Logique & Analyse **150-151-152** (1995), 191-207

MY KNOWLEDGE, OUR KNOWLEDGE, AND APPEALS TO AUTHORITY[1]

Ruth MANOR

I. *The problem with Appeals to Authority*

Suppose Amnon writes an article defending a thesis p. In the first section of his paper, Amnon uses an argument of the general form of appeal to authority (ATA):

(1) A is an authority and A believes that p
Therefore, p

Tamar writes a response to Amnon's article claiming that argument is incorrect. She notes in addition (with reference to a popular critical thinking text book[2]) that this argument is an instance of the Fallacy of Appeal to Authority. Since Tamar's appeal to the textbook is also part of an ATA argument, we conclude that Tamar holds that some ATA arguments are correct (e.g., her own) while others are incorrect (e.g., Amnon's). The question we face is how can Tamar explain and justify her position; that is, when are ATA arguments correct and why.

The common answers to these question are found mainly in the discussion of fallacies within the current U.S. critical thinking tradition. There the fallacy is associated with the falsity of the premises: a fallacious ATA is claimed to be used when A is not really an authority in the domain relevant to p, or when A doesn't really believe that p. Clarification and justification of these views is suggested within a general theory of argumentation (Walton 1987), the interrogative model for knowledge seeking (Hintikka and Bachman, 1991, Hintikka 1987, Bachman 1995), and the pragmatic theory of fallacies (Walton, 1995). These attempt to explain incorrect ATA arguments as violation of rules of the various kinds — rules of argumentation, rules for interrogation or strategic rules in the interrogative models, or

[1] The paper was written during my stay at the Center for Philosophy of Science, University of Pittsburgh, in the summer of 1996.

[2] E.g., Kahane, 1995

the rules of cooperative discourse. They appear to agree that there are correct uses of ATA arguments and that incorrect ones do not involve violations of rules of inference and in that sense, that incorrect ATAs are not fallacious.

These discussions miss an essential feature of ATA arguments, which is important when we think of justifying why ATA arguments are correct, when indeed they are, and why others are incorrect. This question I attempt to address by considering the relationship between individual knowledge and communal knowledge, starting off with the assumption that individuals lack full rational autonomy.[3]

That some ATA arguments are viewed as correct (non-fallacious) reflects the view of knowledge in which the knower is not completely determined and could be singular or plural. The demarcation between what I know as an individual person (who has beliefs as a result of her experiences and inferences) and what I know as part of a community, is not completely determined. What I know is often based on inferences from assumptions which I personally may not really know but are part of what we know as a community.[4]

The question concerning the relationships between communal knowledge and individual knowledge is related to the questions of when ATA arguments are correct (non-fallacious) and why. ATA arguments are correct when the fact that one specific person believes something does provide grounds for another to accept this view. ATA arguments are fallacious when such passage of knowledge is unjustified.

According to this picture then, ATA arguments are the vehicle by which we transfer the knowledge of the (others in the) community to the individual and as such they cannot be viewed as merely weak and insignificant knowledge conferring inferences, as some of the discussion of fallacies seem to imply. Rather, they are an important vehicle in our (individual and

[3] This claim is made by Hardwig, 1988. "Scholars and scientists often can be knowers only because they have deferred and continue to defer to other experts in their pursuit of knowledge. Thus modern knowledge is often possible only through a complex network of appeals to the authority of experts." p. 126.

[4] I am using "community" rather broadly, as referring to any community one belongs. Thus, I see myself as part of the scientific community or a few scientific communities, my family, my universities, the group of people in my neighborhood, etc. "Communal knowledge, is what is known by the community". This, I take may be different from 'common knowledge', which is shared by all or most members of some community. Communal knowledge, like scientific knowledge or our cultural heritage, may not be shared by all or most, but is a rather intricate structure which relates parts of the knowledge bases of individuals and a network of knowledge-seeking activities by some members of the community.

collective) knowledge seeking. This suggests that some ATA arguments are non-fallacious.

Yet viewed in this way, ATA arguments may involve a basic problem, or a fallacy: When A is an authority in the field of p, and when A does believe p, it may still be incorrect to conclude that p. Consider the case when we teach a class, say a class on fallacies. We first introduce the topic, the concepts, the rules, and examples. The rational students listen and study what we say, on the mere basis that we are the teachers, and they study the texts because we, the teachers instruct them to do so. In the course of this stage we may give an example of an argument and claim it exemplifies the Appeal to Authority Fallacy, and the wise student will be efficient in accepting this position on the mere grounds that we the authorities, have asserted it. This is a case of a correct ATA.

But consider the later stage, after reviewing the material, the student is supposed to systematize the information we gave her. She's supposed to internalize and understand the topic. She goes back to that example we gave, and asks whether we really think it to exemplify the Appeal to Authority Fallacy, and why. At this instance, our response "Believe me, I know, I'm an expert in these things", is unjustified, even though we are indeed experts and our claim about the example is in fact true. It would be fallacious for the student to accept the conclusion on the mere grounds that we, the teachers, have asserted it. This case deserves the title of *the Fallacy of Appeal to Authority*, for it presents an essential limitation on an inference rule which is accepted *prima facie*. The question we face in this case is why these are fallacious uses of ATA or why ATA arguments are limited.

So far, I have used 'correct argument', 'knowledge', 'rationality' and 'fallacy', as given, unanalyzed terms in natural language. The long term aim of a theory of fallacies is to provide a systematic account of these terms. My aim here is to point out some of the interconnections between these concepts. In general, I use 'correct argument' to indicate a pragmatic and context dependent concept: the argument is correct when it is rational for us to accept (tentatively, at least) the conclusion on the basis of its premises. I use 'knowledge' in its natural-language form which includes scientific knowledge, and may or may not be objectively true in a realistic sense.

I do not clarify what rationality is, but I use it mainly in a limited form. Mostly I'm concerned with when it is rational for a person to accept the conclusion of an argument on the basis of its premises. I assume that fallacies typically noted in classifications are a species of incorrect arguments. Moreover, I believe that they name flaws in arguments that are, typically, correct arguments, and may even be important inference patterns. For some, there is a theory systematically justifying their correctness. For others, like the ATA, the fallacy literature provides a good starting point for

tracking and forming a theory of the correct forms. In the following, I will not argue the general thesis concerning fallacies. Rather, I will attempt to show that ATA is connected to a basic and important inference pattern which is *prima facie* correct and that the Fallacy of ATA a basic limitation on the inference pattern, i.e., when an argument following the pattern still can be used incorrectly.

The attempt to formulate a theory of fallacies in terms of rationality is not very helpful, unless we have a clear notion of rationality. However, given that we have neither a good theory of fallacies nor a good theory of rationality, but we have some intuitions concerning each, investigating their connection can help in clarifying both concepts. Thus, I attempt to use the ATA Fallacy to illuminate the interrelations between our (singular or plural) knowledge and our (singular or plural) rationality.

II. *Common analyses of* ATAs

Walton (1989) is concerned with a theory of argumentation, and thus the correctness of the ATA depends on when use of such arguments is within the rules of the argumentation game. He provides the following scheme for using for ATA in argumentation:

(2) *A* is an expert in domain *D*
A asserts that *p* is known to be true
p is within *D*
Therefore, *p* may (plausibly) be taken to be true (p. 193)

The use of this scheme can be judged to be weak, erroneous, or insufficiently documented if one or more of these premises is inadequately supported. Indeed, when the premises are not all true, then the argument will be incorrect (though the argument scheme may still be correct, of course). Such mistakes are often described as instances of The Fallacy of Appeal to Authority.[5] Bachman (1995) follows the critical thinking tradition of translating '*Ad verecundiam*' as the Fallacy of Appeal to Authority. But he claims that these arguments are strictly speaking, not to be viewed as fallacious but as arguments which involve errors in strategies for effective reasoning. Walton argues that these mistakes are not an indication of a fallacy, and thus refuses to address them as "the fallacy" of appeal to authority. "To claim that an argument contains a fallacy is a strong form of criticism implying that the argument ... is based on an underlying flaw or misconcep-

[5] e.g., Cederblom & Paulsen, 1982, pp.104-108

tion of reasoning, and can therefore be refuted." (Walton, 1989 p.16.) Yet, he does identify a fallacy of *Ad Verecundiam,* which is committed by an ATA argument if the argument was advanced as part of a strategy of being overly aggressive in trying to prevail and prevent the respondent from advancing critical questions (p.196).[6]

I tend to agree with Walton that these mistakes associated with the falsehood of the premises characterizing the authority of *A*, are inessential and that since the term "fallacy" should be reserved for systematic flaws in reasoning, these should not be viewed as fallacies.

Let us consider Walton's scheme (2) for correct ATA. We note that the second assumption,

(3) *A* asserts that *p* is known to be true.

is problematical. It appears to presuppose an ATA argument which is erroneous by Walton's analysis, as follows. How can *A* know that *p is known to be true*? Presumably, *A* will come to this conclusion after checking the knowledge of all, most, the most expert of all, or some other subset of the experts in the relevant field. This is an erroneous ATA according to Walton (1989, p. 179-80), because of its vagueness: the names of the experts who are claimed to *know* the truth of *p* are not cited.

One may respond to this that the second assumption in (2) does not require that *A* be justified in his assertion but only that he asserts that *p* is known to be true. But Walton needs this assumption to avoid relying on one expert when experts disagree, namely, the above assumption does involve an argument of the form:

(4) *A* asserts that *p* is known to be true
Therefore, *p* is known to be true.[7]

And this argument is an ATA can satisfy Walton's scheme only if the following assumption is added

(5) *A* asserts that *p* is known to be true is known to be true.

Which requires not only that our expert *A* knows that (all/most/most experts/...) experts in the field know that p, but that (all/most/most expert/...)

[6] A similar argument appears also in Walton, 1995.

[7] or: "Therefore, that *p* is known to be true may (plausibly) be taken as true".

experts in the field know that (all/most/most experts/...) experts in the field know that *p*.

As we shall see below, ATA arguments may be justified without assuming that experts know that there is a consensus among them concerning *p*.

There is another problem with Walton's scheme. He refers the readers to Rescher's 1976 Plausible Reasoning for an explanation of his notion of plausible reasoning. Yet, Rescher's notion has to do with reasoning from inconsistent premises. If *A*'s assertion that *p* is indeed known to be true, as is implied by the premises in Walton's scheme (2), then presumably *p* is noncontroversial and could be taken as true without the need to qualify the conclusion as being merely plausible.

Govier 1988 adds another premise concerning the reliability of the authority *A*

(6) *A* does not have a vested interest in *p*'s being true, nor has he been dishonest about matters related to *p* in the past.

She is not satisfied with a weak condition that we have no (reliable) information that implies *A*'s self-interest or dishonesty. Rather, she requires that we know, i.e., at least have some good reasons for concluding that *A* passes this requirement. This seems too strong and does not conform our actual practice when we rely on authority. Moreover, the unrestricted addition of this condition seems to limit the use of appeal to experts significantly. For then we cannot use information (about *A*'s reliability) from other sources, about whose reliability we have no information. This places a severe restriction on our sources, maybe even an impossible one.

Bachman (1995), does not object to Govier's list but in arguing that erroneous ATA arguments are strategic mistakes, he cites an additional case, which in terms of the above could be formulated as an additional requirement and premise

(7) *A* is the best available expert in field *D*.

Now, let us leave the questions concerning the detailed restrictions on ATA arguments open and consider the basic question of how we can justify any ATA in general.

Presumably, what needs justification in the Walton/Govier/Bachman schemes is the assumption that an authority in the field has plausible beliefs concerning the field in question, while in general, non-authorities may not have such plausible beliefs. Namely, I take it that from the Walton/Govier/Bachman point of view, arguments of the form:

(8) *S* asserts that *p*
Therefore, *p* may (plausibly) be taken to be true

(where *S* is any person who may not be an authority at all), is incorrect while (2) is.

Most discussions of ATA concentrate on answering the question when ATA arguments are correct, when they are incorrect, and why they are incorrect when they are. Very little addresses directly the question of why an ATA argument is correct when it is. Walton (1989) claims that ATA arguments are neither deductively nor inductively justified. Moreover he claims that ATA arguments are inherently subjective and judgmental as experts base their judgments on "rules of thumb" and accepted methods they have found useful in their practice. Hence, science has traditionally mistrusted ATA arguments, and these arguments are to be viewed as weak and merely plausible.

Walton does not systematically justify why ATA arguments are plausible at all whenever they are. Why could we, or why should we, rely on expert opinion, when indeed they satisfy the schemes above? Expert systems, according to Walton, are computer programs that duplicate the skills of an expert in a well defined area of expertise are not only widely used, but their practical usefulness has been well established. These rather swooping claims and their relation to his analysis of correct appeal to expert opinion is left unexplained. Moreover, even if we accept the implied claim that expert systems are sources of information which are themselves confirmed empirically, the question of explaining why this is so, is not addressed at all. This may be explained in context for Walton is concerned with argumentation, and the question of justification may not arise in that context at all. He states his theory of argumentation and its rules and he shows that he can derive some common sense (or critical thinking) claims about the fallacies. Thus, I believe he uses the fallacies, that is, what he claims to be facts involving fallacies, to test his theory of argumentation. The present demand for justification questions the factual status of his claims as well as their truth.

Hintikka and Bachman offer the interrogative model as a model for knowledge seeking. The interrogative step is a fundamental ingredient for rational inquiry, in which we seek information from sources. Thus, appeal to any source is, so to speak, accepted as a fundamental step. Hintikka (1987) refers to two different models. In the basic interrogative model, the assumption is that the source's answer is always true. The second model is obtained when the basic model is revised to cases where the source's answers are assumed to be at most probable and fallible. The inquiry itself can concern the evaluation of the reliability of an answer given by a source, possibly by evaluating the reliability of the source itself. According to

Bachman, errors of ATA are strategic errors "concerning the failure to employ strategies of answer and oracle evaluation that would have made our reasoning more effective." (Bachman 1995, p.279).

Thus, the interrogative model implies that successful strategy in obtaining true or probable information is a criterion for distinguishing correct ATA's from the incorrect ones.[8] Hence, if ATA arguments following the Walton/Govier/Bachman schemes have in fact highly probable conclusions, then these schemes could be schemes for correct ATAs. However, none of them suggests that their schemes are such that they have highly probable conclusions nor are they concerned with justifying their schemes using probabilistic considerations.

III. *Justifying* ATA *arguments*

1. I am looking for a justification of ATA arguments. As not all ATA arguments are correct, the justification sought needs to justify only the correct ones, and hence it may provide a criterion for distinguishing correct from incorrect ATA arguments.

I consider ATA arguments as inference patterns that we employ in everyday reasoning. Some such arguments we claim are used correctly, while others are not. The sense of correctness that we employ has to do with a certain form of the argument, e.g., form (1), but it must go beyond the mere form as it depends essentially on the notions of authority or expert. The criteria for satisfying these characterization (of being an epistemic authority or an expert) are social and depend on the community, and hence they are context dependent. I therefore prefer talking of the correctness of the use of an argument in specific contexts, and to do so in terms of rationality, as follows. In certain situations it may be rational for me to accept p on the mere evidence that you told me that p. In this way, a justification of my claim concerning the correctness of the argument, is translated into showing that my accepting p is rational, that is, that accepting p may be a strategy for achieving my goals (in the situation I am in) which is not worse than the opposite strategy, of not accepting p. This, of course implies the Hintikka and Bachman position that incorrect ATA's are strategy mistakes, but it says actually more: that correct ones are strategically efficient. It is this claim that needs justification. In Walton's terms, the question is not why would ATA arguments be merely plausible, but what makes them plausible rather than implausible ones?

[8] Bachman (1995) suggests that a successful strategy in obtaining information in which we have more confidence could also serve as such a criterion.

Walton (1989, p.177) claims that science mistrusts ATA because ATA arguments are subjective and fit neither the deductive pattern nor the inductive pattern, the cannons of scientifically accepted objectively confirmed arguments. Walton does not go into this question further, but we may consider the claim that an expert, or an authority, is one who has a record of correct predictions in the relevant fields. This is somewhat supported by the practice of testing candidates for diplomas extensively, and the fact that we consider the record of successes and failures in performance of a person as relevant to her level of expertese. This suggests that ATA arguments are justified (when they are) inductively because the authority is one who is empirically proven to be a highly reliable source.

Interestingly enough, I could not find any requirement of the correct ATA scheme to have a premise that *A* should be reliable in the sense that he has a record of correct prediction that would make his future predictions probable. As mentioned, none of the above authors addresses this issue. Moreover, I think this is not how we normally use authorities or experts. While the record of success and failure in some tests is relevant to the determination of expertise of a person, it is not the only factor.

We thus hope or expect that an expert's track record in the field makes her predictions probable, but the fact that the authority has no track record of verified predictions to speak of, is no reason not to consider her as an authority. Again, even if the authority has a track record, normally it is available to us indirectly only, via the reports of others. If we require that these sources be experts on the given question, again we have normally no way in empirically confirming that but only indirectly, and so on. Namely, theoretically we can require that an expert in a field be one who is empirically proven to be highly reliable, and thus ATAs can be justified empirically. Yet, this means that any report of others is unjustified unless it is a report of an expert (satisfying the reliability condition above) in the field. So in checking whether *A* is indeed an expert we must have first hand experience of his past record or we must use evidence of other experts. This limits absurdly our notion of expert and cannot be used in an attempt to justify ATA.

In practice, what we check is whether our expert satisfies some set of requirements which seems relevant to the issue, as checked and reported by others. For instance, we choose a doctor using possibly established criteria (her diplomas), advice of other doctors, and possibly by the evidence of a friend who was her patient. These reporters may be experts or non-experts in the field, and we may have no knowledge of the doctor's and the advisers' reliability. We hope and expect that the doctor is reliable, based on the established criteria and the say-so of the reports as well as on the fact that

we have no direct or indirect information concerning the unreliability of the doctor, the diplomas she has, or the reporters.[9]

2. *Appeals to (any) Source* (ATS)

It seems that the above argument is a general argument against any justification of ATA arguments which relies on a special notion of epistemic authority. If (i) I am unjustified in accepting reports by others unless they are epistemic authorities (in the relevant field), and (ii) to be an epistemic authority (in the relevant field) one must satisfy some conditions *C*, then the set of justified ATA arguments will be absurdly small: I can rely on *A*'s report only if I have direct evidence or the evidence of epistemic authority *B* that *A* satisfies *C*. There are very few people I have direct evidence of what they know and whether they satisfy any epistemic conditions. There are very few fields in which I can judge of anyone whether they pass some conditions *C* or not. Given the above, I think I will be unjustified in accepting any report of anyone other than my students and colleagues in my own area.

This suggests that in order to justify ATAs we need to reject the presumption that subjective reports of non-experts and non-authorities are incorrect. Indeed, it is because anybody's report that *p* is *prima facie* grounds for believing *p*, that it may be rational for us to believe a claim of a non-expert in the face of conflicting expert reports. For instance, it may be rational for me to accept the Childs' claim that the emperor has no clothes on, in spite of the expert tailors' contrary claim.

Thus, I suggest that *prima facie* arguments of the following form are correct in the sense and it is rational for us to accept the conclusion on the basis of the premise:

(9) *S* asserts *p*
Therefore, *p*

I take (9), appeal to a source (ATS), to be a *prima facie* rule of inference: given that *S* asserts *p*, it is *prima facie* rational for me to accept *p*. It would be rational for me to not accept *p* only if I have additional reasons for rejecting *S*'s report.

I need to place this rule in a broader picture. I claim that what we know as individuals and what we know as a community (or as parts of many communities) is interconnected and interdependent. When we consider

[9] This expectation that the doctor is reliable can be justified not in view of her own track record but in view of the track record of the general method of choosing an expert in view of evidence like the above. I doubt that this is empirically testable.

what an individual knows or believes in, then though we may be able to identify a piece of information that she has while others in her community do not, (e.g., that she has a headache and how it feels), there is no way to identify her knowledge system as a system without including bits of knowledge that are shared by members of the community, and of which she has no first hand knowledge. When we learn a language we not only obtain factual information, but also methodological information on how we systematize the information. The same is true concerning any other area which we learn, be it science, religion, literature, or anything else: the reports concern specific items of information together with methodologies.

The point I'm trying to make is not just that we share some of our individual information, but that we cannot separate and find the identity conditions for that part of someone's knowledge base that is uniquely their own. One may of course know something that no one else can know in the same way, in the sense that, for instance, no one can know what I feel at a given point of time in the same way that I know it. But, whatever I know about this feeling is so interlaced with other knowledge obtained indirectly from the community, that even this knowledge is not purely mine.

Even in the way we talk about knowledge (and rationality), there seems to be a duality in using "we know" meaning "each of us knows" and when we use "we know" to refer to some meriological collective. Thus, it is true that we know how to get to the moon and that I know that smoking causes cancer, even though, strictly speaking, this is not my knowledge. It is not a piece of information derived using only pieces of information I have acquired without the aid of my community. We often talk about our scientific knowledge, which includes a collection of information systems, built on the basis of what *some* scientists believe. But this picture is distorted. It appears to assume not only that what I know as an individual is closely connected to communal knowledge, but that individual knowledge bases serve as basic and simple components of the meriological communal knower, e.g., science or the scientific community. I claim that community knowledge is in the individual knowledge bases of its members, yet each member develops his own knowledge base on the basis of a view of a communal knowledge base as well as his own interaction with the world and other individuals.

The argument pattern (9) is the vehicle by which we pass information to each other, and hence it is the vehicle by which individual knowledge is passed to the community. This pattern is *prima facie* correct in the sense that it is *prima facie* rational for an individual to accept p on the basis of S's asserting it, rather than to disbelieve S's report. Why?!

— Imagine that the life form discovered recently from Mars has developed to a creature that is rational but it is extremely suspicious. *prima facie* Martians do not trust what other Martians tell them. If they have other grounds for believing p, then of course the fact that a fellow Martian tells

them that p is no reason to question p, but without this added grounds they do not adopt a belief on the mere basis of a Martian's say so. Thus, a Martian who is told that Pythagoras' theorem is provable, will not accept this report as true. Being rational he will, of course, accept the report as true, after reading the proof.

Yet, I cannot see how the Martian's systematic suspiciousness is consistent with his rationality, for he could not learn anything, not even a language. *prima facie* distrust will place Martians at a great disadvantage, compared to us trusting Earthlings: while we gulp information from all sources, we err often and need to correct our mistakes and constantly balance contradictory evidence, the mistrusting Martians don't even have the flow of incoming information. They probably make no less errors than we do and they similarly have to face contradiction, for the world itself, not merely others' reports, may be observed in paradoxical ways. Their disadvantage lies in that they have to create and solve their problems alone, one mind at a time.

3. *Restrictions on* ATS *arguments*

The above outlines a *prima facie* justification of pattern (9) of ATS arguments, which implies but is not implied by ATA arguments. If we accept that pattern (9) is indeed a pattern of correct reasoning, then all ATA arguments are correct, which is too much.

So we need to look at possible restrictions and refinements of ATS arguments. There are two general considerations. One is obvious: different sources may, and often do provide conflicting information, and this presents a problem in constructing consistent belief systems. In these cases, we place some restrictions on the correct use of ATS arguments, in an attempt to maximize the information we receive from others while minimizing the conflict. Thus, in the presence of conflict, sources are restricted to epistemic authorities. To complete the justification in this case, we need to show how accepting ATA and rejecting ATS in cases of conflicting reports will indeed maximize the information and minimizes the conflict. I will discuss this case in greater detail below.

The second consideration is not as obvious but it is essential and concerns the relationship between the individual and the communal knowledge systems. Communal knowledge enables individuals to share particular items of information and more importantly, it enables the individuals to cooperatively develop tools for systematizing knowledge. This sharing is possible first by considering informational items from different sources and comparing the information and the sources.

Second, the systematization requires that the individuals consider, compare, and organize informational items according to their contents and independently of their sources. Once the individual accepts some information

tentatively, he can further evaluate its pieces on the basis of their explanatory power, i.e., their contribution to the systematization of his knowledge base as a whole, raise new hypotheses and reject old pieces of information merely on the basis of systematization consideration (and not on the basis of reliability of its source).

It is this situation in which we are involved in this process of information organization, that questions of the sources of information, where we got the information or how reliable it is, are irrelevant. It is then that it is not rational for a person to accept or reject p merely on the basis of someone's say so, be it an expert or not. It is in this context that the use of ATA arguments is fallacious.

A. *Justified* ATA: *the case of conflicting information*

Without getting into the difficult question of how our knowledge bases are organized and how we use them, we can assume that we seek maximal consistent knowledge, and that our knowledge bases are hence organized systematically. From this assumption it follows that it rational for us (both as individual knowers and as a community of knowers) to act so that our knowledge base has maximal consistent information systematically organized.

Suppose we have two sources, $S1$ and $S2$, who report p and not-p respectively. Accepting both reports (using ATS reasoning) will violate the consistency requirement. Hence, it appears that we need to restrict ATS to be inapplicable in these cases. While rejecting both reports may be required in some cases, its use as a universal strategy is in conflict with the goal of maximizing our information. Hence, we need to develop a method to rank and evaluate ATS arguments, so we can accept one of the two conflicting reports (in many or most cases).

I believe there are two types of considerations. The first relates to existing information concerning the probability of the reports or the reliability of the sources. Though often we have no probabilistic information concerning the reliability of a source, when we have information concerning, say the fact that $S1$ lies often, it is a reason to prefer the information of $S2$. The more common situation is that we have insufficient information about the sources, and then we may use another method to rank the sources. Though the method may not in the final analysis yield the preference of the most reliable source among $S1$ and $S2$, using a specific method systematically for a while will provide us with the data concerning the efficacy of the method, which may in the long run, contribute to our ability to rank sources. For instance, in the case where I have no information concerning p and I have no information concerning the reliability of $S1$ or $S2$, I may be justified in systematically preferring $S1$'s report to $S2$'s for a while, in order to test the reliability of $S1$'s reports. In another situation, in seeking new solutions to a

particularly difficult problem, I may prefer to follow *S*2's report simply because it is part of a new approach.

In general, in our communal culture, a person is considered an expert in a field *D* according to some conventional criteria. These criteria, I believe are constructed combining the two kinds of considerations.

I wish to stress that reliability of one's reports is not the only and possibly not the main reason for a person to be considered an expert in an area *D*, or for us to prefer her report (when it conflicts with someone else's). I believe that concentrating on the reliability in the discussion of ATA and expert witness testimony, distorted the view of how and why we use expert reports. Walton's premise (4), that *A* asserts that *p* is known to be true, in his (1989) scheme (2) for correct appeal to experts, reflects this distortion. It suggests that we rationally seek expert opinion on topics which are not controversial, where *p* is known, by all or some experts. But knowing what all experts agree on is a basic requirement from students studying the material. We often prefer the report of experts because they can give us a new, individual, interesting view of "the big picture", of how the information bits hang together. This has more to do with what we think of the expert's ability to systematize the information in the field than with what we think of his reliability.

This completes the characterization and the justification of ATA arguments. The justification went along the following lines: first, I claim that ATS arguments are *prima facie* correct inference rules, which need further qualifications in case we have conflicting reports. ATA arguments are ATS arguments from sources who are by community convention preferred to other sources, in cases of such conflicts. Assuming we have no way to evaluate the contents of the reports, then in case of a conflict between experts (who are conventionally ranked as equally expert in the field), we may of course utilize different methods for preferring one source to the other but these will not be determined by convention and hence are not classified as an appeal to authority.

B. *Unjustified* ATAs: t*he Fallacy of Appeal to Authority*

Yet even when there are no conflicting reports concerning *p*, or when there are conflicting reports but *A* satisfies all the requirement to make a correct ATA as described above, there are situations in which it would be irrational for me to accept *p* on the basis of *A*'s report that *p*. This has to do with the individual's actions to improve on his own belief system and hence also the communal one. By improvement I'm referring to the systematic organization of the information items and their interrelation. This process may result in changing one's beliefs. It is in these cases that using ATA arguments may be fallacious.

Consider again the case of teaching a class introduced in the beginning. In the initial stages, the students are expected to absorb the information on the basis of ATS or ATA. The teacher is a source, maybe an expert, and her reports are accepted.

After the students have received sufficient information they are able to internalize and understand the material. This means that they reorganize the information, adding their own hypotheses (after all, their own hunch can be added via an ATS argument). Certain items of information accepted initially and tentatively (via ATS or ATA arguments) may then be rejected on the basis that their negation contributes better to the systematization of the whole information in the field. It is at this stage that accepting p on the basis of ATS or ATA arguments is unjustified, and constitutes a fallacy.

The use of ATA in this situation does not constitute a minor blunder, as in the case when the authority is not a "real" expert. The violation involved here is one in which the argument is used in a context which makes it inappropriate. It is the context of the individual's epistemic goals and activities, and not the question how expert is the authority, that makes the ATA fallacious in this case. Indeed, we can imagine a case in class, where in the first stage in which the teacher passes the information in general, the teacher asserts that p and the student accepts it, rationally. Yet, in the second stage, the student may have some reasons for rejecting p, then his accepting p on the basis of the same ATA argument is unjustified. Thus, the Fallacy of Appeal to Authority is the use of ATA by an individual when he's concerned with the systematization of his knowledge base.

The reason this use is fallacious is not merely that at this stage the question of the source of the information is irrelevant to the individual's epistemic goals, nor is it merely that relying on ATA's at this stage confuses the probability that p is true with the explanatory power of p within the knowledge base. The main point is that it is the individual who systematizes information, and this cannot be shared via a passage of information. The final acceptance or rejection of p depends on a process of systematization and it is a private and individual process. Even when the expert teaches a topic systematically and in great detail, the individual student understands it in reinterpreting it, drawing his own connections, associations, and inferences. We share information, we can even share information about how we systematize our knowledge, but the systematization process itself is individual, it is done one mind at a time, and cannot be appealed to anyone.

Whether or not it is rational for me to accept p depends on the one hand on the evidence I have for the truth that p, and on the other hand, on how accepting p contributes to the systematization of my knowledge base. The same in general, holds for the communal knowledge. Yet, while individuals may gain specific information in virtue of their own experiences, there is no communal entity that has experiences, makes decisions, and interacts with

the world. There is no communal entity that knows whatever we (in the plural) know, and that like the individuals, accepts information from various sources, evaluates them and chooses to accept propositions and act accordingly. Communal knowledge is built only of what the individual members of the community know, but it has not the subjective means that the individual has, for ranking and choosing what to accept. My point is not that we cannot assume that a communal entity exists, like the scientific community or whatever may be viewed as the source of correct scientific procedure. Rather, the point is that this entity, unlike an individual, has no subjective experiences and hence no grasping nor interest in the information, beyond the collection of interests and goals of its individual members.

Hence, communal knowledge consists of information which can be shared, that is equally knowable by different individuals, and which is useful for the epistemic goals of different individuals. This information has to do with the systematization of the knowledge, rather than the subjective details. The subjective details are relevant only to the extent that they relate, confirm or disconfirm, a systematization of a body of knowledge. Individuals share in the communal knowledge to the extent that their own experiences and knowledge bases cohere with the communal.

Since it is in the individual's interest to increase his own and hence also the communal knowledge, it is the individual's task to increase systematization of the communal knowledge. It thus may be his task within the community to question propositions even when those are accepted by all the individuals in the community, including himself, to check the effect of this action on the systematization of the rest of his knowledge base. In this case, not only his use of ATA is unjustified, but he may have the task of using (occasionally and tentatively) an opposite rule of inference, concluding p's negation from the fact that everyone says it.

Thus, appeal to a source and appeal to authority may be justified in the realm of the individual knowledge base, as these are the vehicles by which we share maximal knowledge. These are not justified in the realm of communal knowledge, where the issues concern the systematization of the information.[10]

Tel-Aviv University and San Jose State University

[10] It is an honor to write in memory of Leo Apostel. I believe he would have enjoyed the topic, and I regret that I did not have the benefit of his criticism.

BIBLIOGRAPHY

James Bachman, 1995 "Appeal to Authority", in *Fallacies, Classical and Contemporary Readings*, Hans V. Hansen, Robert C. Pinto (eds), Pennsylvania State University Press 274-287.

Jerry Cederblom, David, W. Paulsen, 1982, *Critical Reasoning*, Wadsworth.

Frans H. van Eemeren, Rob Grootendorst, Francisca Snoeck Henkemans, J. Anthony Blair, Ralph H. Johnson, Erik C. W. Krabbe, Christian Plantin, Douglas N. Walton, Charles A. Willard, John Woods, David Zarefsky, 1996, *Fundamentals of Argumentation Theory, A Handbook of Historical Backgrounds and Contemporary Developments*, Lawrence Erlbaum Associates, Publishers.

Trudy Govier, *A practical Study of Argument,* 2nd ed. (Wadsworth, 1988).

C. L. Hamblin, 1970, *Fallacies*, Methuen & Co.

John Hardwig, 1988, "Relying on Experts", in *Selected Issues in Logic and Communication*, Trudy Govier (ed), Wadsworth, pp.125-137.

Jaakko Hintikka, 1987, "The Fallacy of Fallacies", *Argumentation* 1, 211-238

Jaakko Hintikka, James Bachman 1991 *What IF? Towards Excellence in Reasoning*, Mayfield, 1991

Howard Kahane, 1995, *Logic and Contemporary Rhetoric: The Use of Reason in Everyday Life*, 7th ed., Wadsworth.

Nicholas Rescher, 1976, *Plausible Reasoning: an Introduction to the Theory and Practice of Plausibilistic Inference*, Van Gorcum.

Douglas N. Walton, 1989, *Informal Logic A handbook for Critical Argumentation*, Cambridge University Press.

Douglas N. Walton, 1995, *The Pragmatic Theory of Fallacies,* University of Alabama Press.