidence, so that claim to knowledge can be tested by using this external evidence. Such evidence can come in, for example, by observation or by testing of a hypothesis. Second, if a proposition *p* is known to be true, it follows by defeasible logic that it must be accepted as true, unless it can be shown by additional evidence that it is false. In other words, through the argumentation scheme for the argument from ignorance, a proposition can be concluded to be

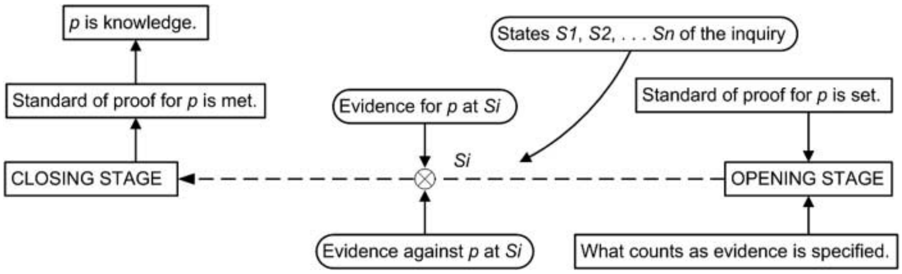

Figure 3. Procedure for evaluating a defeasible knowledge claim.

known to be true provided it has not been proved to be false, subject to the additional assumption that the knowledge base is complete enough to warrant this inference. If the knowledge base is complete enough for the argument from ignorance to meet its standard of proof and go through as an acceptable defeasible argument, the appropriate conclusion can be drawn. The inferential procedure that takes us from the evidence pro and contra a knowledge claim for a proposition *p* to the conclusion that *p* is knowledge (or not) is shown in Figure 3.

As shown in Figure 3, as new evidence comes in, a claim to knowledge should be subject to testing by marshalling evidence both for and against it. As this process is underway, what is knowledge at a given point must yield to new knowledge, as the old knowledge is undercut or defeated by new evidence. This part represents the progress of defeasible knowledge underneath shown by the dotted arrow in Figure 3. As the process of accumulating knowledge is underway, objections and refutations will lead us to reject propositions that we formerly classified as knowledge. They really were knowledge at that time, but now they no longer are. The need for us to reject propositions that are shown to be false by the evidence that is coming in, based on rational argumentation standards appropriate for the inquiry, is required by the goal of the process itself, which is that of obtaining the truth, or at least coming as close to it as we can get, as well as avoiding falsehood, fallacies, dismissal of evidence, other faults of rational inquiry.

It can be argued that this pragmatic model of knowledge is more useful than the prevailing epistemological one, because it arguably shows better how the external standard should be applied to reality in an inquiry (McBurney and Parsons 2001).

It is not possible to set a single standard of proof for every scientific investigation. There is also the question of the conditions under which an inquiry should be reopened. From a fallibilist point of view, it is unrealistic to set a standard of beyond all doubt, and it is necessary even if one wants to set a very high standard, like that of beyond reasonable doubt (BRD), to leave open the possibility that the inquiry can be reopened for further scientific investigation bringing in new evidence. This assumption is based on the defeasibility of scientific knowledge, which is in turn based on falsifiability as a criterion of genuine scientific knowledge.

The procedure outlined in the model represented in Figure 3 aims towards the truth by avoiding the errors revealed during the argumentation brought in by the evidence applied to the proposition claimed to be knowledge during the procedure. In an inquiry to assess knowledge, the goal is to prove that some designated proposition is true or false, or otherwise to determine that it cannot be proved to be true or false based on all the evidence that has been collected. The assumption is that enough evidence has been collected so that this conclusion can be established as meeting an appropriate standard of proof. The following four standards of proof used in Carneades (Gordon and Walton 2009) can be used to give the reader an idea how appropriate proof standards can be set in the kind of argumentation procedure outlined in Figure 3.

- Scintilla of evidence (SE) is met if there is at least one applicable argument for a claim.
- Preponderance of the evidence (PE) is met if SE is satisfied and the maximum weight assigned to an applicable *pro* argument (for the claim) is greater than the maximum weight of an applicable *con* argument (against the claim).
- Clear and convincing evidence (CCE) is met if PE is satisfied, the maximum weight of applicable *pro* arguments exceeds some threshold α, and the difference between the maximum weight of the applicable *pro* arguments and the maximum weight of the applicable *con* arguments exceeds some threshold β.
- BRD is met if CCE is satisfied and the maximum weight of the applicable con arguments is less than some threshold γ.

When I say that I know something, or say that it is knowledge, it should mean that I have strong enough evidence to support and to meet a standard of proof justifying my including it under the category of knowledge. To say that something is knowledge, it is important that the proposition claimed as knowledge be based on evidence of a kind that reaches a level where the proposition passes beyond the level of being accepted as true because it is based on evidence. Only when it is proved by a certain kind of evidence, that is sufficient for the discipline, or more generally the context in which the proposition was claimed, can something be properly said to be knowledge.

The standard has to be high enough in a scientific inquiry to minimise the possibility that the proposition accepted as true will later have to be retracted. On the model, however, since epistemic reasoning that results in knowledge is inherently defeasible, the possibility of retraction can never be excluded entirely. Even once the procedure is closed off, and the conclusion is accepted, there is always the possibility that it can be re-opened, should new evidence come to light.

The Carneades model of inquiry makes sense of the defeasible notion of evidence-based knowledge that does not require the four epistemic reasoning principles set out in Section 1 as parts of the method for proving a proposition to be an item of knowledge.

5. Defeasible logic

In the model, the statement "*p* is known to be true" does not deductively imply the statement "*p* is true". But a comparable inference holds: "*p* is known to be true" defeasibly implies the statement "*p* is true". To say that *p* defeasibly implies *q* is taken to mean that all else being equal, subject to exceptions, if *p* is true, then *q* is true. This defeasible relationship between knowledge and truth does imply a link with external reality. Hence, the model does not entirely forsake external standards. Truth is not guaranteed to be produced by the procedure outlined in the model, but defeasible support for what should be classified as knowledge, based on the evidence, can be. The reason is that the arguments based on evidence that are used to support and attack a claim to knowledge are based on external reality. For example, they can be based on the scheme for argument from appearances. This form of argument, as shown in Section 4, is defeasible, but it is also based on reality as we know it, through our perception of it.

Defeasible logic (Nute 1994) is a rule-based non-monotonic formal system that models reasoning used to derive plausible conclusions from partial and sometimes conflicting information. A conclusion derived in such a system is only tentatively accepted, subject to new information that comes in later, requiring its retraction (Simari and Loui 1992). The basic units of the system are facts and rules. There are two kinds of rules, strict rules and defeasible rules. Facts are indisputable statements that are accepted as true within the confines of a discussion. Statement are denoted by letters, *A*, *B*, *C*, ..., and so forth, using subscripts if we run out of letters. Strict rules are rules in the classical sense: whenever the premises are indisputable (e.g. facts) then so is the conclusion, e.g. “Penguins are birds”. A strict rule has the form of a conditional, $A_1, A_2, A_n, \ldots, \rightarrow B$, where it is not possible for all the A_i to be true and the B false. Defeasible rules are rules that can be defeated by contrary evidence, e.g. ‘Birds fly’. A defeasible rule has the form of a conditional, $A_1, A_2, A_n, \ldots, => B$, where each of the A_i is called a prerequisite, all the A_i together are called the antecedent, and B is called the consequent. One source of defeat for the defeasible conditional is that it is open to exceptions, e.g. “This bird is a penguin”. One rule can conflict with another. A priority relation defined over the set of rules that determines the relative strength of any two conflicting rules.

One purpose of defeasible logic is to resolve a conflict of opinions, but it can also be used to model epistemic reasoning where a knowledge base is incomplete, and hence where we have to reason on the basis of what is not known, as well as on the basis of what is known. In cases where there is some support for concluding A but also support for concluding $\sim A$, defeasible logic determines which conclusion is drawn by using a priority relation. For example, priority can be given to what is better known. If the support for A has priority over the support for $\sim A$, the conclusion to accept A is drawn. Defeasible logic can also be used with temporal reasoning (Riveret et al. 2006). For example, a proposition known at one time may not be known at another time. A definite conclusion is a conclusion that cannot ever be retracted, even if new information comes in. A defeasible conclusion is only a tentative conclusion and might have to be retracted if new information comes in. In addition, defeasible logic is able to tell whether a conclusion is or is not provable.

It is possible to have four types of conclusions (Governatori 2008):

- Positive-definite conclusions: meaning that the conclusion is provable using only facts and strict rules.
- Negative-definite conclusions: meaning that it is not possible to prove the conclusion using only facts and strict rules.
- Positive-defeasible conclusions: meaning that the conclusions can be defeasibly proved.
- Negative-defeasible conclusions: meaning that one can show that the conclusion is not even defeasibly provable.

A defeasible conclusion A can be accepted if there is a rule whose conclusion is A, whose prerequisites are facts, and any stronger rule whose conclusion is $\sim A$ has prerequisites that fail to be derived.

The reasoning process can be explained in terms of argumentation. To prove a conclusion you have to carry out three steps (Governatori 2008).

(1) Give an argument for the conclusion to be proved.
(2) Consider all possible counter-arguments for the conclusion.
(3) Defeat these counter-arguments by either showing that some premises in each of them do not hold, or defeat each of them by producing a counter-argument with a stronger argument supporting its conclusion.

Defeasible logic is most useful where there is a procedure that has the goal of proving or disproving some claim at issue that moves forward by bringing forward the pro and contra arguments with respect to the claim. A conclusion is proved as the outcome if there is an argument supporting it and all the arguments against it are defeated.

An important component of defeasible logic is the notion of a defeater (rebuttal) of an argument. A defeater might be thought to be a rule and a set of facts that proves a conclusion that is the opposite of the original argument. But this definition seems too simple, suggesting the following more complex definition. On this definition, a defeater is a counter-argument of three types directed to a prior argument that has already been put forward. It can be a counter-argument that shows that one of the prerequisites (premises) of the original argument does not hold. It can be a stronger argument that proves the opposite conclusion of the original argument. Or it can be an argument that challenges the inference from the premises to the conclusion.

Simari and Loui (1992), Verheij (1999, p. 115), and Walton (2002, p. 43) have put forward the proposal that many common argumentation schemes fit under a defeasible form of the deductive form of *modus ponens* that we are familiar with in deductive logic. The normal *modus ponens* form of argument is based on the material conditional binary constant → sometimes called strict implication. The variables p, q, r, . . ., stand for propositions (statements).

Major Premise: $p \rightarrow q$
Minor Premise: p
Conclusion: q

This form of argument can be called strict *modus ponens* (SMP). In contrast, there is also a defeasible *modus ponens* (DMP) having the following form, where the symbol => is a binary constant representing the defeasible conditional.

Major Premise: $p => q$
Minor Premise: p
Conclusion: q

This form of argument is called DMP in (Walton 2002, p. 43).[2] To cite an example, the following argument arguably fits the form of DMP: if something is a bird and generally, but subject to exceptions, it flies; Tweety is a bird; therefore, Tweety flies. This argument is the canonical example of defeasible reasoning used in computer science. Suppose we find out that Tweety has a broken wing that prevents him from flying, or that Tweety is a penguin, a type of bird that does not fly. If we find out that in the given case one of these characteristics fits Tweety, the original DMP argument defaults. The argument is best not seen as one that is deductively valid, and that still holds even if new information comes in showing that the argument no longer applies to the particular case in the way anticipated. Instead, it is better seen as an argument that holds only tentatively during an investigation, but that can fail to hold any longer if new evidence comes in that cites an exception to the rule specified in the major premise.

The current trend in applications of defeasible logic in artificial intelligence is to sanction a defeasible form of *modus ponens*, but not to sanction any form of *modus tollens* (Caminada 2008, p. 111). Examples of two systems of defeasible logic that follow this pattern are (Prakken and Sartor 1997) and (Reiter 1980). Caminada (2008) draws a distinction between epistemic reasoning, which is supposed to be based on an objective reality that can support a claim to knowledge, and other kinds of reasoning where contradictions represent soft conflicts that can be dealt with by prioritising defeasible rules. On Caminada's view, *modus tollens* applies to defeasible epistemic reasoning, because it is characterised by hard conflicts. If the second proposition is inferred defeasibly from the first proposition, and the second proposition exhibits a hard conflict so that it cannot be the case, it follows that the first proposition cannot be the case either. However, the same inference principle does not work for soft conflicts, for example, in legal and ethical reasoning, where dilemmas and conflicts of rulings can occur.

6. Reasoning from absence of knowledge

It is clear that defeasible logic is closely related to argumentation schemes of the kind that represent reasoning from knowledge and perception. Two motivating examples given by Nute (2001, p. 89) make some significant connections with such forms of epistemic reasoning. The first example [quoted below] is a typical case of the lack of knowledge argument, or argument from ignorance, as it is more often known.

> The absence of information can sometimes be a positive reason for believing something. Is there any milk in the refrigerator? We look and we do not see any milk. The failure to find evidence of milk in this case is a good reason to believe that there is no milk in the refrigerator.

The second example [quoted below] is a typical case of argument from appearance.

> For another example, I believe that there is a cat in front of me. I believe this because there appears to be a cat in front of me. That seems to be ample evidence. Of course, we can think of situations where I would be wrong. I might be hallucinating, or there might be a hologram of a cat, or there might be a mirror and the cat I think I see in front of me is actually behind me. But I have no reason to believe that I am hallucinating, and there is no evidence of a holographic projector or of a mirror. The absence of evidence that my perceptual circumstances are abnormal provides part of the justification for my belief that there is a cat in front of me.

This example is suggestive, not only in relation to traditional epistemological issues arising from scepticism, but also because part of the structure of the chain of argumentation is an instance of RAK. In the field of argumentation studies, RAK is represented by the so-called argument from ignorance, traditionally taken to be a fallacious form of reasoning. However, recent studies have shown that non-fallacious uses of it are very common.

Indeed, this form of reasoning, called the closed-world assumption (Bondarenko et al. 1993), means that all the information that there is to know or find is listed in the collection of information one already has. The closed-world assumption is met if all the positive information in a database is listed, and therefore negative information is represented by default (Reiter 1980, p. 69). Reiter (1987, p. 150) offers the example of a database for an airline flight schedule. It would be too much information to include in such a database to list all flights and also all city pairs they do not connect. The closed-world assumption tells us that if a positive flight connection between a pair of cities is not listed, the conclusion can be drawn that there is no flight connecting these two cities. This form of argument used in this sequence of knowledge-based reasoning implies that failure to find a proof has sanctioned an inference (Kakas and Toni 1999; Toni 2008). Such a use of the closed-world assumption seems to make it equivalent to argument from ignorance, but before going further let us define argument from ignorance.

The argumentation scheme for argument from ignorance has two premises in which knowledge, or the absence of it, plays a reason-giving role (Walton 1996, p. 254).

> *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).

The scheme for argument from ignorance looks to be the same form of reasoning as RAK, which in turn is the same kind of default reasoning expressed by the closed-world assumption. The last sentence of Nute's example of an argument from appearance is case in point.

The conditions under which RAK is most typically fallacious can be formulated by distinguishing between two forms of reasoning. The fallacious reasoning takes the form $\sim \mathrm{K}p \Rightarrow \sim p$. Here, there is a leap from the absence of knowledge to falsity without taking the conditional premise into account. The non-fallacious instances of this kind of reasoning take a different form: $(\sim \mathrm{K}p \;\&\; (p \Rightarrow \mathrm{K}p)) \Rightarrow \sim p$. Once the depth of search is taken into account by the conditional premise, the inference now clearly has a *modus tollens* form.

Whether the conditional premise holds in a given case depends on how complete the knowledge base is in that case. For example, suppose the question asked is whether Guyana is a major rubber producer in South America, and the knowledge base we have about rubber producers in South

America contains a lot of knowledge about this subject. If Guyana were a major producer, that knowledge would almost certainly be contained in the knowledge base. Suppose we look through the knowledge base, and the proposition "Guyana is a major rubber producer" is not known to be true. We could then conclude, on the basis of an argument from ignorance, that Guyana is not a major rubber producer. Thus argument from ignorance, or argument from absence of evidence as it might less prejudicially be called, is often reasonable even though we must be careful to realise that it needs to depend on what is known as well as what is not known.

Since RAK has the *modus tollens* form, this takes us back to the issue of whether *modus tollens* should be taken to hold for epistemic reasoning. This sequence of reasoning has the following *modus tollens* form: if p then q; $\sim q$; therefore $\sim p$. This example seems to be convincing (Caminada 2004, p. 87), but other authors do not accept defeasible contraposition for *modus tollens*. Brewka (1989) offers this counterexample: men usually do not have beards, but this does not mean that if someone does have a beard, it is usually not a man. In this example, we have the statement that if a person is a man then he usually does not have a beard. The contraposition of the sentence will be the statement that if a person does have a beard then usually that person is not a man. If we had the statement that this person has a beard, and we accept the contraposition of defeasible rules, we can derive the statement that this person is not a man. This conclusion seems wrong. Hence, we have to realise that defeasible rules are different from strict rules based on deductive logic. In the case of a strict rule, we know that whenever the antecedent holds the consequent also holds. Therefore, we know that when the consequent does not hold we can be sure that the antecedent does not hold either. Defeasible rules are quite different. A defeasible rule only says that if you know that the antecedent holds you have some inclination to believe that the consequent holds. Hence, as Brewka's example shows, the relationship between the negated consequent and the negated antecedent could be very different.

Modus tollens is closely related to contraposition. Contraposition for defeasible implication is the rule of inference '$(p => q) \Leftrightarrow (\sim q => \sim p)$', where $\Leftrightarrow$ stands for mutual defeasible implication, $(p => q)$ & $(q => p)$. Mutual defeasible implication is a form of equivalence for defeasible reasoning, enabling us to replace a formula wherever it occurs with a defeasibly equivalent formula. Defeasible *modus tollens* is the inference from premises $p => q$ and $\sim q$ to the conclusion $\sim p$. If we assume that defeasible *modus ponens* holds, and as well that defeasible contraposition holds, defeasible *modus tollens* follows. *Modus ponens* is this form of inference:

$p => q$
p
Therefore q

By contraposition we can replace the first premise with $\sim q => \sim p$. Then by double negation $(p \Leftrightarrow \sim\sim p)$, which we assume holds, we can change the second premise to $\sim\sim p$, and we can change the conclusion to $\sim\sim q$. The resulting inference looks like this.

$\sim q => \sim p$
$\sim\sim p$
Therefore $\sim\sim q$

This inference is clearly a species of defeasible *modus tollens*. In other words, if double negation, defeasible *modus ponens* and defeasible contraposition hold, then so does defeasible *modus tollens*.

In the next section, we will show that there are some distinctions to be made between the closed-world assumption and argument from ignorance. To sum up what we have seen so far, ignorance (lack of knowledge) is as important as knowledge for understanding epistemic reasoning, and the various forms of reasoning used in transitions from ignorance to knowledge need to be understood and modelled in any practically realistic theory of knowledge. Although the models of inquiry of Peirce and Popper postulate that verisimilitude is an important property of the successful marshalling of evidence in this kind of procedure, it needs to be pointed out that the property of moving away from ignorance is equally important in judging the success of the inquiry as a whole. Especially as stressed by Popper, inquiry moves forward through a process of rational criticism that reveals errors and fallacies in arguments and hypotheses that were previously accepted, but are now revealed as simplistic theories or even superstitions based on ignorance. In addition to moving towards the truth, a successful inquiry also moves away from falsehood and error.

7. Defeasible knowledge, evidence and lack of evidence

A perennial problem for the theory of defeasible knowledge concerns statements made in everyday conversational argumentation where somebody says "I know this proposition is true" as opposed to saying "I think it is true", because it is so obvious that there can be no doubt about it. For example, if I look at my two hands in front of me and say "I see two hands in front of my face", I can say I know that there are two hands in front of my face, because I am confident that no amount of new evidence is ever going to make this proposition false. When I see that this proposition is true, it can be taken to mean that I am so confident that it is true that there can be no doubt about it, and therefore I know it is true. The fact is immediately evident. So it needs no evidence, of the kind required by the new model of defeasible knowledge, to justify it.

The speaker in Nute's example says that he believes there is a cat in front of him, and the reason he gives for this belief is that he believes there appears to be a cat in front of him. But when we come to examine how his beliefs can be justified, and we analyse it carefully, it can be seen to be based on an implicit premise. He needs to base his view on evidence, if he were to be challenged to support by someone who is sceptical. As we will show below using the new defeasible model of knowledge, although the speaker does not say so explicitly, this belief is also based on lack of knowledge. As noted in the earlier discussion of the example in Section 2, it can be well represented as an instance of the argument from ignorance or RAK.

The argumentation in Nute's example is represented using the Carneades model in Figure 4. On the left there is a conclusion that there is a cat in front of me. At the top, in the box in the middle at the top with a checkmark in front of it, is the evidence. This evidence is the proposition that there appears to be a cat in front of me. This argument is an instance of the argumentation scheme called argument from appearance. As noted in Section 2, Pollock saw this kind of reasoning as based on his defeasible perception rule.

As shown in reasoning. The two propositions in Figure 4 that have an × in front of them are modelled in Carneades as exceptions. An exception represents a kind of critical question that only defeats the given argument if there is evidence to back it up. For example, if there were evidence that I might be hallucinating, that would undercut the argument from the premise that there appears to be a cat in front of me to the conclusion that there is a cat in front of me. Similarly if there is

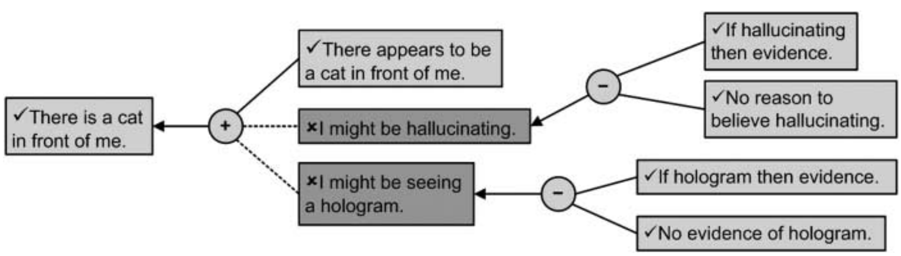

Figure 4. Carneades model of the Cat example.

evidence that I might be seeing a hologram, that would undercut the argument that there appears to be a cat in front of me to the conclusion that there is a cat in front of me. Next, look at the text boxes on the right side of Figure 4. Each of the two arguments displayed on the right is an RAK argument, fitting the argumentation scheme for the argument from ignorance. Each of these two arguments as checkmarks inside the text boxes representing their premises. This means that these premises have been accepted. But the circle representing each of these two arguments has a minus sign contained in it. This indicates that each of the arguments is a contra argument that rebuts the conclusion that it leads to.

What is shown in Figure 4 as a whole is that the evidence standing behind two exceptions that would defeat the original defeasible argument is not there. So what is shown is that the absence of the evidence required to defeat the original argument is actually a reason that shows why it is not defeated. And these evidential conditions show that I am justified in believing that there is a cat in front of me. Given the strength of the original argument from appearances, taken along with the absence of any reason to doubt the veracity of these appearances, I am justified, on balance, in saying that I know there is a cat in front of me. Even though this knowledge claim is defeasible, the evidence displayed in Figure 4, along with the absence of evidence included as part of the evidential situation, justifies the claim that there is a cat in front of me so strongly that I can justifiably say that I know there is a cat in front of me.

When using Carneades to represent an argument, standards of proof have to be inserted for all the propositions that function as premises and conclusions in the argument. In order to reduce the complexities of modelling the example, standards of proof have not been inserted. All that has been represented in the diagram is the structure of the argumentation, and whether each proposition that is a premise or conclusion is accepted or rejected. When using the Carneades graphical-user interface, the user has to insert a standard of proof, like the preponderance of evidence standard for example, meaning that the evidence pro the claim has to be stronger than the evidence contra the claim. The user is also asked to indicate whether each of the premises in the chain of argumentation has been accepted, rejected, questioned, or stated. Ordinary premises, for example are stated, and therefore initially appear in a box containing no colour fill or mark. But once such a premise has been accepted, a checkmark appears in the box containing it. All this information is inserted by the user into the Carneades menu. Once the status of all the premises and conclusions in the chain of argumentation leading to the ultimate conclusion has been indicated, Carneades automatically indicates the status of the ultimate conclusion, depending of course on what the standard of proof is.

Finally, in this section there needs to be a more carefully drawn distinction made between the argument from ignorance and the closed-world assumption. It can be observed that the closed-world assumption does not work well in knowledge bases that include sentences from classical logic. Consider a knowledge base that contains the disjunction $\{A \vee B\}$. Since neither A nor B can be derived from $\{A \vee B\}$, using the closed-world assumption it is possible to derive both the negation of A and the negation of B. From that, using the rule of disjunctive syllogism in classical logic, both A and B can be derived, making the database inconsistent. This result shows that using classical logic along with the closed-world assumption is problematic in a formal system of knowledge inquiry. This result suggests that a careful distinction needs to be drawn between how RAK and the closed-world assumption are applied when modelling the structure of inquiry. As a hypothesis we propose here that one way to deal with the problem would be to only apply the closed-world assumption at the closing stage of the inquiry, and only apply RAK during the argumentation stage. Typically, the argument from ignorance works as a device for shifting burden of proof from one side to the other. During this sequence of speech acts in the argumentation stage when one party makes an assertion the other party requests that the other party give some evidence to support her assertion. This sort of argumentation exchange is typically the kind of case in which the argument from ignorance poses a problem. On the other hand, the closed-world assumption is a vitally important device for marking the point when the argumentation stage should be closed off so that the arguments on both sides can be evaluated to determine which side meets that standard of proof.

8. Objections and replies

In this section we revisit the conflict between the two views of reasoning about knowledge described in Section 1. We need to better understand how either trying to maintain or refute this view makes for a philosophical difficulty, and puts us at odds with ordinary language. We often refer to the contents of the various scientific disciplines as being knowledge of the kind that we possess. However, it seems that these sets of propositions are not really knowledge at all, in the sense required by the four epistemological principles stated in Section 1. The reason is that we cannot *know* that the proposition established by the various scientific disciplines are true, at least if we mean by "know" that we can prove them by such a high standard of proof that there remains no possibility at all, even a logical possibility, that they will be retracted and withdrawn as knowledge at some time in the future. This view has the unfortunate consequence that scientific knowledge is not defeasible and, that is, as we have seen, in conflict with even a weak form of fallibilism. This view represents one horn of a dilemma. The other horn is to take the fallibility approach and embrace the claim that these sets of propositions current in the various scientific disciplines legitimately qualify as being knowledge in a different sense of the word. This fallibilistic view also has a consequence that does not square very well with ways we commonly speak about knowledge. Suppose we have a proposition that is presently accepted as part of established scientific knowledge, but new scientific results show that there is evidence against it, leading the scientists in the field to reject it. On the Carneades model of the inquiry, we can describe the situation by saying that this proposition was formerly knowledge, but is no longer knowledge. To many, it would appear that this way of speaking is not right. They would say that it is much more natural to describe the situation by saying that we thought this proposition was knowledge, but then we found out it was not really knowledge after all.

Although there are the beginnings of a shift away from the traditional epistemological view of knowledge represented by the four principles stated in Section 1 towards some form of fallibilism, neither view is philosophically unproblematic. The conflict between the two views is one of those perennial philosophical problems that will most likely always be with us. So far, the position represented by the four principles is the most widely accepted view in current epistemology. However, the fallibilistic view is a very old one, and can claim two well-known recent philosophers as supporters, Peirce and Popper.

Knowledge statements in everyday language assertions are often ambiguous. A knowledge statement could mean that I am claiming that a proposition is scientific knowledge, or it could mean merely that I am saying I am personally very confident about its being true, independently of whether I have good evidence that it can be taken as scientific knowledge. Another tricky aspect is the shift between the abstract noun knowledge and the verb "know", where the latter may have a much stronger suggestion of personal conviction or belief. The Carneades model of the inquiry postulates an orderly procedure in which a knowledge claim is evaluated by evidence that is collected and tested. But in philosophical discussions of its applicability, we always seem to come back to the same old problem of how it fits with the ways we ordinarily speak about knowledge. What happens, for example, in the history of science where something was taken as knowledge at one point, but later a new theory was developed or some new finding was introduced that defeated this claim to knowledge? From the defeasible knowledge inquiry viewpoint, it was really knowledge at the earlier point, even though it was subsequently defeated so that at the next point some other proposition that was inconsistent with the first one now came to be taken as knowledge.

This view seems paradoxical to some. How can it really have been knowledge before a given point if it was defeated after that point? Surely it is more natural to say that it was not really knowledge after all. It only seemed to be knowledge at that time, but at the later time it can no longer be said to be knowledge.

On this objection there is an alternative description of the situation which is better and removes the inconsistency between saying that something both is and is not knowledge. On this alternative description, we were justified in saying in believing that the proposition in question was knowledge at the earlier point, but that does not mean that really was knowledge. It only means that we were justified in believing it to be knowledge at that point, even though later on it turned out that it was not really knowledge after all. On this alternative view, a proposition may be justifiably believed to be knowledge at a particular point as an investigation proceeds, but the claim that this proposition is knowledge may subsequently be defeated when the proposition is falsified by new evidence. According to this theory, the proper description of the situation is that this proposition was not knowledge at the earlier point in the investigation, even though it was justifiably believed to be knowledge at that point, based on the evidence available at the time.

The Carneades model of the inquiry embraces the assumption that knowledge really is defeasible. It is not only saying that claims to knowledge are defeasible or that what we are justified in believing to be knowledge is defeasible. It rejects the assumption that knowledge always implies truth, because that would take with it the consequence that anything proven false could never have really been knowledge at some earlier point in an investigation, even though we were previously justified in accepting it as knowledge. The Carneades model of the inquiry adopts the view that what is or is not knowledge can change over time. In particular, it allows that something that was knowledge at an earlier time can cease to be knowledge at a later time. The competing approach does not allow for this possibility. Once something is knowledge, it must always remain knowledge forever, during an infinite inquiry, as Peirce showed.

One possibility is that the two approaches model different kinds of knowledge. For example the truth-implying notion of knowledge may best model mathematical epistemic reasoning, whereas the defeasible knowledge approach best models epistemic reasoning in the experimental sciences. However, Peirce even extended fallibilism to mathematical knowledge. He asked if you would bet your life against a penny on the truth of some observational or mathematical statement that you do not in fact doubt. We may be hesitant about observational statements in this regard, but there are reasons for extending this hesitancy to mathematical statements as well. Even the greatest mathematicians, he observed, are susceptible of making a small arithmetical mistake because of a little lapse of attention. He concluded that no rational person would make these kinds of bets on the truth of mathematical statements because "you could not go on making very many millions of such bets before you would lose!" (Peirce 1931, p. 1.150).

Some might think that the fallibilistic view of reasoning about knowledge advocated in this paper implies a pernicious relativism that allows us to say that we can be taken to know all kinds of propositions that are not true. In reply to this objection, it needs to be stated that the inquiry procedure is evidence-based and requires the principle of the falsifiability of knowledge claims. On this model, if a proposition is found to be false, based on sufficient evidence to meet the appropriate standard of proof to support this finding, it has to be retracted. To make this view explicit, some clarifications of its logical implications are helpful.

If we find out that a proposition is false, then we cannot say that it is known to be true. Indeed, if we find out if it is false, by means of some evidence, but before that we had thought was true, then we have to give up the assumption that it is known to be true. Essentially, this statement represents the principle of falsifiability of scientific knowledge proposed by Popper (1963). But it does not follow from this principle of falsifiability that if a proposition is known to be true, then it is true. Even though a proposition is known to be true, it does not follow necessarily that it is true.

In fact, in order for the principle of falsifiability to be tenable, the notion that knowledge implies truth has to be given up. For even though a proposition is part of scientific knowledge, according to the principle of falsifiability, all genuine scientific knowledge must be falsifiable. That is, the reasoning on which it rests must be defeasible.

Consider once again Pollock's example of seeing a red light. Or for that matter, take any statement of common sense knowledge like G.E. Moore's example, "I see a hand in front of my face". I may be very confident that I am seeing a red light, or that I am seeing a hand in front of my face, so that I am entitled to claim in everyday conversational discourse that I know I see a red light, or that I know I see a hand in front of my face. It depends on what the appropriate standard of proof is for the situation in which I am seeing a red light or seeing hand in front of my face, and what the consequences of being wrong might be. But as Pollock noted, if the room is illuminated by red light, since everything looks red when illuminated by red light, the light I am seeing may not be red. For all the various reasons advanced by sceptics since ancient times, as indicated in Section 1, even though I might be very confident that I am seeing a hand in front of my face, so much so that I can even say that I know I see a hand there, it does not follow by deductive necessity that the light is red, or that the hand is there. The reason is that, although it may seem very unlikely, the possibility is there that I am mistaken.

As shown in Section 1, Rescher, outlining the traditional epistemological view of knowledge, wrote that holding that the linkage between knowledge and truth is merely contingent does violence to the way we use the term "knowledge" in everyday discourse. When you make a knowledge claim, you are claiming that a particular proposition is true. If I say that I know this particular proposition is true, it follows that I am claiming that the proposition is true. But as shown in Section 1, this observation does not show that knowledge deductively implies truth. It does not mean that if you know a proposition, it has to be true. It does not follow that the proposition has to be true. It only follows that the person holds it to be true. On the Carneades model, what follows is not truth but acceptance: if a person says that he knows that a proposition is the case, it follows that he must be taken to accept this proposition. What does not follow is that the proposition itself must actually be true, in a sense that it can never be rejected as false once the inquiry has proceeded further.

9. Conclusion

By linking knowledge to an orderly process of inquiry, and by recognising the role that evidence plays in that process, the model represents knowledge as defeasible. The main characteristics of the new model of defeasible knowledge can be summarised as follows.

(1) Knowledge is the result of a process of inquiry in which existing knowledge can continue to be tested, and is subject to retraction. Knowledge must not only be undefeated at present, but must be able to answer new objections, and to survive testing as new evidence comes in.

(2) Knowledge bases can be incomplete and even inconsistent during the procedure where a claim to knowledge is being investigated, but it should be a requirement of knowledge of the kind accepted at the final stage in the defeasible model that the knowledge be consistent.

(3) The reasoning on which knowledge is based is defeasible rather than deductively valid. As such (a) certainty beyond all doubt is an inappropriate standard for knowledge and (b) knowledge claims must be retractable under the right sorts of circumstances.

(4) Because the correction that comes in from the evidence supporting the pro and contra arguments is based on perception of an external reality, there is no need to require an external standard of truth for a proposition to be accepted (or rejected) as knowledge.

(5) The veracity condition K$p \rightarrow p$ is replaced with the weaker assumption Kp => p.
(6) Knowledge bases are not closed under deductive implication. Nor are they closed under defeasible implication. The assumption K(p => q) => [Kp => Kq] fails to hold.
(7) Iteration for defeasible implication Kp => KKp fails to hold.
(8) Whether or not a proposition is rightly classified as knowledge or not depends on evidence of the right kind appropriate for the investigation, and the standard of proof set for the investigation.

The model maintains some but not all of the elements of a justified true belief analysis. It maintains the element of justification, but models justification using defeasible logic. Defeasible logic admits both of strict deductive implication and defeasible implication, but on this model it structures the process of evaluation of knowledge claims using defeasible implication. However, it does not entirely exclude the use of deductive implication in some instances.

It is concluded that, as opposed to the way of defining knowledge as a species of true belief, the model provides a coherent and defensible model of knowledge as a set of propositions based defeasibly on evidence and supported to a specified standard of proof. The specified standard can be set sufficiently high to distinguish between propositions that can rightly be said to be knowledge versus propositions that can be reasonably accepted as based on evidence, but where the evidence is not so strong that we should call the proposition knowledge. An important implication of the new model is that it puts much more of an emphasis on the notion of evidence than one finds in traditional theories of knowledge in philosophy. Evidence, on this model, is based on logical reasoning and on appearances (perceptions) that fit in with other appearances in a consistent manner, that can be tested, and that give reasons to support or refute a claim.

According to the requirements of the model, a proposition may rightly be said to be classifiable as scientific knowledge if it is supported by evidence of the kind used in a particular scientific discipline to the standard of proof appropriate for what counts as knowledge and what does not. A leading feature of the model is that it does not require that for a proposition to be included in knowledge, it must be true. However, it does require that in order to be classified as knowledge, a proposition must be based on input that comes from an external reality as evidential data. It is presumed that these evidential data come from external reality and that the knowledge resting on it is falsifiable. Hence on this new view, even though knowledge does not require an external criterion of truth, it is subject to support and refutation by evidence that can be tested and that comes from external reality. The model allows that knowledge claims that were formerly accepted can be defeated as new evidence comes to light. It is implied by the model that a standard of proving beyond doubt is not appropriate for claims to knowledge. The model is a doubly dynamic one whereby knowledge is not only continually being added to, but is subject to retraction. In the model, known propositions, even properly proved ones, can be retracted as new evidence is acquired.

The model of how knowledge is processed in an inquiry has been made as simple as possible in order to illustrate its central features, but further research would be useful to extend it. It has mainly been applied to representing scientific knowledge, but it could be extended to modelling how a detective conducts an inquiry using everyday conversational reasoning. The model departs most notably from the justified true belief analysis in its rejection of the external truth requirement. The model does not require that there is an agent who believes that the proposition in question is true. Indeed, the model, at least in the simplest form outlined here, does not require either the notion of the agent or the notion of belief, although it represents inquiry as a collaborative procedure undertaken by a group. The model requires weighing of evidence within a procedure in which a knowledge claim can be either accepted or rejected.

Acknowledgements

Research for this paper was made possible by a grant from the Social Science and Humanities Research Council of Canada. I would like to thank David Godden for inspiring me to take up this topic and for collaborating with me on an earlier paper (Redefining Knowledge in a Way Suitable for Argumentation Theory, presented at the OSSA 2007 Conference) that paved the way for this one. I invited him to collaborate with me on this paper as well, but said he was no longer in agreement with the new direction of my investigations, and encouraged me to go ahead and pursue the subject on my own.

Notes

1. In this paper, the terms "proposition" and "statement" are used interchangeably.
2. Verheij (1999, p. 115) (2000, p. 5) called this second form of inference *modus non excipiens*, arguing that it needs to be applied in cases where a general rule admits of exceptions.

References

Bex, F., Prakken, H., Reed, C., and Walton, D. (2003),'Towards a Formal Account of Reasoning about Evidence, Argument Schemes and Generalizations', *Artificial Intelligence & Law*, 11, 125–165.

Bondarenko, A., Toni F., and Kowalski, R.A. (1993), 'An Assumption-based Framework for Nonmonotonic Reasoning', in *Second International Workshop on Logic Programming and Non-monotonic Reasoning*, eds. L.M. Pereira and A. Nerode, pp. 171–189.

Brewka, G. (1989), 'Nonmonotonic Reasoning: From Theoretical Foundations towards Efficient Computation', PhD thesis, University of Hamburg, Hamburg.

Caminada, M. (2004), 'For the Sake of Argument: Explorations into Argument-Based Reasoning', PhD thesis, Free University of Amsterdam. http://icr.uni.lu/~martinc/publications/thesis.pdf

Caminada, M. (2008), 'On the Issue of Contraposition of Defeasible Rules', In *Computational Models of Argument: Proceedings of COMMA 2008*, eds. P. Besnard, S. Doutre and A. Hunter, Amsterdam: IOS Press, pp. 109–115.

Cooke, E. (2006), *Peirce's Pragmatic Theory of Inquiry: Fallibilism and Indeterminacy*. London: Continuum.

Gordon, T.F. and Walton, D. (2009), 'Proof Burdens and Standards', *Argumentation and Artificial Intelligence*, eds. I. Rahwan and G. Simari, Berlin: Springer, pp. 239–260.

Governatori, G. (2008), Defeasible logic, http://defeasible.org

Kakas, A., and Toni, F. (1999), 'Computing Argumentation in Logic Programming', *Journal of Logic and Computation*, 9(4), 515–562.

Lumer, C. (2005a), 'Introduction: A Map', *Informal Logic*, 25, 189–212.

Lumer, C. (2005b), 'The Epistemological Theory of Argument', *Informal Logic,* 25, 213–243.

McBurney, P., and Parsons, S. (2001), Representing Epistemic Uncertainty by Means of Dialectical Argumentation, *Annals of Mathemathics and Artificial Intelligence*, 32, 125–169.

Misak, C. (1987), 'Peirce, Levi and the Aims of Inquiry', *Philosophy of Science*, 54(2), 256–265.

Nute, D. (2001), Defeasible Logic: Theory, Implementation, and Applications, *Proceedings of INAP 2001, 14th International Conference on Applications of Prolog*, IF Computer Japan, Tokyo, pp. 87–114.

Nute, D. (1994), 'Defeasible Logic', in *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol. 3: Nonmonotonic Reasoning and Uncertain Reasoning, Oxford: Oxford University Press, pp. 353–395.

Peirce, C. S. (1931), *Collected Papers,* eds. C. Hartshorne and P. Weiss, Cambridge: Harvard University Press.

Peirce (1984), *Writings of Charles S. Peirce*, Vol. 2, ed. E.C. Moore. Bloomington:Indiana University Press.

Pollock, J.L. (1974), *Knowledge and Justification*. Princeton: Princeton University Press.

Pollock, J.L. (1995), *Cognitive Carpentry*, Cambridge, MA: The MIT Press.

Popper, K. (1963), *Conjectures and Refutations: The Growth of Scientific Knowledge*. Routledge: London.

Popper, K. (1972), *Objective Knowledge: An Evolutionary Approach*, Oxford: Oxford University Press.
Prakken, H. and Sartor, G. (1997). 'Argument-based Extended Logic Programming with Defeasible Priorities', *Journal of Applied Non-classical Logics*, 7, 25–75.
Reiter, R. (1980), 'A Logic for Default Reasoning', *Artificial Intelligence*, 13, 81–132.
Reiter, R. (1987), 'Nonmonotonic Reasoning', *Annual Review of Computer Science*, 2, 147–186.
Rescher, N. (2003), *Epistemology: An Introduction to the Theory of Knowledge*. Albany: Sate University of New York Press.
Rescher, N. (2005), *Epistemic Logic: A Survey of the Logic of Knowledge*, Pittsburgh: University of Pittsburgh Press.
Riveret, R., Governatori, G., and Rotolo, A. (2006), 'Argumentation Semantics for Temporal Defeasible Logic', in eds. L. Penserini, P. Peppas and A. Perini, Third European Starting AI Researcher Symposium (STAIRS 2006), Riva del Garda. *Frontiers in Artificial Intelligence and Applications*, 142. IOS Press, Amsterdam, pp. 267–268.
Simari, G., and Loui, R. (1992), 'A Mathematical Treatment of Defeasible Reasoning and its Implementation', *Artificial Intelligence*, 2(3), 125–157.
Steup, M. (2010), 'The Analysis of Knowledge', *Stanford Encylopedia of Philosophy* (accessed October 18, 2010), http://plato.stanford.edu/entries/knowledge-analysis/
Thorsrud, H. (2002), 'Cicero on his Academic Predecessors: The Fallibilism of Arcesilaus and Carneades', *Journal of the History of Philosophy*, 40, 1–18.
Toni, F. (2008), 'Assumption-based Argumentation for Epistemic and Practical Reasoning', *Computable Models of the Law, Languages, Dialogues, Games, Ontologies*, ed. G. Sartor et al., Berlin: Springer, pp. 185–202.
Verheij, B. (1999), 'Logic, Context and Valid Inference', *Legal Knowledge Based Systems. JURIX 1999*, eds. H.J. van den Herik et al., Nijmegen: Gerard Noodt Instituut, pp. 109–121.
Walton, D. (1996), *Argumentation Schemes for Presumptive Reasoning*. Mahwah, NJ: Lawrence Erlbaum Associates.
Walton, D. (2002), 'Are Some *Modus Ponens* Arguments Deductively Invalid?' *Informal Logic*, 22, 19–46.
Walton, D. (2005), 'Pragmatic and Idealized Models of Knowledge and Ignorance', *American Philosophy Quarterly*, 42, 59–69.
Walton, D. (2006), 'Argument from Appearance: A New Argumentation Scheme', *Logique et Analyse*, 195, 745–777.
Walton, D., and Godden, D.M. (2007), Redefining Knowledge in a Way Suitable for Argumentation Theory, Dissensus & the Search for Common Ground: Proceedings of OSSA June 2007, ed. H.V. Hansen, Windsor: University of Windsor, CD-ROM, 2007, 1–13.
Zagzebski, L. (1999), 'What is Knowledge'? in *The Blackwell Guide to Epistemology*, ed. J. Greco and E. Sosa, Oxford: Blackwell, pp. 92–116.