Research Writing in Computer Science

Version 2 (1998) Kevin B. Korb

Dept. of Computer Science
Monash University
Clayton, Victoria 3168
Australia
korb@cs.monash.edu.au

1 Introduction

Writing is an essential ingredient in the life of anyone doing research. This is true of industrial research and development just as much as of academic research, although engineers and scientists truly bad at (or afraid of) writing can usually find ways of minimizing their exposure to the activity — such as by relying upon peers who are better suited to the task emotionally. This is no benefit to those researchers and their work, however: it is patent that writing up one's research helps one to (re)conceptualize it (cf. Zinsser, 1988). In this paper I shall summarize what research writing is in computer science and how one goes about it. The ideas expressed here should be as applicable to short research notes as to honours theses. I shall assume that the readers are (primarily) honours students in computer science who have had little exposure to the task of writing research, whether or not they have read others' research papers. Nevertheless, my remarks are likely to be as helpful (or as unhelpful) to those working in other disciplines: there is nothing special about good writing in the different technical disciplines except for its rarity.

What is research? *Webster's Collegiate Dictionary* reports that it is "investigation or experimentation aimed at the discovery of facts, revision of accepted theories or laws in the light of new facts, or practical application of such." Such investigation is of little interest, beyond that of the researchers involved, if it is not communicated effectively to others. Research writing is the primary means of communicating the fruits of research to others, for them to assess the value of the research and (as in marking papers or in judging promotions) the abilities and achievements of the researchers.

The key feature of research writing is that it should be reporting *discoveries* — it should be adding, or attempting to add, to the body of human knowledge. It should not (merely) be rehashing others' ideas, nor should it be reporting the outcomes of experiments when those outcomes are already well known. In consequence, a research report must do *something* to contribute to knowledge. The most dramatic form this might take would be to present an original theory of some domain: such as Einstein's theories of relativity. Closer to home (computer science) would be such cases as John Backus's invention of the functional language FFP and his related claim that functional languages are superior to the more common procedural languages (such as C) for designing and implementing algorithms. One does not need to have landed upon some astounding discovery to have the germ of a good research paper, of course. Lesser discoveries qualify. One may also contribute to our collective knowledge by bringing a new perspective to bear upon old issues, by, for example, inventing new arguments for or against theses promulgated by others. Argumentation is a critical part of all research writing, after all. New ideas which are not sustained, developed and argued for will likely be lost to the scientific community, whatever their intrinsic merits. It is obligatory for every research writer to explain not just what his or her thesis is, but why a reasonable person should believe it to

be true. Hence, the selection and presentation of evidence, and the construction of arguments relating the evidence to your conclusions, is an integral part of research writing.

This is true even of survey papers. It is not true of annotated bibliographies: these, however useful generally, present nothing new to an expert audience. Survey papers on the other hand should not merely be lists of books and articles relevant to a subject, they should briefly describe the contents of such research works and their interrelations. In describing relationships between different works survey papers make claims about them: what the relations are; who contributed key ideas to a research area and what is derivative; why one research approach appears more promising than another; etc. All of these claims, even in a survey paper, need to be backed up, if only briefly, by evidence extracted from the writings discussed. A really good survey paper (or textbook), packed with really good mini-arguments, can contribute as much to the activities of a research community as a report of original research (and get more citations as well; cf. Angluin and Smith, 1983, and Oliver and Baxter, 1994).

Perhaps the most widespread misunderstanding that students have about their research papers is that in their cases, because they are amateurs (thus far) at research, the purpose of writing them is to demonstrate their private knowledge, rather than to advance the collective knowledge. In consequence, many student papers read as laundry lists of the intellectual products of others, rather than as sustained arguments for some scientific thesis. But demonstrating technical mastery of a subject is the role of examinations, not research papers. A secondary purpose of a research assignment may be to force students to acquire such background knowledge as the assignment may require. But the primary purpose is to educate students in the process of conducting real research — and that purpose is not served by repeating the conclusions of others, but only by the investigation, analysis, criticism, *thinking* required of true research.

In the remainder of this paper I shall report on some practical techniques for the construction of research papers: the mechanics of paper writing; English style and grammar; analyzing and building arguments; recognizing common fallacies and errors. My discussion is brief and limited, but also pointed. I include references to more complete treatments of specific areas by others in bibliographic notes throughout, so that those who wish may readily pursue the issues in greater depth.

2 The Mechanics of Writing

Douglas Adams (1979) has the best advice on this issue: DON'T PANIC!

My next best advice is: do what works for you. A lot of texts and teachers issue formulae for composition, such as: begin with an outline and iteratively deepen it ("top-down" composition), write introductions first, write introductions last, write your conclusion first, etc. Mostly they are either reporting what they have been taught or what happens to work for them. What happens to work for you may be different. Listen to them, but do what works. (The same applies to my suggestions throughout, of course.)

2.1 Getting Started on a Paper

A great many honours students seem to believe that the "writing up" occurs when the research ends. This leads to two HIGHLY adverse consequences: first, the writing is begun only immediately prior to the deadline, meaning that the paper is hurried, poorly organized and mostly unintelligible; second, the writing is begun only immediately prior to the deadline, meaning that the research being reported has been conducted without the benefit of reflection, self-criticism and the criticism of others — i.e., it has probably been done thoughtlessly. As a good rule of thumb, it is just never too early to begin writing. Even if you are writing before you have conducted any experiment or done any library research or talked over a possible

topic with anyone, writing down your thoughts about a problem cannot hurt and may well help. One professor I know advocates forcing new Ph.D. students to write for their first three months without any reference to what other researchers have done in the chosen field or problem area. As close attention to what others say about a problem may well lead to an unfortunate stultification of original thought, this has considerable merit. Without evidence one can at least speculate intelligently about what the evidence will reveal, or what lines of investigation will be fruitful.

Research is best written up "as you go." In the first instance, research writing is likely to take the form of notes. As more evidence is gathered (experimentally and from others' writings) the notes are likely to become too large to be effectively perused at one sitting. That is a hint that something more formal might be written up. If you seriously undertake to write every step of your research way, you will likely find that putting your thoughts in explicit form has innumerable advantages. Among others: it allows you to criticize your thoughts more effectively; it will change the direction of the flow of your ideas, often into more effective channels; it will subsequently remind you of plans and ideas that otherwise would have been forgotten. In general, we use language to think through our problems, and written language can and ought to be a vital part of the process.

2.2 Getting a Paper Started

The first thing any reader wants to know is: What's in it for me? If your abstract or introduction cannot answer that question, it is unlikely that you will get a second chance. (Of course, your poor supervisor probably does not have the option of just dropping your paper in a can, but the desire to do so may well be reflected in your grade.) Your introductory matter should state your main thesis and briefly describe the approach you will take to defending or investigating it. The main problem to be solved in writing research reports is how to communicate your main ideas to your reader, so the earlier and clearer you state your theses the better. The mysterious paper is an unread paper.

2.3 Organizing Research Papers

Outlines

Personally, I find little value in outlines. Especially, the commonly advocated use for them — write them first thing and then "flesh them out" — is an unlikely idea. People do not think that way, and writing *is* thinking. If you wait to start writing until you have a clear outline in mind, then you have waited far too long to start writing. Furthermore, if you can write a paper without changing your mind in mid-stream about its structure, progress or constituent ideas, then you are writing a trivial paper.

Order

The order in which you write a paper is not important. Whatever you prefer to work on next, or is fresh in mind, should be worked on next. I do find it useful often to begin with my conclusion, even prior to research. No one performs experiments or investigations without some idea of what one expects to find. If you write down your expectations clearly in advance, you are likely to find better ways of testing them. This can give your whole research effort a better focus.

The order in which the paper is to be read is another matter. Obviously, the abstract and introduction come first. The two have related but different purposes. Both of them advertise the contents and importance of the paper. The abstract will be read by those who do not want to read your paper; for these, you need to give them the main message to take away

— your thesis and what problem it solves. Others will read an abstract to decide whether to read your paper; so your abstract also does a sales job. Your introduction should do the same for the same reasons, but also it sets out a context in which the rest of the paper can be read. Usually, a brief literature survey is in order, explaining what prior research is directly relevant, how it is relevant and how your own work goes beyond it.

Just as obviously the conclusion comes last (or, next to last if you want to append some speculation about what the future may bring). This should restate the thesis and summarize the evidence in its favor.

What comes in the middle depends entirely on what research it is that you are doing. The best way to answer questions about how to structure such things is to look at highly regarded papers in your discipline or sub-discipline. Fairly common in computer science would be the following overall structure for an experimental research paper:

1. Abstract.
2. Introduction, including a literature survey.
3. Methodology, problem statement, theoretical approach taken and experimental procedures used.
4. Experimental results, including summary statistics.
5. Interpretation, using the results to argue for a particular hypothesis.
6. Conclusion, summing up the experiment and its interpretation and the wider significance of the study.
7. References.
8. Appendices (if any) report details of the experimental results.

2.4 Rewriting and Rethinking

As writing is thinking, rewriting is rethinking. No one can get a serious research paper right the first time. (I think. Someone has claimed this ability to me, but I fail to believe it.) But this is good news: the pressure is off! You can forget about wording things in just the right way or getting the relation between your evidence and your thesis just right at the beginning. Assuming you are not writing at 3 am the morning the paper is due, you will have plenty of time to rethink such matters later. It is far more important that you make a genuine, if bad, early effort at these things than that you complete your final draft right away. Bad efforts can be improved in all kinds of ways (I suppose that's what makes them bad).

2.5 Research Ideas

Where do research ideas come from? Pretty often they come from research supervisors, but that is a rather lazy answer. Nobody lacks ideas. As Gold (1994) points out, it is a goal of many forms of meditation to free the mind of ideas — and that is extremely difficult to achieve because the mind seems always full of ideas. When you first try to meditate, there is an unending stream of them. The only problem with them is that they are hardly ever ponderings on the deep problems of understanding the universe; they are mostly just a lot of rubbish. So the problem is not one have having ideas, but one of having *good* ideas.

But this is no more of a real problem than the artificial problem of trying to write the final draft on the very first try. There is no point to it. A good research program can as readily be grounded on a bad research idea as on a good research idea — so long as you are

prepared to improve your research idea as you go. As you try to test your (bad) idea, to find out what questions it raises, how others have dealt with it or something similar, and so on, you will have any number of opportunities to revise your guiding idea or to throw it out for something very much different and better.

Puzzlement

It almost goes without saying that many good research ideas come out of someone being puzzled by something, by finding some confluence of events being unusual or unexpected. If you are puzzled by something you hear or read, consider taking on the task of eliminating your puzzlement (or disproving what you heard or read that induced puzzlement). You may be on your way to an important discovery.

Exploratory Reading

Reading for research ideas can be very productive, so long as it is not used as an excuse to procrastinate on starting on a research problem already identified. When you read, if you employ the methods of argument analysis described below, you may well find other researchers pushing lame arguments. This in turn may lead to a productive research idea, for if those researchers' ideas are wrong, something else must be right.

2.6 Bibliographic Notes

Gold (1994) provides an entertaining and useful guide to writing honours theses, even though it is oriented toward philosophy. A copy may be obtained from the Philosophy Department at Monash. Barzun and Graff (1992) is a guide to research and writing for historians, however its advice will benefit researchers of all persuasions. Zinsser's *On Writing Well* (1990) is another excellent guide. Moxley (1992) is oriented specifically to academic researchers and is well informed on recent research in linguistics and communication. Since you will be writing in order to be read, you might want to look at Adler and van Doren's interesting *How to Read a Book* (1972).

3 Style and Grammar

3.1 Clarity

Being clear is, I think, the real point behind every good suggestion regarding writing style and grammar. Clarity of thought in writing is remarkably hard to achieve and easy to lose. The best advice regarding style that I know of is that of George Orwell (1950) in "Politics and the English Language." He points out that people who specifically wish to *hide* or *obscure* their messages, such as many politicians and military leaders, employ long-winded, abstract, jargonistic or euphemistic language. They are largely successful in failing to communicate. The message, of course, is that in order to communicate effectively we should do the opposite: employ concise, vivid, concrete, ordinary language.

Concision aids comprehension by avoiding the introduction of words that do not contribute to meaning and sentences that do not contribute to your argument. For example, instead of "The experiment was not conducive to productive modes of thought for the research project" try "The experiment failed." The point is not, of course, that each sentence should be as short as possible, but that no sentence should be padded out simply for the sake of length and obscurity. Variations in vocabulary, sentence structure and length are all essential to retaining the interest of your reader, but such variations should serve the purposes of providing new and relevant information.

Being *vivid* and *concrete* means means employing imagery, metaphors, examples which people can visualize or imagine, as opposed to abstract ideas which they cannot. Abstractions are, of course, essential to any scientific writing, but when they are supplemented or illustrated with concrete examples, they will be more readily understood. There is a considerable amount of empirical evidence from cognitive psychology supporting this assertion: people remember statements better and longer when they are formulated in concrete terms; people also reason better about problems when they are framed in concrete rather than abstract language (Nisbett and Ross, 1980). In software engineering, it is one thing to be exhorted in the abstract to employ good practice in software design and development, it is quite another educational experience to be directly confronted with software that has been hacked together in an undifferentiated mess.

Using *ordinary* language, instead of jargon and euphemism, renders your writing accessible to a wider audience and often avoids unnecessary ambiguity. If your target audience is quite specifically a technical audience, then using computer science jargon without explanation may be the most effective means of communication. But if the readership is potentially broader than that, the standard jargon you use should be expunged or explained. Euphemism means the use of an imprecise, but supposedly less offensive, expression for a direct, and therefore supposedly offensive, expression. Instead of *dying* people talk about *passing away*; instead of *computer crime* people talk about *hacking*. In writing, I think there is no excuse for using a less precise expression when a more precise expression is meant.

Examples

My favorite example of obfuscatory language, deadening impact by replacing vivid and precise imagery with turgid and imageless prose, is one produced by George Orwell (1950):

> Objective consideration of contemporary phenomena compels the conclusion that success or failure in competitive activities exhibits no tendency to be commensurate with innate capacity, but that a considerable element of the unpredictable must invariably be taken into account.

This is Orwell's "translation" into "modern" (i.e., bureaucratic and lifeless) prose of a sentence from Ecclesiastes:

> I returned and saw under the sun, that the race is not to the swift, nor the battle to the strong, neither yet bread to the wise, nor yet riches to men of understanding, nor yet favour to men of skill; but time and chance happeneth to them all.

It is easy enough to find examples of bad writing in computer science. Indeed, they are more common than examples of good writing, unfortunately. Consider the following from the introduction to *Handbook of Software Engineering* by Charles Vick (Vick and Ramamoorthy, 1984, p. xxxii):

> As the various elements and phases of a software engineering environment are developed, standards for design, development, and testing must also be developed. This will be critical to enforcing the disciplines of the system. Automated checking of adherence to standards must be embedded in the superstructures. Metrics that are meaningful, and (in fact) measurable, must also be identified in support of each phase. Metrics are not only necessary for application to functional and nonfunctional requirements but to the predictive models identified earlier.

You should bear in mind that I found this passage almost immediately upon turning to a shelf of computer science books: such obscure, imageless, redundant and even incoherent

text is not hard to find in technical writing. I shall make a few observations on the quoted text. It is unclear what distinction the author is trying to make in the first sentence between "elements" and "phases", nor is it clear that there *are* phases of an *environment* in normal application development (as opposed to phases of product development). The second sentence appears to mean simply that standards are necessary for the enforcement of standards; this truism might have been better left unsaid. The next sentence invokes "superstructures" when software development processes appear to be meant. The fourth sentence suggests that we should prefer metrics that are measurable over those which are not. As metrics are not themselves typically measured, but used to measure, we could agree with this sentence (as another truism) if it had asserted the need for metrics with which we can measure. The last sentence talks of "functional and nonfunctional" requirements; since that should pretty much cover all requirements it would have been better to say so directly. In general, the writer has used imprecise and convoluted language to express some extremely simple ideas. The reader is either left with the impression of deep thinking (no doubt the intended effect) or with the impression of a near total absence of thinking, which appears to be closer to the truth.

Computer science does have its (small) share of decent writers (thinkers). Contrast the above with the following passage, also on software metrics:

> [A] fallacious thought mode is expressed in the very unit of effort used in estimating and scheduling: the man-month. Cost does indeed vary as the product of the number of men and the number of months. Progress does not.

This simple expression of a simple, but important, insight by Fred Brooks (1975) had — and deserved to have — more impact on the profession of software engineering than the total contents of Vick and Ramamoorthy's handbook.

3.2 Reading Your Writing

Philosophers (so far) by and large stick to the medieval habit of reading their papers out loud when giving public lectures. Hopefully,[1] philosophers will catch up with decades-old teaching technology, such as the overhead projector, sometime in the next century. There is one useful aspect to reading papers out loud, however. This helps the writer to assess the cadence and style of the writing. A paper (paragraph, sentence) that *sounds* bad when read aloud will read badly as well. It is no bad practice, therefore, when working on a final draft, to read the paper, or parts of it, to the wall (or a long-suffering friend or partner); you are very likely to find ways to improve the wording and intelligibility of your writing.

3.3 Punctuation and Grammar

I have two pieces of specific advice on punctuation, which many in Australia will disagree with, presumably because they were taught to disagree with it. The advice here is mine, but is supported by many professional style guides, such as the *Publication Manual of the American Psychological Association* and the *Chicago Manual of Style*. First, quotation marks should generally incorporate commas and periods even when that punctuation was not employed by the source. This is because commas and periods outside the quotation are distracting and ugly; there is little value in making a big point that a comma was not in the source text.[2]

[1] Some people pretend not to understand this use of the word "hopefully". It is, however, perfectly clear in context — meaning "one hopes that" — and has been used in that capacity for more than one hundred years. Compare this usage with that of "fortunately", "conceivably", "happily", "notably" *ad nauseam*. For the contrary opinion see "bunk" under "Grammar" below.

[2] Matters are different with semi-colons, colons, dashes and question marks because they are special-purpose punctuation; quoting them might suggest that your source had some familiarity with what lies outside the quote, which would be odd.

I violate this principle when I am not quoting but using quote marks to indicate a precise syntax, such as a particular word or a logical expression (periods do not occur in logical formulae — unless you are a follower of Bertrand Russell). Second, footnote numbers should occur *after* punctuation, again because they are less distracting there. When using numbers in text it is also less ambiguous, since footnote 2 might otherwise be read as the squaring operation. Whatever you do with your punctuation, you should be consistent about it.

A common, but awful, piece of advice on grammar is to avoid the use of the first person ("I" and "we") when writing scientific prose. A watered-down version of this suggestion is to use "we" in place of "I" even when you are the sole author. Both affectations — avoiding the first person altogether or diluting the singular into a plural first person — fail to inform the reader of something useful and without any compensation, such as simpler phrasing. If you write "It has been discovered that ..." instead of "I discovered that ...," then the reader is likely to misunderstand who has done the discovering. If you write instead "We have discovered that ...," then readers will not know which results are yours alone and which have benefited from some unnamed peers. Humility is a worthy trait, but false humility at the cost of precision is no virtue.

Good grammar is an aid to comprehension. That is its purpose: to aid listeners and readers in parsing and interpreting sentences. All the rest of grammar is bunk. You should not worry about such dicta as "Don't split your infinitives." English was built on the urge to freely emphasize adverbs by splitting infinitives with them; the reformists are wasting your time and theirs. For such reasons I suggest you avoid, for example, Strunk and White's famous style guide (1979). It is full of bad advice. On the other hand you definitely should use a variety of guides to style, punctuation, etc. I recommend some below.

3.4 Bibliographic Notes

Fowler's *Modern English Usage* (1983) is an opinionated, anglocentric and highly valuable review of the accurate use of English. S.I. Hayakawa (1987) demonstrates the sleepy senator's intimate command of nuances of meaning; extremely useful for getting at, for example, the difference between "affect" and "emotion". Robert Graves and Alan Hodges' *Reader over your Shoulder* (1947) elaborate sensible "Principles of Clear Statement." They also dissect various well-known writers in Part II in an excercise that reveals many common stylistic errors, while occasionally becoming stiff-necked and implausible.

Specifically on style guides, I strongly prefer the serious style guides used by professional editors to the popular style guides (such as Strunk and White, 1979). I recommend the APA's publication manual, the *Chicago Manual of Style* or some equivalent (others are in *References* below). Also, good dictionaries offer stylistic advice. For example, the *Oxford English Dictionary* says of the ending *-ise* (for *-ize*): "There is no reason why in English the special French spelling should be followed, in opposition to that which is at once etymological and phonetic. In this Dictionary the termination is uniformly written *-ize*."

4 Argumentation

Here I describe various analytic techniques for understanding arguments, which is a prerequisite to composing good arguments. What I can do within the framework of this introductory paper, however, is quite limited. So, although I believe my remarks will be helpful, far more valuable than merely reading them will be applying them to arguments you find in research articles and newspaper opinion columns.

4.1 Argument Analysis

As I have already mentioned, every research report contains one or more arguments. The paper you are reading now contains numerous arguments even though it is simply an introductory survey. For example, it contains an argument that argumentation is a necessary component of research writing. To employ arguments effectively in your writing you first need to know how to analyze and evaluate arguments, in order either to avoid the construction of poor arguments in the first place or, more likely, to be able to recognize their poverty and repair them prior to the final draft. As an illustration, let us analyze my own argument on behalf of argumentation as it appeared in the introduction and see what we can make of it. The main relevant passage was:

> Argumentation is a critical part of all research writing, after all. New ideas which are not sustained, developed and argued for will likely be lost to the scientific community, whatever their intrinsic merits. It is obligatory for every research writer to explain not just what his or her thesis is, but why a reasonable person should believe it to be true. Hence, the selection and presentation of evidence, and the construction of arguments relating the evidence to your conclusions, is an integral part of research writing.

The first step in argument analysis is to identify the main conclusion. You will not be able to assess an argument if you cannot determine what it is an argument for (or against). In this case, the conclusion is pretty clearly the first sentence — the following sentences appear to support it and also the word "hence" always introduces some form of conclusion, intermediate or final, and in this case it introduces the final sentence, which is just an elaboration of the first sentence. To get to the core of an argument it is often best to ignore the ornamental details writers pack into their sentences and concentrate on the *propositions* (such as the meanings of the sentences) they are asserting, whatever their form. (Of course, some details can be crucial to identifying the propositions.) I shall label the conclusion of the above argument (1) and report it thus:

> (1) Argumentation is a critical part of all research writing.

The next question is: what premises are offered in support of this conclusion? The main idea of the second sentence appears to be:

> (2) New ideas which are not supported by argument will be lost to the scientific community.

The third sentence supports (2) by suggesting that, lacking an argument, an idea is lost because readers ("reasonable persons") will be left unpersuaded. Hence,

> (3) Unargued new ideas will not be believed.

The intended inferential structure of the argument can now be seen: (3) $\rightarrow$ (2) $\rightarrow$ (1). We can assess this argument as it is in order to locate the holes in it. The relevant concept of logic for such assessment is *validity*. An argument is valid if and only if its premises logically imply its conclusions — that is, if it is *impossible* for its conclusions to be false if its premises are true; an argument is also *sound* if additionally its premises are true. Not very many arguments outside of the professional mathematical literature will stand up to such a severe test as validity, and certainly mine does not. Consider the inference (2) $\rightarrow$ (1). Is it *possible* for (2) to be true while (1) is not? Of course. We can invent a counterexample. If, for example, no one cared whether research writing aided the retention of new ideas, then the loss of ideas due to a failure to argue on their behalf would be of no interest. Implicitly,

then, my argument rests upon the claim that it is a goal of research writing to introduce and maintain new ideas within the scientific research community. Actually, this is not very implicit since nearly that same statement was made earlier in my introduction. Still, if we were assessing only the argument in the sentences quoted, it would be clear enough that it relied upon this implicit claim, which, in a careful assessment, ought to be made explicit:

> (4') It is a goal of research writing to introduce and maintain new ideas within the scientific research community.

I indicate that this was not explicitly within the text by the apostrophe. Note that naive commentators in public life will sometimes jump upon the failure to make all premises explicit as a failure of the argument or a failure of the thought processes of someone arguing the other side of a dispute. When the implicit presuppositions of an argument are as plain and non-controversial as this, however, such a move is unlikely to be effective rhetorically. Much better advice is to apply the *principle of charity*: if the step from an invalid argument to a valid argument requires a suppressed premise which is known to be true, or likely to be true, or non-controversial, then assume that the argument contains that premise. At the same time, it is prudent to apply a *principle of minimality*: if the argument needs only some particular assumption to become valid, then do not assume anything beyond that.

By adopting these principles you may forgo a few opportunities to get cheap shots in on an opponent; however, you will have a much better idea as to what the real argument is about.

Getting back to my argument, we now have the subargument: (4') and (2) $\rightarrow$ (1). Is this valid? If we are to avoid pedantry, we should probably concede that it is. It appears to be a basic principle of everybody's theory of action that if you have a goal X [introduce new ideas] and an action (or a failure to act) Y [failure to argue] will defeat it, then you should do not-Y [argue].[3] The only reason to deny this reasoning would be if a more important goal supported Y; but it is plausible that such a goal would simultaneously override the writing of a research report.

We could now perform a similar analysis of proposition (3) and its relation to the rest of the argument. Were this a crucial argument in a research report it would be useful to continue the process until a complete graphical representation of a valid version of the argument were available (one can use AND/OR graphs for this purpose; see Winston, 1992). However my illustrative purposes have been served.

There are a number of points to this exercise, in addition to the two principles of analysis introduced above. One (implicit thus far) is that any argument can be augmented until it is deductively valid. Furthermore, careful analysis will reveal that almost every argument you encounter is initially invalid — there is *some*, perhaps weird, circumstance in which the premises could be true and the conclusion false (given always a standard, or explicitly requested, understanding of the terms used: it is not fair to use non-standard meanings of the words, unless the author invokes them). In some sense, then, it is strictly incorrect to claim that anyone's argument cannot be interpreted as a valid argument — this reveals your failure to recover the suppressed premises rather than your interlocutor's failure to have them. More realistically, someone can be charged with presenting an invalid argument in case the missing premises needed are obviously false. For example, it would not be outlandish to describe the following as invalid: "Aborigines don't deserve any compensation. They weren't able to protect their land; tough luck." Such an argument requires a denial of the legitimacy of the rule of law to get anywhere close to validity.

Another point of this exercise is that the process of reformulating an argument as a valid deduction forces you to think explicitly about its presuppositions. This is the main reason

[3] To make the argument *formally* valid, one would introduce this principle of action as an additional axiom schema.

such an exercise is worth doing. It clearly is a lot of trouble to fill in the presuppositions of an argument, but when applied to a key argument for some research work the effort of a detailed exposure of its presuppositions is well worth while. Note that a valid argument can fail to be sound in only one way: its premises can fail to be true. Therefore, by explicitly formulating all the premises required for an argument to become valid, one will automatically be making any weaknesses in the argument explicit — as weak, explicit premises, rather than hidden inferences. This, of course, is useful if your goal is to criticize someone else's argument; but it is equally useful in criticizing your own arguments, which is an essential step in producing good arguments in your own research writing. Just as it is difficult to write computer programs longer than 20 lines that are bug free, it is difficult to write arguments that are bug free. Self-criticism is useful in overcoming defects in arguments. The analogy with programming can be continued: it is very difficult to find bugs in your own code, because that code in fact reflects your way of thinking. *Looking over* your own code is unlikely to induce a new way of thinking about it and so is unlikely to reveal bugs that are right in front of your face. So it is with argumentation. The process of making your full argument *explicit* can serve to overcome this difficulty in part, for rendering argumentative steps explicit is not likely to be your (or anyone's) natural way of thinking.[4] Finally, as with bug catching, showing your ideas to other people for their criticism is *very strongly recommended*. Most academics are sensible enough to do this prior to submitting papers for publication; students should be equally sensible and swap critiques of each others' papers prior to turning them in for marking.

4.2 Inductive Arguments

Although any argument can be turned into a deductively valid argument, in many cases it is most natural to formulate an argument as an inductively "sound" argument. Deductive soundness requires the extremely strong condition that it is strictly impossible for the conclusion to be false, because of the validity of the argument and the truth of the premises. Inductive soundness is a looser notion, usually reported as the high probability of the conclusion, which depends upon the truth of the premises and a strong inductive relationship between them and the conclusion. Sherlock Holmes, in one of the many stories of Conan Doyle, concludes of a train passenger sitting opposite that he is a teacher, because he has chalk on his hands. The form of the (partially implicit) argument is presumably:

(1') This man has white powder on his hands.

(2') White powder which looks like chalk is chalk.

(3) This man has chalk on his hands.

(4') Most people with chalk on their hands are teachers.

(5) This man is a teacher.

The first two premises make explicit that an induction has already occurred in Holmes's "observation" that the man has chalk on his hands. Premise (4') is the main suppressed premise; it is the most plausible way of interpreting the argument. To turn (3) $\rightarrow$ (5) into a strictly valid argument would require the adoption of a manifestly false premise: that all

[4]Incidentally, it is of course not my intention to be seen as advocating the rendering of all your arguments explicit in every detail in your *paper*. Such a style of writing would be pedantic and actually more difficult to read than leaving straightforward steps implicit. The goal of making your arguments explicit is to understand them better and take notice of their flaws so as to improve your arguments, not to make them more cumbersome reading.

people with chalk on their hands are teachers.[5] In this, as in other, cases there is no need to attribute to Holmes the strong conviction that his conclusions are *certain* and so there is every reason to understand Holmes's reasoning as inductive rather than deductive — despite Conan Doyle's repeated confusions on this matter. (He frequently had Holmes talking about his "deductive powers" of reasoning.)

The analysis of inductive arguments begins in the same way as that of deductive arguments: identifying the basic argument structure and filling in missing premises. Criticism of an inductive argument, however, occurs in different terms from criticism of a deductive argument. It is enough to reject a deductive argument as valid that a single counterexample be found where the premises are satisfied but the conclusion is not. Clearly, counterexampling will not work against inductive arguments, for the conclusion is specifically intended to "go beyond" the premises. To criticize the inductive strength of an argument you must instead show that the truth of the premises would fail to make the conclusion highly probable.

There are various kinds of inductive argument. One type involves a *direct inference* from a proportion of individuals in a class having some property to the probability of a particular individual having that property. If the individual has been selected from the class without any (known) bias, then the probability should be equal to the proportion within the class. Thus, if 95% of Swedes are Lutheran, then being told that a Mr. Peterson is a Swede, it would be sensible to conclude that Mr. Peterson is a Lutheran. Such an inference would be defeated by knowing some additional biasing factor, say that Mr. Peterson had just left a Catholic mass.

A second form of inductive argument proceeds in the reverse direction from a sample of individuals to characteristics of the population from which they are drawn, in an *inverse inference*. Thus, political survey samples typically inquire of a small group of voters about their voting preferences; these preferences, within a range of probable error, are then ascribed to the total voting population. Again, such inferences can be defeated if the sampling procedure can be shown to be biased. In an infamous incident, the *Literary Digest* predicted that U.S. presidential candidate Alfred Landon would defeat Franklin D. Roosevelt. This was based on a sample of voters who were found in telephone directories and automobile registration lists. What was the bias? (Hint: this occurred during the Great Depression.)

Analogies provide the basis for many inductive arguments. By drawing attention to similarities between structural or lawlike features of two systems, some support may be found for claiming that a further, unobserved feature of one is likely to be similar to a corresponding feature of the other. Thus, some argue that artificial neural networks are more likely to succeed in AI than rule-based systems because of the analogy between them and the only physical systems we know of that support cognition, namely brains. An analogical argument may be undercut by pointing out *disanalogies* between the two systems. Disanalogies may suggest that the property in question is more likely to be within the class of properties that differ between the two systems. Artificial neural nets and biological neural nets are alike in being built out of heavily interconnected but individually simple units, but are unalike in so many other ways (complexity, hormonal processing, developmental processes, metabolism, etc.) that the analogical argument is not really taken seriously, even by most proponents of artificial neural nets.

Finally, in science, most inductive arguments of the above forms find their place within larger inductive arguments for or against various causal theories about how some physical process works. Some specific kinds of inductive reasoning which may support or undermine causal theories invoke the existence of (or failure of) correlations between observed variables and again the existence of (or failure of) correlations between a variable subject to experimental intervention and another variable. Initially, the presence or absence of the correlations

[5] Instead of strengthening the premise we could obtain validity by weakening the conclusion, but to investigate such refinements you will have to read more extensive treatments of inductive reasoning than that in this paper.

must be the subject of an inverse inference from sample to population (the second form of argument above). The further inference, from the correlation (or absence of correlation) to a causal structure that would explain that correlation, is again subject to various possible complaints; however, the structure of such arguments is beyond the scope of this paper (see Giere, 1991, or Korb, 1997).

In general, inductive arguments are attacked by showing that the premises are not representative of the facts put forward in the conclusion, i.e., that some bias is operating in the argument. This implies that inductive arguments, unlike deductive arguments, are sensitive to background information. In deductive argument, if the premises *ever* render the conclusion necessary, no additional knowledge can undermine that valid relationship. However, in inductive argument one may come to learn of a biasing factor which undermines an argument that otherwise is perfectly fine. For example, in Holmes's case it may be that in late 19th-century England the majority of people with chalk on their hands were school teachers. However, if we came to know that the train was carrying a convention of chemists, who mix many medications containing white powder, Holmes's argument would be severely undermined, for premise (2'), that what appears to be chalk is chalk, while true in general, would no longer be accurate in the case at hand.

There are many other special considerations which apply to assessing the strength of an inductive argument — that is, how probable the conclusion is, given the truth of the premises — especially when the premises involve statistical samples or outcomes of scientific experiments. I discuss some of these issues in "Common Fallacies and Errors" below. For more on such matters see Giere (1991) or Baird (1992).

4.3 The Credibility of Your Evidence

Since the object of any argument is to persuade rational persons of some given level of expertise (depending upon your target audience) to believe your conclusions, validity (or inductive strength) is not enough. The reader must also believe your premises. Such belief may be achieved by recursively arguing on their behalf, but at some point the recursion must end and you must simply present premises which will be believed independently of argument.[6]

How can that be achieved? The primary means is by offering authoritative evidence, evidence that comes from sources that are known or expected to be reliable. Ultimately, evidence derives from observation or experimentation. If you are reporting on your own observations or experiments, then you need to report enough information so that others can reproduce the experiments or observations. This means that in addition to the data themselves, or summary statistics and tables, your report needs to describe the methods employed in collecting them. Such reportage will be used by readers to assess whether your experimental and observational technique incorporated any hidden biases, for example — which goes directly to the reliability of your evidence. If the nature of the report does not allow the inclusion of all such details, they should be made available in an appendix or in a technical report available separately. In addition, the raw data produced experimentally (e.g., by a computer program) should be retained, and made available to inquirers, if any significant claims are being based upon them.[7]

Of course, you may also (or instead) be reporting evidence from secondary sources.[8] Readers will assess the reliability of the data and your use of them according to such factors as the research reputations of their authors and their institutions, the research reputation of the venue in which the data were published (e.g., *Science* will be given more credence than *Scien-*

[6]Lewis Carroll (1895) makes a similar point, that justification of an inference can be never-ending.

[7]Indeed, in some circumstances this is a legal requirement.

[8]Primary sources are those which provide first-person accounts of events, such as laboratory notes. Secondary sources are derivative compilations which may be based upon such evidence, such as textbooks.

tific American), the audience to whom the publication was originally directed (e.g., popular books will be given less credence than conference proceedings), and the adequacy with which you appear to understand and interpret the data. The reputations of different publications are zealously defended in academia; they are of no small account. Once lost, reputations — individual or institutional — are rarely recovered. They are built up both by the quality of the research produced and also (not coincidentally) by the thoroughness of the professional refereeing process through which the research is put.

Secondary sources need to be cited properly. This means, at least, that all the information necessary for a reader to obtain a copy of the source needs to be made available, including page numbers wherever relevant and known. See my references at the back for examples.

Questions worth asking about proffered evidence include:

- Is it a primary or secondary source?
- Does the source have a particular interest in the conclusion drawn (monetary, personal, political, etc.)?
- What is the reputation and/or reliability of the source of evidence?
- How detailed is the description of the evidence? How precise is the evidence? Is the claimed precision justified?
- How believable is the evidence relative to background information?
- Is the report of the evidence backed up by detailed records? What amount of time passed before the events described by the evidence occurred and records were made?
- If the evidence is quoted or paraphrased text, does it take the context into account? (This is a question worth asking, even if you cannot answer it directly. The popular press frequently abuses people by quoting and paraphrasing out of context.)
- Is the evidence based upon the impressions of human witnesses? Are the humans expert at reporting such evidence?

4.4 Why Sound Argument?

Perhaps some will have been wondering why anyone should go to all of the trouble implicit in the discussion so far. If the goal is simply to persuade readers to believe one's conclusions, there are much easier methods of doing so than the heavily analytic ones described above. Advertisers, politicians and lawyers use them all the time: illegitimate appeals to authority, rabble-rousing slogans, negative insinuations about one's opponents, etc.

There is no denying that such techniques work. Furthermore, there is no denying that they are used by academics, at least on occasion (and perhaps especially when dealing with politicians or other "outsiders"). But with scientific writing the intentional use of psychological and rhetorical tricks is likely to be less effective than elsewhere; indeed it is more likely to spoil your reputation than to serve it. Scientific knowledge is grown by a process of critical review and scientists are very explicit about this process. Cheap tricks exposed are often the end of a scientific career. Where the process of scholarly criticism and review is the primary one by which new ideas are accepted into the collective body of belief, the use of evidence and argument that can withstand such review is far more persuasive than any other approach.

4.5 Argument Evaluation

The quality of your final assessment of an argument depends upon the quality of all the steps taken leading up to that assessment. What I have described above is an abbreviated version

of a method of analysis developed by Michael Scriven (1980). His complete list of analytical steps is:

1. Clarification of meanings, including the use of dictionaries and encyclopedias where appropriate. Removing ambiguities of terms. Identifying and labeling propositions (not sentences: there may be multiple propositions per sentence; other sentences may have no [relevant] content).

2. Identification of conclusions, stated and unstated. Just as premises may be left implicit, so too may conclusions be left implicit, even the main point of an argument. Included here are intermediate conclusions of subarguments.

3. Portrayal of the argument structure in graphical form. This will help you find out what is missing to make the argument valid.

4. Formulation of unstated premises. This is the most difficult step: it requires creativity and relevant background knowledge to find the premises most appropriate to produce a valid argument. This is also the step most prone to bias and misrepresentation if you are emotionally responsive to the argument.

5. Criticism of the inferences and premises. Try to find counterexamples — cases where the premises are true, but the conclusion false. For inductive inferences try to determine whether the evidence is biased or unrepresentative. Check whether the samples are too small to offer reliable support for the conclusion. Premises may be criticized by constructing new arguments against them which attack their credibility.

6. Introduction of other relevant arguments. This involves all the other steps combined in order to generate entirely new arguments in opposition to your conclusion — exploring reasons and ways in which the thesis in question could go wrong. Also, you should examine arguments on the subject that others have published.

7. Overall evaluation. How good is your argument compared to the others? How strong are your premises and their relation to the conclusion? How much of your bank account would you bet on the conclusion? At what odds?

4.6 Bibliographic Notes

An interesting guide to the visual presentation of data is Tufte (1990).

The process of recovering missing premises is called *abduction*, a concept introduced by C.S. Peirce (1940). For an elaboration of a computer architecture for implementing argument analysis and generation see Zukerman, Korb and McConachy (1996).

There are a great many books on the construction and analysis of arguments. Most of them overemphasize formal logic and underappreciate inductive forms of reasoning. One which does not that I especially recommend is Scriven (1980) — a lucid and practical guide. For an introduction to the basic concepts of logic see Salmon (1973). Three good books which discuss the construction of inductive arguments, including those based upon scientific experiments, are Giere (1991), Baird (1992) and Skyrms (1986). An alternative approach to analyzing arguments that is quite popular, giving rise to many textbooks, is due to Toulmin; see Toulmin et al. (1979) and Toulmin (1958). The Toulmin-line of textbooks share most of the weaknesses of the majority, specifically including an underappreciation of the role of induction (it has some additional weaknesses as well, discussed in Freeman, 1991).

The critical reasoning literature represents an educational movement to foster critical and analytical approaches to argumentation. It is represented, for example, in Ennis (1962) and Paul (1984) and in the journal *Informal Logic*. Whereas the emphasis this tradition puts

on the skeptical review of arguments is laudable, it shows an unhealthy tendency to dismiss the role of creative thinking. Indeed, according to the reformist Siegel (1988), the "Informal Logic Movement" (as he calls it) simply repeats the error of formal logicians in dismissing any role for anything other than logic in argument analysis, albeit in this case they are talking of *informal* logic: "[the Informal Logic Movement] *identifies* (informal) logic with critical thinking" (p. 2; my italics). I believe Scriven's approach by contrast highlights the key role of creative thinking *within* critical thinking, for example in the abductive generation of missing premises and in building alternative arguments.[9]

5 Common Fallacies and Errors

5.1 Bayesian Reasoning

In order to have a framework for assessing the degree to which some evidence supports a conclusion I shall now introduce some basic principles of Bayesian evaluation theory. I will not defend them here, as their proper defence is a complicated story. If you are curious about that story, you should look at Korb (1995), Howson and Urbach (1993) or the marvelous paper, launching modern Bayesianism, by Frank Ramsey (1931).

Bayes's theorem, a discovery of Thomas Bayes (1763), reports the relation between the probability of a hypothesis (conclusion) given its evidence — $P(h|e)$ — and its probability prior to any evidence $P(h)$ and its likelihood $P(e|h)$, the probability that the hypothesis implies for the given evidence. In particular,

$$P(h|e) = \frac{P(e|h) \times P(h)}{P(e)}$$

The theorem is just that and is not controversial. Bayesian evaluation theory is controversial and asserts that the proper way to evaluate hypotheses in science just is to apply Bayes's theorem (the process is called "Bayesian conditionalization"); that is, we should judge hypotheses according to the relation above between likelihood and prior probability when the evidence is taken to be the total relevant set of observations and experimental outcomes available.

It is an implication of Bayesian conditionalization that evidence will support a hypothesis if and only if the *likelihood ratio* — $\lambda(e|h) = P(e|h)/P(e|\neg h)$[10] — is greater than one, and indeed the degree of support just is the degree to which the likelihood ratio exceeds one (see Korb, 1994). Thus, for example, the claim that a coin has a probability $p = 0.8$ of heads, where the only alternative hypothesis being entertained is that the coin is fair, will be supported more by 18 heads out of a random sequence of 20 coin flips than by 15 heads because the likelihoods given by

$$\left(\begin{array}{c} 20 \\ k \end{array} \right) p^k (1-p)^{20-k}$$

lead to the likelihood ratio in the case of 18 heads:

$$\lambda(18|p = .8) = \frac{P(18|p = .8)}{P(18|p = .5)} = 760.55$$

and for 15 heads:

$$\lambda(15|p = .8) = \frac{P(15|p = .8)}{P(15|p = .5)} = 11.80$$

[9] Note that Siegel himself disputes the adequacy of the Informal Logic Movement's view of critical reasoning and argues that creativity and criticism are interdependent (note 21, p. 152).

[10] $\neg h$ being the negation of h.

Note that in both cases the hypothesis of a biased coin is supported, but the support is far greater in the first case.

The concept of the likelihood ratio provides a simple but effective tool for analyzing the impact of evidence on conclusions (hypotheses). For example, it makes clear why Karl Popper's (1959) insistence that scientific hypotheses be subjected to severe tests makes sense. Intuitively, a severe test is one in which the hypothesis, if false, is unlikely to survive; that is, whereas the hypothesis predicts some outcome e, its competitors do not. Since the hypothesis predicts e, $P(e|h)$ must be high; since its competitors do not, $P(e|\neg h)$ must be low. These jointly imply that the likelihood ratio is very high. Therefore, a severe test will be highly confirmatory if passed and highly disconfirmatory if failed — and so provides the most efficient approach to testing a hypothesis.

5.1.1 Bibliographic Notes

Howson and Urbach (1993) and Korb (1992) both survey the recent successes of Bayesian evaluation theory in accounting for various aspects of scientific method. Other interesting accounts of Bayesian inference may be found in Polya (1968, chapter XV) and I.J. Good (1985).

5.2 Logical Fallacies

Logical fallacies are prime targets for logicians seeking stories of error and shame with which to discomfit their students. Of course, the fallacies they identify do describe forms of argument that are often suspect — the fallacies are hardly *valid* argument forms. In certain cases, however, Bayesian evaluation theory suggests that there may be more merit to arguments employing them than the logicians customarily allow. This is probably due to the extreme attention most logicians and their textbooks pay to valid deductive argument, which in turn is probably due to the historical belief that correct inference could be understood as a matter of *syntactical form* alone of the sentences involved.[11] As a result, in logic textbooks inductive varieties of argument are generally given a brief acknowledgement and passed over, or left to the unread end of the text. Such treatment is a serious disservice, since inductive argumentation is a key to the majority of practical arguments deployed in science, industry, politics and ethics. This likely contributes to the nearly universal failure to find skills learned in logic classes having any payoff outside of those classes (e.g., Pollard and Evans 1980). Here I review a few fallacious forms of argument, noting both merits and demerits.

Equivocation

Scriven's first step in the analysis of argument is the clarification of meanings. At this stage it will be helpful to look for equivocal uses of terms, that is to say, for inferences which depend for their *appearance* of validity upon using the same term in two different ways. One must be sensitive to the ways in which one context activates one meaning or connotation of a term or phrase and another activates an importantly different meaning. Consider,

1. Killing is immoral; every religion and system of law acknowledges this.

2. Every system of law and religion treats complicity in a killing as immoral also.

3. Eating meat is complicity in a killing.

[11] Note, therefore, that my criticism here of the teaching of logic is restricted to teaching logic as a means of understanding or aiding human argumentation and inference. None of my criticism is relevant to the teaching of logic in formal mathematics, or in support of metamathematics, or in support of the construction of philosophical logics for various special purposes.

4. Therefore, eating meat is immoral.

Of course, the kind of killing banned by religion and law is the killing of humans, whereas that of proposition 3 is not. This argument on behalf of vegetarianism is no more legitimate than would be its equivocal extension into an argument in favor of starvation (i.e., against also the eating of vegetables).

Affirming the Consequent

Perhaps the most basic form of valid deductive inference is *modus ponens*:

1. If A, then B.
2. A.
3. Therefore, B.

This is certainly intuitively a correct form of argument, at least when A and B are replaced by instances (e.g., let A = "I go to the store" and B = "I will spend too much money"). It depends for its validity upon a perfectly ordinary understanding of the conditional "if ... then ...". Other uses of the conditional are less clear intuitively, such as *modus tollens*:

1. If A, then B.
2. not-B.
3. Therefore, not-A.

This runs the conditional in reverse, by denying its consequent, and is just as valid as *modus ponens*. In many cases it is also just as intuitive. For example,

1. If you wish to compete, you must have completed a negative drug test.
2. You did not complete a test.
3. Therefore, you cannot compete.

In less usual settings *modus tollens* may be harder to handle. Consider the claim: Every card with a vowel on one side has an even number on the other (if vowel, then even). If you are shown four cards displaying C, E, 8, 3, which cards need to be turned over to check whether the claim is true? Most people have no difficulty seeing that the E card needs to be turned over (*modus ponens*), but few notice that the 3 card also needs to be turned over (*modus tollens*) (see Wason, 1966, for an exploration of such difficulties in reasoning).

Affirming the consequent attempts to use the conditional in reverse, but *without* denying the consequent. Here is a compellingly atrocious case, by Joseph McCarthy, after whom "McCarthyism" was named (McCarthy, 1952):

> I do not tell you that Schlesinger, Stevenson's number one man [Stevenson was the Democratic presidential candidate at the time], number one brain trust, I don't tell you he's a Communist. I have no information on that point. But I do know that if he were a Communist he would also ridicule religion as Schlesinger has done.

Despite the introductory denials, of course, it was precisely McCarthy's intention to encourage his hearers to infer that Schlesinger was a Communist.

Affirming the consequent is perhaps the most blatantly illogical form of argument in the logicians' collection of fallacies. Nevertheless, there are circumstances where affirming the

consequent makes sense. If there are very few possible explanations for a newly discovered fact e, and h is one, then this discovery surely raises the probability that h is true. In Schlesinger's case, there are any number of explanations possible of why he may have ridiculed religion, so the implied connection with Communism was pathetically poorly supported. On the other hand, if the recently claimed discovery of life on Mars is confirmed, then surely the probability that life originated independently on Mars is greatly increased, for such an independent origin of life is one of only a very few plausible explanations of its discovery there (another being that life forms may have been carried from Earth to Mars, for example).

If the newly discovered fact e is surprising, then a hypothesis h which is a possible explanation of it is confirmed by the discovery. This follows, again, from considering the Bayesian likelihood ratio: e being originally unlikely (surprising) means $P(e)$ is low; since h is a possible explanation $P(e|h)$ is high; the combination of these conditions implies that $\lambda(e|h)$ is high and h is confirmed. A further condition that there are few competitors to h in explaining e would imply also that the posterior probability of h is not low (since the total posterior probability must be distributed among only a few candidate hypotheses).[12]

More generally, affirming the consequent is the basic logical form of confirmatory arguments in science — described by many philosophers of science as *hypothetico-deductive* confirmation, because the confirmatory evidence e is (ideally) deduced from the hypothesis (plus background knowledge). Of course, from the evidence we cannot *deduce* the truth of the hypothesis: that there is no such deduction we can determine from the logical form of the propositions involved. But whether or not the argument offers inductive support is not a matter of logical form. The argument that Schlesinger is a communist and that life originated on Mars have the same *form*; what they do not share is a semantic property: the likelihoods. Any case of affirming the consequent that satisfies the Bayesian requirements for a severe test will provide noticeable confirmation for a hypothesis.

Red Herrings and Straw Men

Red herrings and straw men are two of the strange beasts that logicians frequently bag. There is nothing much to say in their defence; they simply serve to avoid real argument. *Red herrings* are arguments or facts that are irrelevant to the main issue under discussion but are brought in in order to distract people from that issue. They are particularly effective when they are highly emotive. Thus, if McCarthy had been challenged on his fallacious appeal against Schlesinger, perhaps he would have claimed that Communists are taking over in Washington and that it was imperative that we do everything in our power to counter their influence. Whether true or false, though, it would have been an utter irrelevance to the argument at issue — a red herring, in other words.

Straw men are what you invent when you violate the principle of charity: the attribution to your opponents of arguments which they in fact do not endorse. This is usually done in order to have a target of attack that is easier to demolish than the real argument. But if you cannot offer serious argument to the views actually held by your opponents, perhaps you should reconsider your opposition to those views.

Appeals to Authority

The only way in which research can progress collectively is by a division of labor, some people must become expert in some areas and others in different areas. Inevitably, then, we shall have to appeal to the expertise of others either in validating some evidential claim or in supplying an expert opinion as a premise for our arguments. The extent to which such appeals support

[12] Affirming the consequent is also one way (among many) of implementing Peircean *abduction*. Abduction, however, is not normally taken to offer *support* to a hypothesis, so this need not concern us here.

our arguments depends upon the extent to which those cited actually have relevant expertise and are properly considered reliable sources.

Illegitimate or improper appeals to authority occur when: the authority or expertise of the people invoked pertains to some other domain (as when sports stars publicly endorse headache remedies), or the expert opinion is widely disputed by other acknowledged experts in the domain, or the expert or authority has a history of offering unreliable opinions. The reliability and integrity of an authority can be judged by knowing something of the research history of the person and the person's institution. Of course, if the expert has a monetary interest (or other direct interest) in one side of the dispute, as various researchers supporting the "safety" of the use of tobacco do for example, then neither the integrity nor the reliability of the source can be considered unimpeachable.

The recently proposed theory that HIV/AIDS passed to humans through polio vaccine trials in the 1950s has some significant circumstantial evidence on its side. For example, the vaccine trials occurred in Zaire immediately prior to the first known human infections, which occurred in Zaire. The trials used vaccines grown in monkey kidneys and it is widely believed among experts that HIV jumped species from African monkey to Homo sapiens. The vaccine trials involved tens of thousands of monkeys and humans, thus providing a plausible opportunity for a new infection to jump the species barrier. Nevertheless, scientific journals refused to publish this theory. (It has only come to the public's attention through a book, Cribb, 1996). Various epidemiologists and AIDS researchers have spoken out against it. Whatever the merits of their views, or the case generally, it is a relevant consideration that the medical establishment has a direct monetary interest in sustaining cross-species medical work, including the use of non-humans in medical transplants, and confirmation of the HIV theory could put this into jeopardy.

Ad Hominem

Ad hominem attacks are those directed against the person rather than the argument. It is something like shooting the messenger when you have received bad news. This is a very popular device in argument because it is often easy to apply and can be greatly distracting from the real issues. It is, in fact, a common and effective form of political campaign "debate" in Western democracies: political campaign organizations spend considerable energy attempting to find juicy stories about their opponents and relatively little energy either developing their own policies or finding problems with those of their opponents. Although character may be a legitimate issue in a campaign, the private sexual life of a politician is probably not. Note that people do not complain of John Kennedy that he was a bad president because of his adultery.

Regardless of the dangers of *ad hominem* attacks, in many cases they are legitimate because they undermine the authority or reliability of an opponent relying upon his or her own expertise to support parts of an argument. Thus the fact that Richard Nixon was exposed as a liar who, as president, arranged for forged letters to discredit political opponents and covered up various felonies by subordinates, surely undermines Nixon's credibility in general, and therefore also in particular, such as raising doubts about evidence against Alger Hiss developed by Nixon during his red-baiting days as a U.S. Congressman.

What you need to distinguish in assessing an *ad hominem* argument is whether the person under attack has simply put forward an idea or argument for your consideration or whether he or she is *vouching* for some evidential statement used by the argument. In the first case an *ad hominem* argument is inappropriate: the idea or argument can, and should, be assessed independently of its origin. In the latter case whatever argument legitimately undermines the authority of the person also undermines the evidence put forward.

Evidence and inference: Legitimate appeals to authority and their opposite, legitimate

ad hominem attacks, pertain to the reliability of the evidence. That is to say, for some evidence e explained by a hypothesis h they raise or lower $P(e)$. In consequence, they alter the posterior probability of h, supporting it and undermining it, respectively.[13] As already argued, affirming the consequent is also analyzable in Bayesian terms: whether such an argument is successful or not depends directly upon the relevant likelihood ratio. In general, assessing the merit of arguments requires an understanding of the relationship between hypothesis and evidence, the scope and viability of competing hypotheses, and the relevant background to the argument. This will be true regardless of whether these factors can, or cannot, be explicitly quantified.

Stereotyped Reasoning and Prejudice

Everyone reasons in part by stereotypes. We stereotype almost everything. We label other people on the basis of superficial encounters ("He's boring" and "She's cool"), books on the basis of their covers and movies on the basis of their advertisements. This is not such a bad thing when nothing much hangs on our superficial judgments. It is also not a bad thing when much does hang on them, if we are constrained by time or other pressures to skip a careful inquiry and judgment. Indeed, evolutionary pressures to survive and reproduce would strongly favor the use of stereotypes in such situations, for they often are reasonably informative for decision making, even if not optimal.

Stereotypes are one relevant and appropriate source of prior probabilities, given a readiness to acquire new information and its appropriate use in adjusting one's probabilities. Inappropriate uses of stereotypes, on the other hand, involve the failure to adjust them in the light of new information. Such behavior is what is generally, and rightly, decried as *prejudice* (and not merely the use of stereotypes to literally pre-judge people or things). This prejudicial bias can take two forms. First, some people refuse to acquire new information when it is readily available and appropriate to the problem. A refusal to test a woman for the physical capacities required for fire-fighting would be an example of this kind of prejudice. Second, many people refuse to take account of information overriding a stereotype even when it is given to them. Most people either know or ought to know, for example, that the intellectual differences between races of Homo sapiens, if they exist, are far smaller than differences between individuals of the same race. So, for the purpose of judging individual humans for some role where intelligence is an issue, race is irrelevant. This knowledge does not prevent many people from bringing race up in the context of judging individual intellectual performance, and that is an example of prejudice. (See also the discussion of "underadjustment" below.)

5.2.1 Bibliographic Notes

I know of no other Bayesian analyses of the logical fallacies.[14] Some orthodox treatments of logical fallacies are Copi (1972), Mackie (1967) and Walton (1989). The psychology of deduction as a research area began with Peter Wason; see Wason (1960, 1966). Various papers in this and related areas are collected in Johnson-Laird and Wason (1977). A good overview, relating the psychology of deduction to the psychology of induction, is Evans (1989).

[13]The explanation of this impact lies in a generalization of Bayes's rule for conditionalization, called Jeffrey conditionalization (Jeffrey, 1983). Where $P'(e)$ is the probability of e after some new evidence or argument has been considered, when the new evidence or argument does not make e itself certain, its impact can be assessed via:

$$P'(h|e) = P(h|e)P'(e) + P(h|\neg e)P'(\neg e)$$

[14]However, Koertge (1993) has discussed judging books by their covers from a Bayesian viewpoint.

5.3 Inductive Error

Since Phillips et al. (1966), cognitive psychologists have been demonstrating that inductive cognitive error is widespread even among people trained in statistics.[15] This psychological work has yet to find its way into the literature of informal logic and argument analysis. However, it is highly plausible that by identifying the more common errors to which people are prone and the conditions that give rise to them, we can hope to stop them from creeping into our own thinking and so improve the quality of our written argumentation. The best way to improve our inductive thinking in response to these errors appears to be a combination of the study of the relevant probabilistic rules and their use in working through realistic cases in which the errors arise (see Nisbett et al., 1987; Johnson-Laird, 1983). I will do neither of those here, but will describe the errors and their circumstances in sufficient detail to allow the interested reader to pursue the matter.

The Base Rate Fallacy

This is the primary inductive fallacy: people's tendency to ignore prior probabilities, the "base rate fallacy" (where "base rate" means prior probability in the jargon of the psychologists). This tendency manifests itself in a variety of circumstances. For example, if people have a particular image of what a stereotypical member of a class is like, then something matching that image will very likely be considered to be in the class — regardless of any explicit prior probability of its being in the class. In a notorious example of this Tversky and Kahneman (1982) found that most people consider a woman described as a politically active feminist more likely to be a feminist bank teller than to be a bank teller. This is despite the fact that such an assertion is strictly incoherent. The reason for this incoherence, according to Tversky and Kahneman (and they offer some experimental support for this), is that people's stereotype of bank tellers is incompatible with their stereotype of feminists, so anyone described as a feminist is unlikely to be associated with the class of bank tellers without further qualification.[16] According to Tversky and Kahneman, humans use a variety of heuristics to solve problems more quickly and simply. Amongst these is the *Representativeness Heuristic*, which says that to judge whether individual A is a member of set X assess the similarity of A to a stereotypical member of X. Nisbett et al. (1976) propose a different explanation of the phenomenon. They claim (and again offer experimental support) that people's tendency to ignore base rates is a consequence of base-rate information normally being presented in an abstract and pallid way. When base-rate information is presented in a vivid way, using concrete examples to communicate the information, then people *do* take the base rate into account in their problem solving. Cognitive psychologists have dubbed the tendency to employ information which is vivid, salient or recently presented the *Availability Heuristic*.

Regardless of what the right story is, we can observe: these heuristics are likely to be very useful in situations where rapid decision making is necessary, as must have occurred frequently in human evolution; the writing of research papers is not (or, ought not to be) one of those times. Although we may all have tendencies to suppress the proper treatment of prior information, when we have the time and other resources to consider prior information and its relevance we should do so.

[15] Interestingly, this chain of research was set off by the publication of the seminal Bayesian work of Savage (1954), according to Edwards and von Winterfeldt (1986).

[16] Of course, it is also true that in the experimental setup Tversky and Kahneman did not require their subjects to put the two incompatible probability assessments side-by-side, which would have made the incoherence obvious.

The Law of Small Numbers

The Law of Large Numbers says that as sample sizes increase without bound the sample mean will converge upon the true mean to an arbitrarily close degree. The "Law of Small Numbers" was introduced by Tversky and Kahneman (1971) to describe human cognitive errors in dealing with small samples and asserts that estimates based upon small samples are just as reliable as estimates based upon large samples. This is again ascribed to the Representativeness Heuristic: that samples that are representative of the population are likely to be found, independently of the sample size. This is, of course, wrong.

Suppose, for example, that the probability of a live birth being male is 0.51. Then from the binomial function we can determine that the probability that a large hospital with 1000 births per year will have more male than female births in a given year is about 0.74. What is the probability that on an average day the number of males born is greater than the number of females? If you are inclined to say "About 75%," then you are inclined towards believing in the Law of Small Numbers. The average number of births per day will be around 3. Examination of the binomial function reveals that the probability that 2 or 3 of three random births will be male is 0.51.[17] The point is that large samples are far more likely to be representative of the underlying population (the population of births, having more males than females) than small samples. The failure to note this also leads experimental scientists to be overconfident that repeating a successful significance test will again produce a significant result (Tversky and Kahneman, 1971).

The "gambler's fallacy" is a related phenomenon: gamblers often expect a short run of luck (good or bad) to be immediately rectified — if you have just had a good run, quit now before it turns (or the opposite excuse for mortgaging your house) — because the small sample of recent outcomes and the near future must together resemble the long run outcome of a large sample. But it just ain't so. Small samples are simply less likely to be representative.

Overgeneralization and Bias

Another common error is to overgeneralize on the basis of a small or biased sample. For example, job interviews are usually given far more weight in deciding whom to hire than is rational. How a person performs under the stress of an interview is not all that closely related to how the person will perform in day-to-day job situations. Furthermore, job interviews offer at best an *extremely* small sample of a person's total behavior during a career. First impressions may be important, but they *should* be a good deal less important than they are. In a study of academics, Nisbett and Ross (1980) report that evaluations of interviews and presentations were utterly irrelevant to predicting future academic research performance; simple linear functions of grades and prior publication statistics, while imperfect, were at least not orthogonal to future performance.

In general, when presented with evidence offered as support for a generalization, you should ask whether the evidence has really been collected from the full range of the generalization. For example, Benjamin Whorf, an erstwhile influential writer on linguistic anthropology, claimed that the Hopi (Native Americans) have nothing like the "Western" concept of time, that they lack the notion of a flow of time from past into future. His stated reason for this is that their language has "no words, grammatical forms, constructions, or expressions that refer directly to what we call 'time'" (Whorf, 1956). This may be (if you put the stress on *directly*), but the Hopi language does have tense, units of time, metaphors for time, etc. (Malotki, 1983)

[17] The formula is:

$$\binom{3}{2} .51^2 \times .49 + \binom{3}{3} .51^3 = 0.51$$

and the observance of various rituals at specific times is very strict in Hopi culture.

The reverse side of overgeneralization is that Bayesian evaluation theory supports the gathering of diverse evidence in testing a hypothesis, as opposed to the repeated collection of similar evidence. This follows immediately from Bayes's theorem. Suppose there are three items of evidence that would support h, say d, e, and f, and suppose that the *a priori* likelihoods are equal: $P(d|h) = P(e|h) = P(f|h)$. Suppose also that d and e are more closely related than d and f. We can represent this probabilistically by: $P(e|d) > P(f|d)$. If d is learned to be true, then other things being equal (e.g., the cost of experimentation) which of e and f should we investigate next — i.e., which of them, as additional evidence, would have greater confirmatory power? The answer is trivially the more disparate evidence f. Bayes's theorem requires, in the use of e versus f, that the numerators be identical (given the assumptions above) but the denominator will be $P(e|d)$ in the first case and $P(f|d)$ in the second, whence it follows that $P(h|e,d)$ is less than $P(h|f,d)$. (See Franklin, 1986, for a different proof.) For an example, if we are testing the claim that one algorithm is better than another at letter recognition, a second positive test against a second sample from the US postal database would provide less support than a similarly positive test using a European database.

Anchoring

People tend to anchor their judgments of probability upon some salient initial value and then underadjust in the light of new evidence (the underadjustment is called "conservatism"; see Phillips et al., 1966). This occurs, for example, in making judgments about compound events (Bar Hillel, 1973). To take a particular example, this appears to occur in the scheduling of software projects. Schedules are typically produced by managers requesting their subordinates to produce schedules for their small piece of the whole, estimated to achieve some high confidence level — say, having a 90% chance of being met. Once 90% gets into the picture it hardly ever departs. Thus, if ten engineers produce ten 90% schedules and the schedule of the whole is dependent upon all of its parts, what is the probability that the composite schedule will be met? If we assume independence of schedule slips (and if we do not, the probability will be even smaller), that probability is $(0.9)^{10} = 0.35$. Most managers will report the composite as a "90% schedule." The great majority of software schedules are not met.

The Fundamental Attribution Error

The fundamental attribution error in social psychology is the common mistake of attributing responsibility for an individual's behavior to the individual (fair enough, so far) *without reference to the individual's environment* (which likely leads to mistakes). Thus, many non-Germans claim that German complicity in the Nazi crimes against humanity was due to some intrinsic flaw in German character. This is a claim which is oddly resonant with the German ideology that backed those crimes. It ignores the social environment surrounding Germans at the time and it supports the easy dismissal of general human tendencies to commit or tolerate such crimes under circumstances where individual responsibility is likely to be evaded (as in Cambodia, ex-Yugoslavia, the Soviet Union, Lebanon, etc.).

How much responsibility to attribute to the environment and how much to attribute to genetics — i.e., the nature-nurture controversy — is a related and ever-difficult issue. It is made more complex by the fact that any sensible measure of degree of heritability (e.g., of IQ) is itself dependent upon the range of environments considered.[18] There is unlikely to be

[18]For a simple example: an environment of extreme malnourishment during development will tend to render all IQs alike, so the measured heritability of IQ will drop. This, of course, does not mean that the "true" heritability of IQ is so low, but that any heritability report is properly made relative to a distribution of environments within which it has been measured. People who simply assert that the heritability of IQ is 80% (or whatever) across the board (e.g., across all cultures or human environments) are simply ignorant.

a fixed answer to such issues, other than the obvious injunction to pay proper respect to both factors, environmental and hereditary, when attempting to understand a causal story.

Overconfidence

According to Joseph Butler (1736), "probability is the very guide to life." For our subjective probabilities to aid us as much as possible they should reflect the physical probabilities of events as closely as possible.[19] Thus, someone who bets on football but knows absolutely nothing about the game will do better in the long run making only bets at even odds (assuming she or he is setting the odds and others, more knowledgeable, are deciding whether to accept the other side of the bet) — in other words, such a person's maximum expected profit is zero. Of course, to maximize profit (above zero) from gambling on football one ideally has to have as much knowledge about the game as possible and incorporate this knowledge in the choice of odds. That is, it would be ideal to have probabilities that are as *extreme* as possible (near to 1 or 0), while at the same time tracking one's understanding of the game. In fact, someone who is very knowledgeable may be able to profit while occasionally setting odds that are more extreme than those of the professional book-makers (i.e., winning from the professional bookies). In short, to maximize returns one must combine maximum knowledge of the game with the precise *calibration* of probabilities relative to that knowledge (see theorem 6.1.2 of Cover and Thomas, 1991).

Humans rarely show good calibration, however. We can measure calibration by, for example, asking people to assign probabilities to some collection of assertions pulled out of a fact book, half of them being negated. Perfect calibration would mean that a subject assigned, say, a 0.6 probability of truth to a subset of sentences of which 60% turned out to be true, etc. The tendency across a fair range of cases is for people to be overconfident of the truth (or falsity) of statements that lie towards the extremes; that is, if the "objective" probability of a statement is 0.8, people are likely to assign it a probability greater than 0.8, whereas roughly equally likely statements tend to be more accurately assessed (see Lichtenstein et al., 1982).

Overconfidence is enhanced when the judgment is made more difficult or complex (Pitz, 1974). Another way of expressing this is that easier tasks are assessed more accurately. This suggests that training in particular domains will improve calibration on problems within the domain, by making people more knowledgeable and therefore judgments easier; and, indeed, experimental studies support this conclusion (see Pitz, 1974, and Edwards and von Winterfeldt, 1986). In general, the literature on overconfidence suggests that we should tend to qualify predictions put forward by non-experts (including ourselves), particularly when they are extreme.

Hindsight and Memory

At the end of our inductive reasoning comes hindsight. In this case psychologists do not mean the 20-20 vision of outcomes which all of us retrospectively have. Rather, they mean the remarkable *foresight* which we all seem to have in *retrospect*. That is, people after the fact are prone to attribute to themselves and others a greater ability to predict what in fact happened than anyone actually had at the time. Indeed, people are prone to misremember what they in fact predicted, exaggerating their understanding of the situation at the time (Fischhoff, 1980, 1983). This should suggest some caution in assessing such claims, including self-attributions (for example, your own recollections), when they report predictions or interpretations for which no objective evidence is available.

[19] Here I am relying upon the distinction between subjective and physical probabilities widely endorsed, for example by Nagel (1938) and Carnap (1962). For an illuminating discussion of the relation between the two kinds of probability see Lewis (1980).

There is a long and substantial literature reporting the unreliability of eyewitness identifications in the courts (see Schachter, 1996, chapter 4), which supports caution in acceding to the conviction which most of us give to our memory claims. A recent meta-analysis of that literature has found little or no correlation between the confidence of eyewitnesses and the accuracy of their identifications (Sporer et al., 1995). In general, unaided memory is a malleable and fickle thing. Again, significant doubts have been raised recently over the alleged phenomena of "facilitated communication," "recovered memories" and "multiple-personality disorder." Facilitated communications are those supposedly elicited by facilitators when the subjects are themselves unable to use language unaided. The facilitators, in some cases, place their hands over those of the subjects on a keyboard and "amplify" the subjects' hand movements until some meaningful text appears. In a case in Melbourne in 1990 a severely intellectually disabled woman "alleged" in a facilitated communication that her father and brother had repeatedly raped her. This resulted in action by the authorities which was later reversed. In another Victorian case a staff member in a home for the disabled was fired on the basis of similar testimony. Even if the amplified testimony in such cases is based upon some underlying real impulses, it is clear that the technique is open to serious abuse by incorporating ideas from the facilitator, whether or not such incorporation is intentional (Hudson, 1995). Recovered memories are those which people have revealed during special therapy sessions, often decades after the alleged events being recalled. Many of these cases also have involved allegations of sexual abuse and have led to great social disruption and multi-million dollar court cases. In the way of therapeutic influence on mental phenomena, the ancestor of these developments may be multiple-personality disorder (MPD). MPD is a syndrome which some psychiatrists have reported in which people apparently display multiple, distinct personalities. These personalities supposedly control the individual at different times and have limited awareness of each other. Unlike other severe psychological disorders, such as schizophrenia, depression, and epilepsy, MPD has no history prior to the development of psychiatry in the nineteenth century. The psychotherapeutic treatment involved is detailed in Putnam (1989), which includes an extensive discussion of methods of eliciting alternate personalities, using "truth drugs" and hypnosis — which are well known to increase people's suggestibility. Merskey (1992) claims that there is no case of MPD in which the disorder has been observed prior to such therapeutic intervention, and he suggests that MPD is iatrogenic (therapy induced) rather than caused by repressed memories of abuse, as is the standard account. It is worth pointing out in support of Merskey's view that the incidence of MPD has grown explosively just as the therapeutic techniques involved have grown in usage. In 1972 Horton and Miller (1972) reported fewer than a dozen cases from the preceding fifty years; by 1986 six thousand cases had been diagnosed (Coons, 1986). As Hacking (1995, note 14 p. 270) remarks, "the most reliable predictor of the occurence of multiple personality is a clinician who diagnoses and treats multiples"; Hacking also notes the curious geographical correlation between diagnosis and therapeutic techniques (p. 14).

The psychological literature is very clear that reported memories are manipulable,[20] and, so, such evidence, produced in response to therapy, is suspect in much the same way that digitized records, such as electronic mail, audio and video recordings, ought to be held suspect when produced by a computer laboratory. Despite the fact that much of the evidence so reported is likely to be true, *all* such evidence is suspect.

My own opinion is that the sum of these "memories" makes for a very sad story in the history of science and society, one in which existing research on the fallibility and manipulability of human memory has been largely ignored. Clearly, the failure to acknowledge our own, human cognitive weaknesses can have disastrous consequences.

[20] In a literature review Lindsay and Reed (1994) report "the ground for debate has shifted from the question of the possibility of therapy-induced false beliefs to the question of the prevalence of therapy-induced false beliefs."

5.3.1 Bibliographic Notes

Nisbett and Wilson (1977) had a major impact on thinking about thinking, revealing how unreliable humans are as witnesses to their own cognitive processes. Many other influential papers in the psychology of induction were collected in Kahneman, Slovic and Tversky (1982). Nisbett and Ross (1980) provide an excellent overview. Holland et al. (1986) present an interesting theory of induction for artificial intelligence with reference to the psychological work on induction.

A general, and highly readable, review of the science of memory is Schachter's *Searching for Memory* (1996). A recent collection of research articles on the false memory debate is Pezdek and Banks (1996), including reports on the subject by the American Psychological Association and the British Psychology Society.

6 Research Ethics

Here I describe some of the more important practices in the conduct of experimental research. I shall not attempt to justify the standards which these practices exemplify; that is a problem for ethics as a branch of philosophy. However important and interesting the study of ethics may be, my purposes here are more practical, namely to provide guidance or advice to those new to research and so possibly not aware of the standards against which their actions are likely to be judged. Failure to understand those standards can lead to, and has led to, serious and unhappy disputes.

These guidelines are aimed more at those pursuing academic research, rather than industrial research; still, industrial researchers must frequently examine the results of academic research, and so may find the practices which lead to them of interest. The following subsections are organized alphabetically.

6.1 Authorship

The authorship of a research report should reflect the intellectual contribution of the named individuals to the report. Anyone who has contributed substantially to the research outcome reported upon should be invited to be an author. No one who is uninvited (or who is unwilling) should be listed as an author. Also, anyone else who has contributed more than trivially should be acknowledged in a footnote or an acknowledgements section. Clearly, the choice between, say, ignoring someone's casual but relevant remark, acknowledging someone's help, and inviting someone to co-authorship is a matter of judgment — and may be a proper matter of discussion between the potential co-authors.

It is an ethically dubious, and substantially misleading, practice to "reward" a supervisor or laboratory research leader with co-authorship of a paper simply for existing.[21] If a supervisor has contributed significant ideas and motivation for the research, then co-authorship is most appropriate. But if your supervisor has made no significant intellectual contribution, then his or her forcing a co-authorship from you is an abuse of power. A plausible standard for authorship is knowledge and responsibility (Luey, 1995, p. 16): an author "should have direct knowledge of the conduct and results of the entire study and should be willing to take responsibility for its conduct, data and conclusions." A relevant practice some journals have adopted is to request that each author identify his or her exact contribution to the paper. This makes for a natural standard for authorship: any "author" who would be embarrassed by a public statement identifying his or her contribution should withdraw. Naturally, many

[21] Monash University regulations report that "honorary authorship" is unacceptable: " 'Honorary authorship' occurs when a person is listed as an author of a publication when they have not participated in any substantial way. . . ." (*Monash University Education and Research Policies*, Part III §2.1(b).)

cases will not be clear-cut, and good judgment will be called for, preferably based upon an open discussion of the issues between the parties involved.

Since authors are involved in, or at least aware of, the entire process of research leading up to the completion of the research report, they should also be full participants in its publication. It is unacceptable to submit papers for publication on behalf of other authors without their knowledge and consent; likewise, it is unacceptable to present conference papers in their names without their consent.

Somewhat related issues are involved in determining author order. It is natural for readers to assume that the first author contributed at least as much to authoring the paper as the second author, etc. Given that, it is basically dishonest to violate such expectations in selecting an author order. In some cases the dishonesty might be forgivable, or least understandable. Some supervisors will prefer that their students' names come first, in order to help them get started on a career; or, a research group might have an agreement to rotate author order, with the overall impression of equality being accurate, even though when applied to a particular paper it may be inaccurate. In other cases, author order reflects the political power of the leading author(s) rather than the intellectual power of his or her contribution. Since the outcome of research is an intellectual product, and since that is its sole proper role in assessing the work of the authors, it is only fair and reasonable to weigh authors' contributions in terms of their intellectual significance. Thus, a single idea, if it is important enough, can easily outweigh all the other work needed to put the idea in a publishable form. Political power is irrelevant to the proper evaluation of authorship.

6.2 Faking Results

It is obvious that faking research results is unethical; it also violates all the guidelines of professional practice, such as those codified in university regulations on research. Flagrant violation leads to dismissal and the end of careers.

Unintentional fakery is also a possibility. "Cleaning" data, the biased selection of problems favoring one's algorithm, testing against data that were in part used to develop or train one's algorithm — any of these and other mistakes may lead to misleading experimental results. In principle, the error should be discovered when a subsequent unbiased study fails to confirm your results. For this reason it is reasonable to expect that authors will include in their reports, if this is possible, enough information about the methodology employed for the study to be repeated and the same or similar results reproduced. If inputs and parameter settings are necessary for this, they should be retained in an archive.[22] In practice, the person most likely to discover such an error is the researcher her- or himself, when pursuing further research in the same area. In order to avoid wasting the time and money of other researchers, it is appropriate then to make a public retraction and correction of the prior report.

6.3 Freedom of Inquiry

Society, in many guises, continually asserts pressure to direct and constrain scientific research. The history of science shows clearly and repeatedly, however, that scientific progress is greatest when both inquiry and criticism are free to find their own way, relatively unhindered by political, religious, cultural and ideological restrictions. Soviet genetics, for example, was severely hampered for decades by the imposition of Lysenkoism by the political authorities. As a result, genetics and related scientific areas were and are extremely weakly represented fields in the ex-Soviet states, even while those states export expertise in other areas, such as physics and mathematics. Italy was the major intellectual center of Renaissance Europe —

[22] Monash University regulations require the retention of research data or the inputs necessary to regenerate the data for a minimum of five years.

until the Pope condemned Galileo, when scientific leadership passed to Reformation countries, such as England. Since leadership in scientific research produces leadership in technology and economy, the defence of freedom of inquiry is not merely an academic problem.

Some research is condemned because of the uses to which it may, or perhaps probably will, be put, such as military uses. However, if those uses are bad, then the fault lies with the user of the ideas rather than their originator. It may be that the originator, and the rest of us, have a moral obligation to combat those unethical uses of research; it does not follow that the research itself is or was unethical. A relevant consideration in such cases is that scientific inquiry uncovers the true nature of the world (perhaps fallibly, with only a slow convergence on that nature): since that nature will likely reveal its secrets to sustained inquiry, the attempt to control unethical behavior by spreading or maintaining ignorance is very likely to fail.

6.4 Intellectual Property Rights

Any software which you produce, and any report which you write, is automatically under copyright protection; that is, *you* own the copyright, unless you have specifically signed away that right.[23] However, honours research is often undertaken under an explicit understanding between the supervisor and the student that the work is a part of a broader research effort, and either the student will be making use of pre-existing software or making her or his software available to subsequent researchers. In such a case, although the student may be under no legal obligation to allow the software to be used subsequently, the student cannot plausibly feign ignorance and clearly is under a moral obligation to allow its subsequent use. Better than relying upon anyone's moral sense would be a written agreement at the beginning of research to allow the supervisor to use the resultant software without giving away copyright.

If anything which you design is potentially patentable, and if you wish to attempt to patent it, then, despite the scientific preference for openness, discretion is called for. If you publish your design prior to filing a patent application, it will no longer be patentable (or, a defence of the patent will likely fail). What constitutes publication is unclear, but you should probably avoid making your work available on the internet or through technical reports until a patent application has been filed. Simply including material in an honours report would not constitute publication, so long as the report is not distributed beyond its examiners. Since higher degree theses are made available to a wider community, inclusion of patentable ideas within such a thesis would more likely constitute publication.[24]

6.5 Overpublishing

Overpublishing is not a crime, but it probably should be. It comes in two major forms: publishing, republishing and rerepublishing the same information *ad nauseam*; and, slicing and dicing a research report into numerous "least publishable units." Unfortunately, both practices are fairly common. Both practices are motivated by the nearly universal administrative method of assessing research performance: counting refereed publications. Since the reward for having "high research productivity," as measured by bean counters, is direct and substantial, the temptation to "play the game" by generating intellectually worthless publications is simply too great for many to resist.

[23] However, programs and written works that constitute teaching materials produced by you in the employment of Monash University are owned by Monash University as a condition of your employment. This restriction, by the way, is in my opinion unfortunate: it undoubtedly leads some academics to teach with inferior materials, in order to preserve their copyright over better materials not yet published.

[24] Students and staff of Monash University are obligated by regulation to take any patentable ideas to the university intellectual property officer, as the university retains the right to patent applications in its name.

Republication is normally effected by making trivial alterations to a prior publication's format, title, section titles, introduction and conclusion, and perhaps regrouping data so as to give them a different form as well. The result will often not be identified by the referees as essentially a duplicate of a prior publication, and so may be accepted. This can lead to a significant problem when the outcomes of experiments are reported, since a meta-analysis reviewing the results of multiple related experiments may overweight the research, counting it as many times as it has been published. In medical research, for example, this could cause improper medical treatments to be given, killing or injuring patients. Eventually, however, the informed community of researchers will notice the duplication, and so there is also personal down-side for the author, namely that one will have achieved notice in the community primarily by producing a bad smell. Bean counters, however, have no sense of taste or smell, so the up-side remains.

Publishing in least publishable units maximizes publications while minimizing intelligibility. In order to get a reasonable view of the product of the research project, the author is forcing readers to track down and read a large number of related articles. The only way in which this might be defensible is when page limits on conference contributions more or less force this outcome.

6.6 Plagiarism

Science thrives on openness, robust criticism and honesty and withers under secretiveness, dishonesty and political interference. Therefore, in addition to the usual societal sanctions, the scientific community applies its own sanctions against those who violate the trust accorded researchers by presumption. In particular, those who are found to have plagiarized others' works will normally not get any second chance either to continue in their dishonesty or to redeem themselves. There is no reason that you should not take full advantage of others' work, including the work of your peers; but it is both prudent and a courtesy to acknowledge that assistance when it has any significant impact on your own research.

6.7 Refereeing

An active researcher will encounter both sides of refereeing, having his or her own papers reviewed and reviewing the work of others. Ideally, refereeing is done double blind, with neither the author nor referee knowing the identity of the other. Many journals and conferences do not delete author identification, and some authors ensure by citations or other devices that their identities are known to reviewers. This practice tends to discriminate against new researchers and researchers outside the particular group organizing the conference or running the journal. Some referees are sufficiently intellectually dishonest to deliberately bias outcomes in favor of their students, collaborators and friends. I have even heard such behavior defended with the remark that competing schools of thought will have their day when its representatives rise to sufficiently senior positions! There is, of course, a direct and immediate benefit to oneself and to one's supporters. But, since the consequence — more or less by definition — is the advancement of inferior research ahead of superior research, the net gain to the community at large is negative; this is an extremely anti-social act. Double blind refereeing, while being an imperfect answer to such bias (identities can often be guessed, schools of thought can always be guessed), at least minimizes the opportunities for bias, whether the bias be intentional or not.

When receiving a referee's report, whether the outcome is acceptance or rejection, it is well to take the criticisms seriously, even when they reflect the referee's misunderstanding. After all, the referee is likely to have considerable expertise in your research area ("likely" because this is certainly not guaranteed); so, if the referee has misunderstood your paper,

then your targeted reader is likely to have very little chance of understanding it. In other words, the fault is yours, not the referee's. Occasionally, it will be clear that the fault is indeed with the referee, either because the referee does not have sufficient background to evaluate your work or because the referee has not bothered to read your paper with even moderate care or attention. This is obviously a disservice both to you and to the journal or conference. The only saving grace is that by experiencing such a rejection you might be encouraged when acting as a referee yourself to do so with care; and if you agree to referee a paper which turns out not to be one you are capable of refereeing, then perhaps (unlike some) you will simply say as much to the editor and return the paper unreviewed.

6.8 Submission

When a report is complete, the next step is to choose a venue for its dissemination. It is generally acceptable to distribute it as a technical report, or via a pre-publication electronic site, and this will not normally jeopardize prospects for publishing the paper in a journal or conference (it may, however, jeopardize patent prospects; see § *Intellectual Property Rights*). Once you have assigned copyright to a publisher, however, you must cease distributing your paper via the Internet. Note that you can assign copyright only explicitly: if you publish your paper in a proceedings, and the publisher has not formally requested that you assign copyright, then you retain copyright to your paper. This is true even if the publisher asserts copyright over the proceedings. In that case, the publisher can defend the proceedings as a whole, its organization and content, from being republished by someone else; however, the publisher has no right to also prevent you from republishing your individual article.

For guidance in selecting a journal or conference for publication consult your supervisor and peers. If you approach those which have published some or many of the papers which you cite, you will not go far wrong. Having selected a venue, you should go to the trouble of looking at a recent issue and formatting your report according to their current style requirements.

By default you should *not* submit a paper to multiple venues simultaneously. Some conferences allow multiple submissions, but if so, their instructions for authors will state this explicitly, and they usually require authors to report the multiple submissions in advance — very likely disadvantaging the submission during its review. If your preferred venue rejects your report, you should take the rejection in good grace (it happens to everyone), use the referee reports to improve your paper, and resubmit the result to another journal or conference — unless your conclusion is that the paper is not salvageable.

6.9 Bibliographic Notes

Frankena (1973) is a clear introduction to the philosophical study of ethics. Miller and Hersen (1992) is a useful collection of articles on research ethics, concentrating on cases of fraud, and includes a number of case studies. Erwin, Gendin and Kleiman (1994) is an excellent anthology on ethical issues in research, including human and animal experimentation, and reprints the best paper I know on value judgments in science (by Michael Scriven). The Panel on Scientific Responsibility and the Conduct of Research (1992, 1993) report various useful guidelines and practices, including those for the preservation of research data.

7 Conclusion

An understanding of the techniques of argument analysis and the principles of hypothesis evaluation will help you to assess arguments as they appear in research writing and will help you to create effective arguments in your own research. An understanding of the inferential weaknesses to which humans are prone will help you locate errors in others' reasoning and

instill a useful caution in your own reasoning. Your research is both developed by and finds its final expression in your research writing, which may be either aided or impaired by your ability to deal with the English language. The development of all of these skills is not easy and never really ends, so I hope my discussion here encourages you to work at these skills and to search for whatever aids may be available in that effort.

Acknowledgement

I thank John Crossley for copious comments on an early draft.

References

Achtert, W.S. and Gibaldi, J. (1985). *The MLA style manual.* New York: Modern Language Association of America.

Adams, D. (1979). *The hitch hiker's guide to the galaxy.* London: Pan.

Adler, M. and van Doren, C. (1972). *How to read a book* (2nd ed.). New York: Simon and Schuster.

American Psychological Association (1994). *Publication manual of the American Psychological Association* (4th ed.). Washington: APA.

Angluin, D. and Smith, C.H. (1983). Inductive inference: Theory and methods. *ACM Computing Surveys, 15*, 237-269.

Baird, D. (1992). *Inductive logic: Probability and statistics.* Prentice Hall.

Barzun, J. and Graff, H.F. (1992). *The modern researcher* (5th ed.). Harcourt, Brace Jovanovich.

Bayes, T. (1763) An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society, 53.* Reprinted in *Biometrika, 45* (1958), 293-315.

Bransford, J., Sherwood, R., Vye, N., and Rieser, J. (1986). Teaching thinking and problem solving. *American Psychologist, 41*, 1078-1089.

Butler, J. (1736/1906) *Analogy of Religion.* London: Methuen.

Carnap, R. (1962). *The logical foundations of probability* (2nd ed.). University of Chicago.

Carroll L. (1895). What the tortoise said to Achilles. *Mind, 4*, 278-280.

The Chicago manual of style (13th ed.). (1993). Chicago: University of Chicago Press.

Coons, P.M. (1986). The prevalence of multiple personality disorder. *Newsletter of the International Society for the Study of Multiple Personality and Dissociation, 4(3)*, 6-8.

Copi, I.M. (1972). *Introduction to logic* (4th ed.). New York: Macmillan.

Council of Biology Editors (1990). *Ethics and policy in scientific publication.* Bethesda, MD: Council of Biology Editors.

Council of Biology Editors (1994). *Scientific style and format: The CBE manual for authors, editors, and publishers.* Cambridge University Press.

Cover, T. and Thomas, J. (1991). *Elements of information theory.* New York: John Wiley and Sons.

Cribb, J. (1996). *White death.* Harper Collins.

Edwards, W. and von Winterfeldt, D. (1986). On cognitive illusions and their implications. *Southern California Law Review, 59*, 401-451.

Ennis, R.H. (1962). A concept of critical thinking. *Harvard Educational Review, 32*, 81-111.

Erwin, E., Gendin, S. and Kleiman, L. (Eds.) (1994) *Ethical issues in scientific research: An anthology.* New York: Garland Publishing.

Evans, J.St.B.T. (1989). *Bias in human reasoning: causes and consequences.* London: Lawrence Erlbaum and Associates.

Fischhoff, B. (1980). For those condemned to study the past: Reflections on historical judgment. In R.A. Shweder and D.W. Fiske (Eds.) *New directions for methodology of social and behavioral science: No. 4. Fallible judgment in behavioral research* (pp. 79-93). San Francisco: Jossey-Bass.

Fischhoff, B. (1983). Reconstructive criticism. In P.C. Humphreys, O. Svenson and A. Vari (Eds.) *Analysing and aiding decision processes* (pp. 515-523). Amsterdam: North Holland.

Fowler, H.W. (1983). *A dictionary of modern English usage*, 2nd ed., with corrections (revised by Sir Ernest Gowers). Oxford: Oxford University Press.

Frankena, W. K. (1973) *Ethics*, second edition. Englewood Cliffs: Prentice-Hall.

Franklin, A. (1986). *The neglect of experiment.* Cambridge University.

Freeman, J.B. (1991). *Dialectics and the macrostructure of arguments.* Berlin: Foris Publications.

French, S. and Kamminga, H. (Eds.) (1993). *Correspondence, invariance, and heuristics: essays in honour of Heinz Post.* Dordrecht: Kluwer Academic.

Giere, R.N. (1991). *Understanding scientific reasoning*, 3rd ed. Fort Worth: Holt, Rinehart, and Winston.

Gold, I. (1994). *Writing an honours paper in philosophy.* Unpublished ms.

Good, I.J. (1985). *Good thinking: The foundations of probability and its applications.* University of Minnesota.

Graves, R. and Hodges, A. (1947). *The reader over your shoulder: A handbook for writers of English prose* (2nd ed.). London: Cape.

Hacking, I. (1995). *Rewriting the soul: multiple personality and the sciences of memory.* Princeton University.

Hayakawa, S.I. (1987). *Choose the right word: A modern guide to synonyms* (2nd ed.). New York: Harper and Row.

Holland, J.H., Holyoak, K., Nisbett, R. and Thagard, P. (1986). *Induction: processes of inference, learning, and discovery.* MIT Press.

Horton, P. and Miller, D. (1972). The etiology of multiple personality. *Comparative Psychology, 13*, 151-159.

Howson, C. and Urbach, P. (1993). *Scientific reasoning: The Bayesian approach* (2nd ed.). Chicago: Open Court.

Hudson, A. (1995). Disability and facilitated communication: a critique. *Advances in clinical child psychology, 17.*

Jeffrey, R. (1983). *The logic of decision* (2nd ed.). NY: McGraw-Hill.

Johnson-Laird, P.N. (1983). *Mental models.* Harvard University.

Johnson-Laird, P.N. and Wason, P.C. (Eds.) (1977). *Thinking: Readings in cognitive science.* Cambridge University.

Kahneman, D., Slovic, P. and Tversky, A. (Eds.) (1982). *Judgment under uncertainty: heuristics and biases.* Cambridge University Press.

Koertge, N. (1993). Ideology, heuristics, and rationality in the context of discovery. In French and Kamminga (1993) (pp. 125-136).

Korb, K.B. (1992). *A pragmatic Bayesian platform for automating scientific induction.* Ph.D. dissertation, Indiana University.

Korb, K.B. (1994). Infinitely many resolutions of Hempel's paradox. in R Fagin (Ed.) *Theoretical aspects of reasoning about knowledge V, Asilomar, California* (pp. 138-149). Morgan Kaufmann.

Korb, K.B. (1995). Inductive learning and defeasible inference. *Journal of Experimental and Theoretical Artificial Intelligence, 7*, 291-324.

Korb, K.B. (1997). Probabilistic causal structure. Submitted for the forthcoming volume *Causation and natural laws: Australasian studies in the history and philosophy of science* edited by Howard Sankey.

Lewis, D. (1980). A subjectivist's guide to objective chance. In R. Jeffrey (Ed.) *Studies in inductive logic and probability, vol. II*, University of California Press.

Lindsay, D.S. and Reed, J.D. (1994) Incest resolution psychotherapy and memories of childhood sexual abuse. *Applied Cognitive Psychology, 8*, 281-338.

Lovell, D.W. and Moore, R.D. (1992). *Essay writing and style guide for politics and the social sciences.* Canberra: Australasian Political Studies Association.

Luey, B. (1995) *Handbook for academic authors*, third edition. Cambridge: Cambridge University Press.

Mackie, J.L. (1967). Fallacies. In P. Edwards (Ed.) *The encyclopedia of philosophy* (pp. 169-179). NY: Macmillan.

Malotki, E. (1983). *Hopi time: A linguistic analysis of temporal concepts in the Hopi language.* Berlin: Mouton.

McCarthy, J. (1952). The red-tinted Washington crowd. Campaign speech at Appleton, Wisconsin, November 3, 1952. Reported in A. T. Rottenberg *Elements of argument*, 1991, Bedford Books, Boston.

Mencken, H.L. (1963). *The American language* (4th ed.). NY: Knopf.

Merskey, H. (1992). The manufacture of personalities. *British Journal of Psychiatry, 160*, 327-340.

Miller, D.J. and Hersen, M. (Eds.) (1992). *Research fraud in the behavioral and biomedical sciences.* New York: Wiley.

Moxley, J.M. (1992). *Publish, don't perish.* Wesport, CT: Praeger.

Nagel, E. (1938). *Principles of the theory of probability. International encyclopedia of unified science, v.1; no. 6.* Chicago University.

Nisbett, R., Borgida, E., Crandall, R., and Reed, H. (1976). Popular induction: Information is not necessarily informative. In J.S. Carroll and J.W. Payne (Eds.) *Cognition and social behavior* (pp. 113-133). Hillsdale, NJ: Erlbaum.

Nisbett, R., Fong, G.J., Lehman, D.R., and Cheng, P.W. (1987). Teaching reasoning. *Science, 238*, 625-631.

Nisbett, R., and Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment.* Englewood Cliffs, N.J.: Prentice-Hall.

Nisbett, R., and Wilson, T.D. (1977). Telling more than we know: Verbal reports on mental processes. *Psychological Review, 84*, 231-295.

Oliver, J. and Baxter, R. (1994). *MML and Bayesianism: Similarities and differences.* Technical Report TR 207, Dept. Computer Science, Monash University.

Orwell, G. (1950). Politics and the English language. In G. Orwell *Shooting an elephant.*

Passmore, J. (1969). *Talking things over* (4th ed.). University of Melbourne.

Panel on Scientific Responsibility and the Conduct of Research (1992) *Responsible science: Ensuring the integrity of the research process, vol. 1.* Washington: National Academy Press.

Panel on Scientific Responsibility and the Conduct of Research (1993) *Responsible science: Ensuring the integrity of the research process, vol. 2.* Washington: National Academy Press.

Paul, R.W. (1984). Critical thinking: Fundamental to education for a free society. *Educational Leadership (Sept)*, 4-14.

Pezdek, K. and Banks, W.P. (Eds.) (1996). *The recovered memory–false memory debate.* San Diego: Academic Press.

Peirce, C.S. (1940). The probability of induction. In J. Buchler (Ed.) *Philosophical writings of Peirce.* London: Routledge and Kegan Paul.

Phillips, L.D., Hays, W.L. and Edwards, W. (1966). Conservatism in complex probabilistic inference. *IEEE Transaction on Human Factors in Electronics, 7*, 7-18.

Pitz, G. (1974). Subjective probability distributions for imperfectly known quantities. In L.W. Gregg (Ed.) *Knowledge and cognition* (pp. 29-41). New York: John Wiley and Sons.

Pollard, P. and Evans, J.St.B.T. (1980). The influence of logic on conditional reasoning performance. *Quarterly Journal of Experimental Psychology, 32*, 605-624.

Polya G. (1968). *Mathematics and plausible reasoning.* Princeton University.

Popper, K.R. (1959). *The logic of scientific discovery.* New York: Basic Books. Translation, with new appendices, of *Logik der Forschung* (1934), Vienna.

Putnam, F. (1989). *Diagnosis and treatment of multiple personality disorder.* New York: Guilford Press.

Ramsey, F.P. (1931). Truth and probability. *The foundations of mathematics and other logical essays.* London: Routledge and Kegan Paul.

Salmon, M. (1990). Informal reasoning and informal logic. In J.F. Voss, D.N. Perkins, J.W. Segal *Informal reasoning and education* (pp. 153-168). Hillsdale, N.J.: Lawrence Erlbaum.

Savage, L.J. (1954). *Foundations of statistics.* John Wiley and Sons.

Schachter, D. (1996). *Searching for memory.* New York: Basic Books.

Shrader-Frechette, K.S. (1994). *Ethics of scientific research.* Lanham, MD: Rowman and Littlefield.

Siegel, H. (1988). *Educating reason: Rationality, critical thinking, and education.* New York: Routledge.

Skyrms, B. (1986). *Choice and chance: An introduction to inductive logic* (3rd ed.). Belmont, Calif: Wadsworth.

Sporer, S.L., Penrod, S., Read, D. and Cutler, B. (1995). Choosing, confidence, and accuracy: A meta-analysis of the confidence-accuracy relation in eyewitness identification studies. *Psychological Bulletin, 118*, 315-327.

Strunk, W. and White, E.B. (1979). *The elements of style* (3rd ed.). NY: Macmillan.

Toulmin, S. (1958). *The uses of argument.* Cambridge.

Toulmin, S., Rieke, R., and Janik, A. (1984). *An introduction to reasoning* (2nd ed.). New York: Macmillan.

Tversky, A. and Kahneman, D. (1971). Belief in the Law of Small Numbers. *Psychological Bulletin, 76*, 105-110.

Tufte, E. (1990). *Envisioning information.* Cheshire, Conn: Graphics Press.

Tversky, A. and Kahneman, D. (1982). Judgments of and by representativeness. In Kahneman et al. (1982), pp. 153-160.

Vick, C. and Ramamoorthy, C. (Eds.) (1984). *Handbook of Software Engineering.* Van Nostrand Reinhold.

Walton, D.N. (1989). *Informal logic: a handbook for critical argumentation.* Cambridge University Press.

Wason, P.C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology, 12*, 129-140.

Wason, P.C. (1966). Reasoning. In B.M. Foss (Ed.) *New horizons in psychology I.* Harmandsworth: Penguin Books.

Whorf, B.L. (1956). *Language, thought, and reality.* MIT Press.

Williams, J.M. (1990). *Style: Towards clarity and grace.* University of Chicago.

Winston, P. (1992). *Artificial intelligence* (3rd ed.). Addison-Wesley.

Zinsser, W. (1988). *Writing to learn.* New York: Harper and Row.

Zinsser, W. (1990). *On writing well* (4th ed.). Harper Perennial.

Zukerman, I., Korb, K. and McConachy, R. (1996). Perambulations on the way to an architecture for a nice argument generator. In *Notes of the ECAI-96 Workshop Gaps and bridges: New directions in planning and natural language generation,* pp. 31-36, Budapest, Hungary.