How Convinced Should We Be by Negative Evidence?

Ulrike Hahn (hahnu@cardiff.ac.uk)
School of Psychology, Cardiff University, Tower Building, Park Place
Cardiff CF10 3AT, UK

Mike Oaksford (oaksford@cardiff.ac.uk)
School of Psychology, Cardiff University, Tower Building, Park Place
Cardiff CF10 3AT, UK

Hatice Bayindir (bayindirh@cardiff.ac.uk)
School of Psychology, Cardiff University, Tower Building, Park Place
Cardiff CF10 3AT, UK

Abstract

Since John Locke, the so-called argument from ignorance has been considered to be a fallacy, and is widely represented in informal logic textbooks as an example of incorrect reasoning. This might seem surprising to researchers in many scientific disciplines who routinely draw inferences from negative evidence. Oaksford and Hahn (2004) argued that this discrepancy can be explained within a Bayesian framework. We present here experimental evidence for this view.

Introduction

Fallacies, or arguments that seem correct but aren't, have been a longstanding focus of debate. Catalogues of reasoning and argumentation fallacies originate with Aristotle and populate books on logic and informal reasoning to this day. One such classic fallacy, which dates back to John Locke, is the so-called argument from ignorance, or argumentum ad ignorantiam:

(*1*) *Ghosts exist, because nobody has proven that they don't*

This argument does indeed seem weak, and one would want to hesitate in positing the existence of all manner of things whose non-existence simply had not been proven, whether these be UFO's or flying pigs with purple stripes.

However, is it really the general structure of this argument that makes it weak, and if so what aspect of it is responsible? Other arguments from negative evidence are routine in scientific and everyday discourse and seem acceptable:

(*2*) *This drug is safe, because no-one has found any side effects*

Should all arguments from negative evidence be avoided, or can a systematic difference between the two examples be recognized and explained?

The classic tool brought to the analysis of fallacies such as the argument from ignorance is formal logic and it is widely acknowledged to have failed in providing a satisfactory account. Testament to this is the fact that fallacies figure in logic textbooks under the header of 'informal reasoning fallacies' (see e.g., Hamblin, 1970) – an acknowledgement of the absence of a sufficient formal logical treatment. In particular, logical accounts have proved unable to capture the seeming exceptions to fallacies that arise with simple changes in content that leave the structure of the argument unaffected. This suggests that either it is not formal aspects of fallacies that make them fallacious, or else that the relevant formal aspects are not being tapped into by classical logics.

The so-called pragma dialectical approach (see e.g., van Eemeren & Grootendorst, 2003; Walton, 1995) is a more recent approach to the fallacies which eschews the idea that fallacies can be explained purely through reference to their inherent structure. Rather, fallacies need to be viewed within the wider argumentative context in which they are embedded. Arguments are fallacies because they fall short of standards of rational discourse.

This approach has its roots in pragmatics (e.g. Grice, Searle) and seeks to distinguish different types of argumentative discourse (e.g. 'information seeking') for which normative rules are then established. An example of such a rule is: "the discussant who has called into question the standpoint of the other in the confrontation stage is always entitled to challenge the discussant to defend his standpoint" (Rule 2, van Eemeren & Grootendorst, 2003). The argument from ignorance on this account then, is fallacious wherever, and because, it violates the discourse rules of the current context.

What such an account cannot explain, however, is variations in the strength of different arguments from ignorance within the same type of discourse context. Oaksford & Hahn (2004) provide evidence of such variation and put forth an alternative, Bayesian account: individual arguments such as (1) and (2) are composed of a conclusion and evidence for that conclusion. Both conclusion and evidence have associated probabilities which are viewed as expressions of subjective degrees of belief. Bayes' theorem then provides an update rule for the degree of belief associated with the conclusion in light of the evidence. Argument strength, then, on this account is a function of the

degree of prior conviction, the probability of evidence, and the relationship between the claim and the evidence, in particular how much more likely the evidence would be if the claim were true.

A Bayesian account captures, among other things, the difference between positive and negative evidence and allows one to capture the intuition that the positive argument (3a) is stronger than the negative argument (3b):

(*3a*) *Drug A is toxic because a toxic effect was observed* (*positive argument*).

(*3b*) *Drug A is not toxic because no toxic effects were observed* (*negative argument, i.e., the argument from ignorance*).

However, (3b) too can be acceptable where a legitimate test has been performed, i.e.,

> *If drug A were toxic, it would produce toxic effects in legitimate test.*
> *Drug A has not produced toxic effects in such tests*
> *Therefore, A is not toxic*

Demonstrating the relevance of Bayesian inference for negative vs. positive arguments involves defining the conditions for a legitimate test. Let e stand for an experiment where a toxic effect is observed and $\neg e$ stand for an experiment where a toxic effect is not observed; likewise let T stand for the hypothesis that the drug produces a toxic effect and $\neg T$ stand for the alternative hypothesis that the drug does not produce toxic effects. The strength of the argument from ignorance is given by the conditional probability that the hypothesis, T, is false given that a negative test result, $\neg e$, is found, $P(\neg T|\neg e)$. This probability is referred to as negative test validity. The strength of the argument we wish to compare with the argument from ignorance is given by positive test validity, i.e., the probability that the hypothesis, T, is true given that a positive test result, e, is found, $P(T|e)$. These probabilities can be calculated from the sensitivity ($P(e|T)$) and the selectivity ($P(\neg e|\neg T)$) of the test and the prior belief that T is true ($P(T)$) using Bayes' theorem:

$$P(T \mid e) = \frac{P(e \mid T)P(T)}{P(e \mid T)P(T) + P(e \mid \overline{T})(1 - P(T))} \quad (4)$$

$$P(\overline{T} \mid \overline{e}) = \frac{P(\overline{e} \mid \overline{T})(1 - P(T))}{P(\overline{e} \mid \overline{T})(1 - P(T)) + P(\overline{e} \mid T)P(T)} \quad (5)$$

Sensitivity corresponds to the "hit rate" of the test and 1 minus the selectivity corresponds to the "false positive rate." There is a trade-off between sensitivity and selectivity which is captured in the receiver operating characteristic curve (Green & Swets, 1966) that plots sensitivity against the false positive rate (1 – selectivity). Where the criterion is set along this curve will determine the sensitivity and selectivity of the test.

Let n denote sensitivity, i.e., $n = P(e|T)$, l denote selectivity, i.e., $l = P(\neg e|\neg T)$, and h denote the prior probability of drug A being toxic, i.e., $h = P(T)$, then positive test validity is greater than negative test validity as long as the following inequality holds:

$$h^2(n - n^2) > (1 - h)^2(l - l^2) \quad (6)$$

Assuming maximal uncertainty about the toxicity of drug A, i.e., $P(T) = .5 = h$, this means that positive test validity, $P(T|e)$, is greater than negative test validity, $P(\neg T|\neg e)$, when selectivity (l) is higher than sensitivity (n) and $n + l > 1$. As Oaksford and Hahn (2004) argue, these are conditions often met in practice for a variety of clinical and psychological tests. Therefore, in a variety of settings, positive arguments are stronger than negative arguments.

Oaksford and Hahn (2004) also provide experimental evidence to the effect that positive arguments such as (3a) are indeed viewed as more convincing than their negative counterparts under the conditions just described. The evidence from their experiment further shows that people are sensitive to manipulations in the amount of evidence (one versus 50 studies or tests) as predicted by the account. Finally, participants were sensitive to the degree of prior belief a character in a dialogue initially displayed toward the conclusion, as the Bayesian account predicts. This finding captures the 'audience dependence' of argumentation assumed in the rhetorical research tradition (e.g., Perelman & Olbrechts-Tyteca, 1969).

Though these results are encouraging, they were drawn from a single experiment using only two topics of argument. It is consequently important to test the generality of the account with other materials. We do this in Experiment 1. Experiment 2 then examines further structural variants of arguments with negative evidence.

The experimental tests of the Bayesian account we provide here have a dual role. With regards to the development of a normative account of argument strength, participants' data provide basic modal intuitions about argument strength to supplement our own. This is important as it is only too easy to mistake one's own judgments for universal. At the same time, our Bayesian account provides (only) a computational level theory. A detailed psychological account of how people actually evaluate arguments is still required. Experimental data is essential for such an account as well. To this latter end, it is of interest not only whether or not people are sensitive to the basic factors posited by the account, but also how sensitive they are to their interactions and what limitations people show in practice. There are numerous finer interactions between prior belief, polarity and evidence predicted by the Bayesian account (for details see Oaksford and Hahn, 2004, in particular Figure 1); however, as these are of interest primarily for more detailed modeling they will not be considered here.

Experiment 1

The goal of the first experiment was to broaden the evidence for the Bayesian account in several ways. First, we wanted to see whether the same patterns would be obtained with a different set of scenarios. Second, we wanted to test a different evidence manipulation. Contrasting the number of experiments or studies that have failed to find an effect as we did in Oaksford and Hahn (2004) is a rather simplistic change that seems impossible to ignore. Consequently it is not obvious that other more subtle, in particular non-numerical manipulations will also affect people's judgments in the predicted way. As an alternative, the reliability of the source of the evidence instead of simply its quantity is manipulated in the present experiment. Third, and finally, we wanted to present people with a more naturalistic rendition of the differences in prior belief than used in Oaksford and Hahn (2004) where the participants of the fictitious dialogues said things such as "I weakly believe that this drug has no side effects". Though maximally clear with regards to what is being manipulated, the phrase is somewhat stilted und unnatural which might in itself call attention to it where more naturalistic phrases would simply be ignored.

For our materials, we devised four dialogues on topical issues: one on the benefits of privatizing public transport, one on the dangers of clone technology, one on the efficacy of international environmental laws, and a final one on the respective importance of nature and nurture in language acquisition. For example:

Brenda: *Do you think it is beneficial to privatise public transportation?*
Adam: *I am fairly convinced that it is beneficial to privatise public transportation.*
Brenda: *You can be more than fairly convinced; you can be certain that it is beneficial.*
Adam: *Why do you say that?*
Brenda: *Because I read a newspaper interview with the members of a non-governmental research body and they said that it is beneficial considering the improved service quality and the reduction in the overall operating costs.*

To allow comparisons between negative and positive evidence, each dialogue existed in a positive and a negative version. For example, privatization was argued to be beneficial in one variant and not beneficial in the other. For the prior belief manipulation, the addressee of the argument (here Adam) was either "fairly convinced" or "sort of believed" the proposition in question. Finally, there were variants with high and low reliability evidence. For these, we manipulated the source of the evidence. In the above example, a non-governmental research body was contrasted with a TV street interview of passersby. Each participant received all variants of each topic to rate.

Method

Participants 73, predominantly female, Cardiff University students and staff.

Materials & Procedure The dialogues were presented to participants in booklets, with 10 different random orders. Each sample dialogue was followed by a ratings scale. For example:

> "How convinced do you think Adam should *now* be that it is beneficial to privatize public transportation? Please indicate your response by putting a tick (√) in the corresponding box in the 0 (not convinced at all) to 10 (totally convinced) scale below."

The booklet took about 15 minutes to complete and participants were tested individually or in small groups (without talking to each other).

Results & Discussion

In a 4 (Topic) × 2 (Polarity) × 2 (Reliability) × 2 (Prior Belief) within subjects analysis of variance with acceptance rating as the DV, we found significant main effects of all three manipulated factors. First, there was a main effect of polarity, positive arguments ($\overline{m}$ = 4.53, SE = .08) were more convincing than negative arguments ($\overline{m}$ = 4.27, SE = .08), $F(1,72)=11.58$, p = .001, MSE = 40.35. Second, the arguments with a higher reliability source ($\overline{m}$ = 5.58, SE = .07) were more convincing than those with a less reliable source of evidence ($\overline{m}$ = 3.21, SE = .07), $F(1,72) = 198.63$, p < .0001, MSE = 3272.79. Third, arguments with a higher degree of prior belief ($\overline{m}$ = 4.47, SE = .08) were more convincing than arguments with a lower degree of prior belief ($\overline{m}$ = 4.32, SE = .08), $F(1,72) = 5.96$, p = .017, MSE = 11.94. Like Oaksford and Hahn (2004), this experiment also showed differences between topics that are consistent with a Bayesian account of content effects and with Toulmin's (1992) position that the criteria for argument acceptance varies with subject matter.

In summary, the present experiment replicated with different topics, a different evidence manipulation, and a different wording for the prior belief manipulation, all three main findings of Oaksford and Hahn's (2004) study, suggesting that people's assessment of how convincing an argument is, is indeed influenced by the three main factors posited by the Bayesian account.

Experiment 2

The previous experiment provides evidence that arguments based on negative evidence are not always unacceptable and that their degree of acceptability can be explained by a Bayesian account. However, one might still query whether this amounts to a satisfactory treatment of the argument from ignorance. This is because the textbook example of the ghosts (1) differs from all of our experimental materials so far in one, possibly important way. The argument for ghosts not only involves negative evidence, but also a flip in polarity between evidence and conclusion: negative evidence is

provided to support the *positive* existence of something. In other words the inference is of the form:

(*7*) *not proven* (*not exist*) → *exist*

as opposed to merely:

(*8*) *not proven* (*exist*) → *not exist*

All our examples so far, both in Experiment 1 and in Oaksford and Hahn (2004) arguably have the latter structure not the former. But it may be the opposite polarity case (7) that constitutes the true fallacy of the argument from ignorance.

Classical logic licenses an inference from *not*(*not p*) to *p*, but not the inference underlying (7) which might be rendered as:

(*9*) *not says* (*not p*) → *?*

This is because when one has not said '*not p*,' one can either have said '*p*' or not spoken about '*p*' at all. For example, in an argument one might defend oneself with the claim "I didn't say you were rude", which could be true either because one had specifically claimed the opposite or because one had not mentioned rudeness at all. So maybe nothing at all can be inferred in such cases?

Walton (1992) first drew attention to parallels between the argument from ignorance and the negation-as-failure procedure (Clark, 1978) within AI. Knowledge-based systems frequently support the inference that a proposition is false—so its negation is true—because it can not be proved from the contents of the data base. This type of inference relies on the concept of epistemic closure (De Cornulier, 1988; Walton, 1992) or the closed world assumption (e.g., Reiter, 1980).

Walton (1992) also provides a real-world example of negation by failure: suppose the point at issue is whether the 13:00 train from London, Kings Cross to Newcastle stops at Hatfield. If the timetable is consulted and it is found that Hatfield is not mentioned as one of the stops, then it can be inferred that the train does not stop there. That is, it is assumed that the timetable is epistemically closed such that if there were further stops they would have been included.

But should epistemic closure ever license a negative evidence inference to the positive existence of something such as required by (1), and, moreover, can epistemic closure be captured within a probabilistic treatment of argument strength?

The case for both points will be made with an informal example: imagine your colleagues at work are gathering for a staff picnic. You ask the person organizing the picnic whether your colleague Smith is coming, to which you receive the reply that "Smith hasn't said that he's not coming". Should this allow you to infer that he is in fact coming, or has he simply failed to send the required email reply. Your confidence that Smith will be attending will vary depending on the number of people that have replied. If you are told that no one has replied so far, assuming Smith's attendance seems premature; if by contrast you are told that everyone has replied, you would be assured of his presence. In between these two extremes your degree of confidence will be scaled: the more people have replied the more confident you will be. In other words, the epistemic closure of the database in question (the e-mail inbox of the organizer) can vary from no closure whatsoever to complete closure, giving rise to corresponding changes in the probability that *not says* (*not p*) does in fact suggest that *p*.

To demonstrate and test this idea we devised four separate topics that varied in the amount of epistemic closure they involved. For each of these topics, we generated four possible combinations of evidential and conclusion polarity. One topic, for example, concerned the existence of a secret treaty between two countries, the evidence for which stemmed from newspaper archives. At stake could be either the existence or non-existence of the treaty. The evidence could either be positive or negative (says vs. *not* says) and could either affirm ('exists') or deny ('*not* exists') the conclusion, giving rise to the following cases, concerning newspaper reports of a secret treaty:

(a) Article says: exists → treaty exists
(b) *not* (Article says: *not* exists)) → exists
(c) Article says: *not* exists → *not* exists
(d) *not* (Article says: exists) → *not* exists

Case (d) corresponds structurally to (8) and the negative evidence cases tested experimentally so far, whereas (b) is an instance of the opposite polarity negative evidence case (7) exemplified by the ghosts example.

The other topics concerned the presence or absence of a book on a library shelf as function of what it says in the electronic catalogue, whether or not a train stopped depending on what it said in the timetable, and whether or not part of a routine hospital procedure had been carried out depending on whether or not it was mentioned in the medical notes. The examples were chosen to give variation in the degree of closure, but also to include cases where opposite polarity negative evidence inferences can be non-fallacious.

Method

Participants 72**,** predominantly female, Cardiff University undergraduates.

Materials & Procedure Though the individual arguments were readily comprehensible when read in isolation, reading all four structural variants of a topic was quite confusing. Consequently we used a Latin square confounded design whereby each participant was presented with only one structural variant (a,b,c, or d) of each topic (1,2,3,4), seeing for example, 1a, 2b, 3c, and 4d. We generated all 24 different combinations, making 24 different booklets containing the arguments.

Results The main results can be seen in a qualitative way from Fig 1. First, there is considerable variation across the four topics, not just in the overall ratings, but in the patterns among the 4 structural variants for each topic. In other words, it is not just that firmer conclusions overall can be drawn from say, train time tables than from newspaper reports; rather, the relative strength, for example, of inferences from positive vs. negative evidence also varies. This is confirmed statistically by significant main effects of both topic (hospital, paper etc.) and structural variant (a,b,c,d), with $F(3,217) = 35.55$, $p < .0001$, and $F(3,217) = 19.05$, $p < .0001$, respectively, and a significant interaction between topic and structural variant $F(9,217) = 3.71$ $p < .0001$, and MSE = 5.59. In keeping with the Latin squared confounded design, as in Kirk (1982), all these overall analyses were conducted on the residuals once participant variance has been removed, not on the raw ratings, because different participants contribute to different cells.

Second, both negative evidence cases (b and d) are always weaker than straightforward positive evidence (a), confirming that negative evidence is typically seen as less convincing. However, the ratings for the opposite polarity negative evidence case are significantly different from 0, $t(71) = 9.16$, $p < .0001$, with $\overline{m} = 3.36$, and SE = .37. Moreover, the opposite polarity negative evidence case (b) is not always less compelling than the same polarity negative evidence case (d) which we had examined in Exp. 1. While (b) is worse than (d) for two of the topics, it is actually significantly better for Topic 1, the library case $t(34) = 2.39$, $p < .05$ (with $\overline{m} = 6.16$, SE = .65, and $\overline{m} = 4.0$, SE = .64 for b and d respectively). There was also no significant difference between (b) and (d) for Topic 3 about the newspaper archive and the existence of a secret treaty.

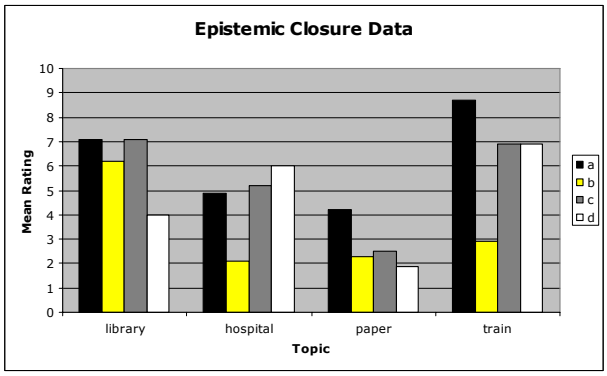


Figure 1: Mean ratings from Experiment 2 across topics.

Discussion

Limiting the argument from ignorance to negative evidence cases with opposite polarity such as in the ghosts example (1) no more gives a structure that is *always* fallacious, than did the wider interpretation based on negative evidence in general. In the opposite polarity case, we also observed graded differences in argument strength and these differences seem amenable to a probabilistic treatment. Epistemic closure can be a matter of degree so that the probability of *p* given *not says* (*not p*) can vary throughout the interval from 0 to 1. Under the typical conditions described in the *Introduction*, negative evidence arguments will be weaker than their positive counterparts whether they involve a switch in polarity or not. However, whether or not such an argument is fallacious depends not on its logical structure but the probabilities involved.

The Burden of Proof

A third and final class of arguments from ignorance distinguished in the philosophical literature are characterized as illegitimate attempts to shift the burden of proof (Walton, 1992). Indeed, several authors wish to restrict the argument from ignorance to this final case (e.g., Copi & Cohen, 1990). Consequently, it merits further inspection. In light of the fact that some arguments from negative evidence can be reasonable, and in light of the fact that even inferences to the positive existence of something can be inferred from negative evidence under appropriate conditions, violation of the burden of proof is invoked to explain arguments such as the classic ghost example (1). The idea is that the pragmatics of argument (at least for the 'information-seeking' discourse relevant here) demand that whoever makes a claim has to provide reasons for this claim when challenged. Pointing out that no-one has failed to disprove their existence as a reason for believing in ghosts is an illegitimate attempt to shift that burden onto the other party instead of providing an adequate reason oneself. However, this idea of shifting the burden of proof does not *explain* why the ghost example is a fallacy, it merely *labels* it. As we've seen in the preceding sections, negative evidence *can* constitute a good reason for believing something. What's more, there are combinations of test sensitivity, specificity and priors where negative evidence is *more compelling* than positive evidence. This means one has to be able to explain why negative evidence *vis* a *vis* ghosts is not of this kind. Without such an explanation it remains entirely unclear why it is not an adequate reason in this case also and as such does not shift the burden of proof. Consequently, 'shifting the burden of proof' doesn't explain an argument's weakness, it presupposes it.

Lacking independent definition 'shifting the burden of proof' is not, in fact, a separate category of arguments from ignorance, but merely a catchphrase to label ones that are weak. The real reason we consider negative evidence on ghosts to be weak is because of the lack of sensitivity (ability to detect ghosts) we attribute to our tests as well as our low prior belief in their existence.

Conclusions

We have presented analyses and evidence to suggest that the argument from ignorance need not constitute a fallacy, and that the conditions under which it is a good as opposed to a bad argument are well captured by a Bayesian account.

It is an advantage of the Bayesian framework that it has the potential to provide a normative account of argument strength not just for the argument from ignorance, but for argumentation more generally. Acceptance as a normative

standard would require demonstration of the fact that key intuitions about the relative strength of arguments are preserved within a Bayesian reconstruction. For this, it is clearly desirable to look at more than one kind of argument.

The traditional catalogue of argument fallacies provide an ideal testing ground for such an endeavor, and we have recently sought to extend the account to other fallacies (Hahn & Oaksford, submitted). We have found that intuitive variation in other fallacies such as circular arguments or slippery slope arguments also seem well-captured, recommending the Bayesian account as a general framework for evaluating argument strength.

Moreover, such an attempt to capture everyday, informal reasoning in Bayesian terms links well with several other developments over the past decade or so. For one, it can trade on similar approaches to scientific inference within the philosophy of science (e.g., Howson & Urbach, 1989; Earman, 1992; but see also e.g., Miller, 1994). It also connects with the trend in AI knowledge representation away from attempts to extend logical approaches to account for uncertain reasoning in favor of Bayesian approaches to uncertainty (see e.g., Pearl, 1988).

Currently, work on argumentation within AI, for example in the context of legal argumentation (of many see e.g. Gordon, 1995 or Sartor, 1995), is still dominated by logic–based approaches. Likewise, the philosophical literature on argumentation still has a strong logic-based strand (see e.g, Prakken, 2002). However, none of this work has yet provided a satisfactory treatment of the fallacies.

While the development of normative accounts is of central interest, acquiring a broad empirical base is at least as important for a fuller study of human argumentation. While well-chosen examples and associated intuitions are essential to theory development, it is crucial to supplement these with experimental data. At present, there is considerable experimental work that is of general interest to anyone engaged with argumentation such as the sizeable literature on persuasion (see e.g.,Johnson, Maio & Smith-McLallan, 2005). However, the volume of experimental work on the comparative strength of arguments is extremely limited (Rips, 2002; Neumann & Weizman, 2003; Oaksford & Hahn, 2004). This, we hope, is set to change.

References

Clark, K. L. (1978). Negation as failure. In H. Gallaire & J. Minker (Eds.), *Logic and databases* (pp. 293-322). New York: Plenum Press.

Copi, I. M., & Cohen, C. (1990). *Introduction to logic* (8th Ed.). New York: Macmillan Press.

De Cornulier, B. (1988). *Knowing whether, knowing who, and epistemic closure.* In M. Meyer (Ed.), Questions and questioning (pp. 182-192). Berlin: Walter de Gruyter.

Earman, J. (1992). *Bayes or bust*? Cambridge,MIT Press.

Eeemeren, F.H. van & Grootendorst, R. (2003). *A systematic theory of argumentation. The pragma-dialectical approach*. Cambridge: Cambridge University Press.

Gordon, T.F. (1995). *The pleadings game*: *An artificial intelligence model of procedural justice*. Dordrecht: Kluwer.

Green, D. M., & Swets, J. A., (1966). *Signal detection theory and psychophysics.* New York: Wiley.

Hahn, U. & Oaksford, M. (submitted). *A Bayesian approach to informal argument fallacies*.

Hamblin, C.L. (1970). *Fallacies*. London: Methuen.

Howson, C., & Urbach, P. (1989). *Scientific reasoning: The Bayesian approach*. La Salle, Illinois: Open Court.

Johnson, B. T., Maio, G. R., & Smith-McLallan, A. (2005). Communication and attitude change. In D. Albarracin, B. T. Johnson, & M. P. Zanna (Eds.), *The handbook of attitudes and attitude change: Basic principles*. Mahwah, NJ: Erlbaum.

Kirk R.E. (1982). *Experimental design: procedures for the behavioral sciences*. Pacific Grove: Brooks/Cole.

Miller, D. (1994). *Critical rationalism: A restatement and a defence*. La Salle, IL: Open Court.

Neuman, Y., & Weizman, E. (2003). The role of text representation in students' ability to identify fallacious arguments. *Quarterly Journal of Experimental Psychology*, 56A, 849-864.

Oaksford, M., & Hahn, U. (2004). A Bayesian approach to the argument from ignorance. *Canadian Journal of Experimental Psychology*, 58, 75-85.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems.* San Mateo, CA: Morgan Kaufman.

Perelman, C., & Olbrechts-Tyteca, L. (1969). *The new rhetoric: A treatise on argumentation.* Notre Dame, Indiana: University of Notre Dame Press.

Prakken, H. & Vreeswijk, G. (2002) Logical systems for defeasible argumentation. In D. Gabbay and F. Guenthner (Eds.), *Handbook of philosophical logic*, 2nd Edition, Vol. 4, (pp. 219-318). Dordrecht: Kluwer

Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13, 81-132.

Rips, L. J. (2002) Circular reasoning. *Cognitive Science*, 26, 767-795.

Sartor, G. (1995) Defeasibility in legal reasoning. In Z. Bankowski, I. White, & U. Hahn (Eds.), *Informatics and the foundations of legal reasoning* (pp. 119-157). Dordrecht: Kluwer.

Toulmin, S. (1992). Logic, rhetoric, and reason: Redressing the balance. In F. H. van Eemeren, R. Grootendorst, J. A. Blair, & C. A. Willard (Eds.), *Argumentation illuminated* (pp. 3–11). Amsterdam: Sic Sat.

Walton, D. N. (1992). Nonfallacious arguments from ignorance. *American Philosophical Quarterly*, 29, 381-387.

Walton, D.N. (1995). *A pragmatic theory of fallacy*. Tuscaloosa/London: The University of Alabama Press.

Walton, D. N. (1996). *Arguments from ignorance*. Philadelphia, PA: Pennsylvania State University Press.