

Critical reading



**Hannah Lewis, Dana Macgregor
and Hilary M Jones**

Contents

1. Why are you reading?
2. What is critical reading?
3. Papers, review articles and textbooks
4. How to read a paper
5. The peer review process
6. Bias

Much of the information in this handout is taken from:

1. Critical Reading, making sense of research papers in life sciences and medicine. Ben Yudkin. Routledge, London and New York, 2006.
2. How to read a paper, the basics of evidence-based medicine. 3rd Edition. Trisha Greenhalgh. Blackwell publishing, Oxford, 2006.
3. A Guide to the Critical Reading of Scientific Research Papers
<http://www.greenwich.ac.uk/~bj61/talessi/tlr51.html>

1. Why are you reading?

It may sound strange, but before you start your reading, you should ask yourself the question

“Why am I reading?”

This is because it's so easy to get absorbed and side-tracked. So if you're to make the most efficient use of your time, then you need to stay focussed and make sure that you get the information that you are looking for.

Some answers to the question might be:

- For general interest or background information
- To find out what the latest developments are in a field
- To seek evidence to support or refute your ideas
- To broaden avenues of research
- To find out exactly how a certain piece of research was done
- Because I have to

How you answer this question should dictate how and what you read, and should also save you a lot of time. Even if you are reading solely because it was assigned, to effectively learn the information that was meant to be imparted in the assignment, determine **why** the reading was assigned and approach it with that question in mind.

There are several reasons you would want to read to obtain background knowledge. Some of these reasons include writing a report on a subject, in preparation for a seminar you are going to attend, or needing a general overview to inform your own research interests.

On the other hand, if you are embarking on experimental work in a specific area, you will need to be fully versed in the current state of knowledge. It will be important to read as widely as possible, not only about the background but also about the latest developments that might affect your work.

To be a good scientist, it is important to be aware of the latest developments in your field. Knowing what is happening in your field allows you to use that information to guide your own research toward topics of interest and away from things that have already been done. In a large or fast-moving field, papers will be appearing as rapidly as you can keep up with them, and you will need to be selective about what to read. Reviews are a useful place to start, however they do not get published on a regular basis and are one author's summary of the

papers they found interesting and relevant. As such, reviews can be biased and myopic and should be used as a place to start collecting primary literature rather than an endpoint in themselves. means to an end. Supervisors or colleagues with more experience in the field can help by suggesting important papers or researchers that they admire. There are also numerous services available that will e-mail you relevant abstracts as they appear. For example, PubCrawler (<http://pubcrawler.gen.tcd.ie/>) advertises itself as a '...free "alerting" service that scans daily updates to the NCBI Medline (PubMed) and GenBank databases. PubCrawler helps keeping scientists informed of the current contents of Medline and GenBank, by listing new database entries that match their research interests.'

2. What is critical reading?

Taken from: http://www.criticalreading.com/critical_reading.htm
Accessed 30/04/2009

Facts v. Interpretation

To **non**-critical readers, texts provide facts. Readers gain knowledge by memorizing the statements within a text. This is a surface-level of learning and won't get you very far at all in your research, which has to be an in depth understanding of your subject.

To the **critical** reader, any single text provides but one portrayal of the facts, one individual's "take" on the subject matter. Critical readers thus recognise not only **what** a text says, but also **how** that text portrays the subject matter. They recognise the various ways in which each and every text is the unique creation of a unique author.

A non-critical reader might read a history book to learn the facts of the situation or to discover an accepted interpretation of those events. A critical reader might read the same work to appreciate how a particular perspective on the events and a particular selection of facts can lead to particular understanding.

What a Text Says, Does, and Means: Reaching for an Interpretation

Non-critical reading is being satisfied with recognising what a text *says* and restating the key remarks.

Critical reading goes two steps further. Having recognised what a text **says**, it reflects on what the text **does** by making such remarks. Is it offering examples? Arguing? Appealing for sympathy? Making a contrast to clarify a

point? Finally, critical readers then infer what the text, as a whole, **means**, based on the earlier analysis.

These three steps or modes of analysis are reflected in three types of reading and discussion:

- What a text **says** – **restatement**
- What a text **does** – **description**
- What a text **means** – **interpretation** .

You can distinguish each mode of analysis by the subject matter of the discussion:

- What a text says – restatement – talks about the same topic as the original text
- What a text does – description – discusses aspects of the discussion itself
- What a text means – interpretation — analyses the text and asserts a meaning for the text as a whole

Goals of Critical Reading

Textbooks on critical reading commonly ask students to accomplish certain goals:

- to recognise an author's purpose
- to understand tone and persuasive elements
- to recognise bias

Notice that none of these goals actually refers to something on the page. Each requires inferences from evidence within the text:

- recognising purpose involves inferring a basis for choices of content and language
- recognising tone and persuasive elements involves classifying the nature of language choices
- recognising bias involves classifying the nature of patterns of choice of content and language

Critical reading is not simply close and careful reading. To read critically, one must actively recognise and analyse evidence upon the page.

Analysis and Inference: The Tools of Critical Reading

These pages are designed to take the mystery out of critical reading. They are designed to show you **what to look for (analysis)** and **how to think about what you find (inference)**.

The first part —what to look for— involves recognising those aspects of a discussion that control the meaning.

The second part —how to think about what you find— involves the processes of inference, the interpretation of data from within the text.

Recall that critical reading assumes that each author offers a portrayal of the topic. Critical reading thus relies on an examination of those choices that any and all authors must make when framing a presentation: choices of content, language, and structure. Readers examine each of the three areas of choice, and consider their effect on the meaning.

Critical Reading v. Critical Thinking

We can distinguish between critical reading and critical thinking in the following way:

- Critical *reading* is a technique for **discovering** information and ideas within a text.
- Critical *thinking* is a technique for **evaluating** information and ideas, for deciding what to accept and believe.

Critical reading refers to a careful, active, reflective, analytic reading. Critical thinking involves reflecting on the validity of what you have read in light of our prior knowledge and understanding of the world.

For example, consider the following (somewhat humorous) sentence from a student essay:

Parents are buying expensive cars for their kids to destroy them.

As the terms are used here, **critical reading** is concerned with figuring out whether, within the context of the text as a whole, "*them*" refers to the parents, the kids, or the cars, and whether the text supports that practice. **Critical thinking** would come into play when deciding whether the chosen meaning was indeed true, and whether or not you, as the reader, should support that practice.

By these definitions, critical reading would appear to come before critical thinking: Only once we have fully understood a text (critical reading) can we truly evaluate its assertions (critical thinking).

The Two Together in Harmony

In actual practice, critical reading and critical thinking work together.

Critical thinking allows us to monitor our understanding as we read. If we sense that assertions are ridiculous or irresponsible (critical thinking), we examine the text more closely to test our understanding (critical reading).

Conversely, critical thinking depends on critical reading. You can think critically about a text (critical thinking), after all, only if you have understood it (critical reading). We may choose to accept or reject a presentation, but we must know why. We have a responsibility to ourselves, as well as to others, to isolate the real issues of agreement or disagreement. Only then can we understand and respect other people's views. To recognize and understand those views, we must read critically.

The Usefulness of the Distinction

If critical thinking and critical reading are so closely linked, why is this still a useful distinction?

The usefulness of the distinction lies in its reminder that we must read each text on its own merits, not imposing our prior knowledge or views on it. While we must evaluate ideas as we read, we must not distort the meaning within a text. We must not allow ourselves to force a text to say what we would otherwise like it to say—or we will never learn anything new!

Reading Critically: How Well Does The Text Do What It Does

We can think of a writer as having taken on a job. No matter what the topic, certain tasks must be done:

- a specific topic must be addressed
- terms must be clearly defined
- evidence must be presented
- common knowledge must be accounted for
- exceptions must be explained
- causes must be shown to precede effects and to be capable of the effect
- conclusions must be shown to follow logically from earlier arguments and evidence

As critical readers and writers, we want to assure ourselves that these tasks have been completed in a complete, comprehensive, and consistent manner. Only

once we have determined that a text is consistent and coherent can we then begin to evaluate whether or not to accept the assertions and conclusions.

Thinking Critically: Evaluating The Evidence

Reading to see what a text says may suffice when the goal is to learn specific information or to understand someone else's ideas. But we usually read with other purposes. We need to solve problems, build roads, write legislation, or design an advertising campaign. We must evaluate what we have read and integrate that understanding with our prior understanding of the world. We must decide what to accept as true and useful.

As readers, we want to accept as fact only that which is actually true. To evaluate a conclusion, we must evaluate the evidence upon which that conclusion is based. We do not want just any information; we want reliable information. To assess the validity of remarks within a text, we must go outside a text and bring to bear outside knowledge and standards.

3. Papers, review articles and textbooks

Textbooks

These are best for answering the more basic kinds of query. Don't get too bogged down with them or try to read them all the way through! Look through the contents page to identify potential chapters of interest and then skim through the sub-headings of the chapter to locate exactly what will be useful to you. Textbooks can also be a very useful tool when designing the introduction to your own presentations. They can give a good example of how to explain your topic to others without using field-specific jargon or 'lab-isms' that inevitably develop.

Some textbooks have detailed tables of contents at the front that list all the sub-headings. Reading the detailed contents is often an excellent way to obtain a brief summary of the chapter. Once you have honed in on a relatively small number of pages, reading them all in detail is not so laborious.

Reviews

Review papers give an overview of an area of research by summarising research papers in the field. Typically, many tens or some hundreds of papers are reviewed. This saves you the work of finding and reading all the original research, and gives a summary of trends in the results.

"However, be careful; reviews are the author's interpretation of what they alone have read. As we all know, one person cannot read every paper there is on the topic and as such, reviews can be myopic and/or biased. Furthermore, it is important to decide for yourself whether the review's author was a critical reader when they read the primary research paper. Therefore, it is essential to check the critical facts in the primary research paper for yourself before you believe them.

Also, this warning would apply equally well for text books. My opinion is that text books and reviews are essentially on the same plane... they both are reviews of a lot of material summarised into digestible bites. Reviews are aimed at a specialist audience whereas text books need to be more generally useful, but they have many of the same drawbacks."

Dana MacGregor, research associate, University of York

Research papers

Textbooks and review papers will give you a good background to what's going on in your area of research, but you will sooner or later need to get down to reading original research papers. Here you will find all of the specific

experimental details together with the rationale for the research and the author's interpretation of the results.

The goal of this handout and the training session is to give you some tips on how to make your reading and understanding of these papers most efficient.

4. How to read a paper

Whatever your reason for reading the paper, your own critical assessment of it is very important. One useful trick is read the experimental results – the tables and figures together with their legends – at least as closely as the main text. Another is to avoid reading the discussion section until you have come up with your own ideas about what the research findings actually are. Don't rely exclusively on the author's views.

The following pages discuss the various parts of a paper and what they are for.

The Abstract

"The abstract on the other hand is what the author wants the take-home message to be. Because most of us only read the abstract on pubmed before we even open up the paper... it has to attract the right audience, give a concise and precise summary of what the main conclusions are from the data, and still intrigue you enough so that you read the rest of it -- and therefore hopefully cite it in your own work! The abstract is to the paper what a movie trailer is to a film. It has to advertise and give you the plot without giving away the big chase scene."

Dana MacGregor, research associate, University of York

When you read the abstract, you should ask yourself the question:
"what controversy or orthodoxy does this research take as its starting point?"

While the Abstract should give you a brief summary of the paper's main findings, the Introduction give some context to the research...

The Introduction

"The introduction can be a means by which the authors guide the reader to a specific question... in my own work we had to do this very carefully. If you introduce your paper with background material on the mechanisms controlling root system architecture, the readers will be thinking about it as a 'root paper.' If you introduce the SAME data instead as how sugar effects the development of root system architecture, then the reader is thinking of the data as a 'sugar

effect.' Very different audiences, and very different implications are targeted in the introductions of papers."

Dana MacGregor, research associate, University of York

Here are three reasons to read the introduction with some care before you read the rest of the paper:

1. it gives you some idea what background information you need before starting
2. it gives you an insight into the authors' starting point and approach to the subject
3. it helps to focus attention on whether the results presented in the paper actually address the question that they are supposed to answer

Once you have read the abstract and the Introduction you may feel that you know all there is to know about this paper, but take care! Although these sections contain a lot of information, what they omit is crucial. How were the experiments done? Do the raw data actually support the conclusions? Are there some confusing findings that don't appear in the Abstract?

The only way to discover the answers to these questions is to read the paper properly; the highlights chosen by the authors are no substitute. If you don't have time to read the paper, or you decide that it is beyond the scope of what you need to know, that's fine. But don't deceive yourself by thinking that you fully understand what the paper is actually about!

Materials and Methods

Unless your reason for reading the paper is to replicate the methods used, it's often tempting to skip this section. However, you should remember that if there's one thing that distinguishes science from non-science, it is the methodology. In order to be considered scientific, an investigation must adhere to certain principles. There has to be a logical reason why the method can or may answer the question. Defined and reproducible protocols must be followed. Controls must be in place in order to rule out extraneous influences on the results. A detailed Methods section is not just a conventional but essentially arbitrary part of the paper. Rather, it constitutes the evidence that you are reading a scientific paper and not a work of science fiction.

Even if you have no intention of performing an experiment yourself, you still need to know how it was done in order to understand what it means. You should ask yourself 'how would *I* go about answering this question? Does their method do the job? If the method is not scientifically ideal, why not? Is there a good excuse?'

Ideally a scientific method
Should be:

- Direct
- Robust

“You cannot judge results without
judging methods”

A **direct** method is one that tells us directly what we want to know. A direct method for estimating the density of bacteria in a culture would be to take a sample and count the bacteria. An **indirect** method is one that generates an observation from which we can infer what we want to know. An indirect way to estimate the density of bacteria in a culture would be to measure the optical density of the culture (how cloudy it is) – and use that to calculate the density of the cells.

An indirect method requires deduction or inference, which in turn rely on assumptions, and assumptions may sometimes be wrong. The more assumptions that must be made in converting an observation into a conclusion, the more room there is for error.

A **robust** method is one that tries to eliminate misinterpretations by approaching a problem from several different angles – e.g. running two different experiments that answer the same question.


Finally, you should ask yourself whether the technique has used proper controls for the experiment they are trying to do and the conclusions they draw from that particular experiment.

The following text taken from the University of Greenwich's 'Guide to Critical Reading of Scientific Research Papers' illustrates:

The Essentials

Although the methods and technology used in scientific experiments might vary tremendously, from particle accelerators in physics to chainsaws in ecology, the basic design of an experiment should, where possible, follow a few simple rules. Here is a description of 'the essentials':

Drugs Education - does it work? Many millions of pounds are spent on health education programmes every year. But do these campaigns actually work? The only way to find out is to adopt an experimental approach. Suppose you were asked to determine if a new way of teaching about the dangers of drug abuse to teenagers had an effect. You could set up an experiment as follows:



✖ What is wrong with this design?

Consider how you would interpret the results you might get:

Result	Interpretation
Percentage using drugs goes down	1. The education has been effective. 2. Something else (e.g. a highly publicised death, like that of Leah Betts) has changed behaviour. 3. There is bias e.g. students are lying.
Percentage using drugs goes up	1. Education has encouraged drug use. 2. Education has discouraged drug use, without it, even more students would be using drugs.

So the experiment cannot tell you anything useful, because you could interpret the results in many ways. The first thing you need to do to improve the design is to include a **control** as well as a **treatment** group, that is a group which does not receive the treatment, to act as a comparison.


✖ Write down now four characteristics of a good control group for this experiment

Now consider your interpretation of the possible outcomes:

Result	Interpretation
Drug use in the treatment group goes down, drug use in the control group stays the same	1. The education has been effective. 2. Some other factor has acted to reduce drug use in the treatment, but not control group (e.g. perhaps one member of the treatment group was arrested, thus scaring his friends off drugs).
Drug use in treatment and control groups stays the same	1. The education has been ineffective. 2. Some factor would have caused the treatment group to increase their consumption of drugs, but they did not do so because of the education.

Although the design is now better, it still cannot answer your question. This is partly because the question itself is too broad. It is much harder to answer the question "Does this intervention have an effect on illegal drug use?" than to address the more focused question, "Does this intervention reduce illegal drug use?" In general, it is best to make sure that the hypothesis you are testing (in this case, that this education reduces drug use in teenagers) is as focused as possible.

The other main flaw is the lack of **replication**, both of treatments and controls. How do we know that the groups we have chosen are representative? If the treatment group happened to have a charismatic teacher, drug abuse might reduce because of this, and not because of the particular educational technique of interest. However, if we had ten treatment and control groups, chosen randomly, it is most unlikely that the treatment groups will all happen to have particularly charismatic teachers, and the control group particularly poor ones. So this would be a much better design:



Increasing the number of replicates will make the experiment more convincing, but will increase the costs and work involved. One way round this might be simply to subdivide the treatment and control groups. For example, if all the students in a class (rather than the class itself) were regarded as replicates, a single class would provide perhaps 20 replicates.

✂ Write down what is wrong with this approach.

Finally, a good experimental design will ensure that the replicates (of treatments and controls) are **independent**. Consider running your experiment in a single large school, with five classes as treatments and five as controls. It is likely that students in the control groups will have sisters, brothers and friends in the treatment groups, and will therefore hear from them details of the education that the treatment groups are receiving. This could affect the behaviour of students in the control groups, making the experiment invalid (for example, if the control-group students reduced their drug use because what they heard from their friends was so effective, the experimental results would suggest that the education was *not* effective, since there would be no, or little, difference between control and treatment groups).

Results

This represents a summary and analysis of the data. Usually graphs and tables will be included here, if they make presentation of the data clearer. Note that the results section should simply present the results of the work described, without discussing them.

Discussion

This is the most 'open' section of the paper, where the author(s) draws conclusions from the work described. Have the initial aims of the investigation been achieved? Have the hypotheses of interest been tested? How do the results fit-in with other people's work, and what further work needs to be done?

Discussion sections often contain speculations and generalizations. This is fine, as long as the author(s) is careful to point out what conclusions his research can definitely support, and what is more speculative

5. The Peer Review Process

Once a manuscript has been submitted to a journal, it will be assessed by an editorial board and be published if it is deemed to meet criteria for relevance, interest, originality and so on. Most scholarly journals use the initial editorial review to eliminate a large proportion of manuscripts, and then send promising ones out to be reviewed by the author's peers – other experts in similar fields. Manuscripts are sent to three independent reviewers who return their reports to the editor. Reviewers' comments are passed, usually anonymously to the authors who have a chance to adapt their papers accordingly and to comment on the reviewers' reports.

Not all articles in a journal are peer reviewed. Non-refereed articles may include:

- news
- editorials
- opinion pieces
- commissioned reviews of a field
- letters to the editor (e.g. critiques of other people's work)

Shortcomings of the peer review process

- Reviewing is a chore – reviewers frequently get little or no payment for the review work they are asked to do, and would rather be getting on with their research anyway. So they may not pay such careful attention as they should do
- Reviewing can be difficult – One of the main ideas behind the peer review process is that researchers cannot always spot the flaws, ambiguities or loose ends in their own work. Sometimes the reviewers miss those shortcomings too. Godlee *et al.* (1998) deliberately introduced eight areas of weakness into a manuscript and sent it to reviewers. Each reviewer spotted an average of two problems. Clearly, readers cannot sit back and assume that all the work of critical reading has been done for them!
- Reviewers don't have the last word – The content of a journal is determined by its editors. Reports from reviewers are advisory and editors have no obligation to accept them. As well as being scientific gatekeepers, editors are also intimately involved in the production of a commercial product, and have an interest in publishing exciting or revolutionary papers. Of course, no editor likes to be caught out by accepting a paper which later turns out to be unsound; nevertheless, the advice of cautious referees may not always be given the weight it deserves.

6. Bias

So you've checked the materials and methods carefully, understood the results and drawn your conclusions about the findings of the paper. There are still a few things to consider around bias. The list below gives some common areas, but you may be able to think of more.

- *Authors' backgrounds and points of view influence their writing*
Readers need to familiarise themselves with a field as much as possible before interpreting novel research findings. Try asking the following questions:
 - o Is the paper in line with conventional thinking or does it represent a departure from the accepted view?
 - o Who are the authors – do they have a particular axe to grind?

The process of authorial interpretation is not some sinister plot to distort the truth or to uphold a theory that experiments have falsified. The point is that experimental results often allow more than one interpretation. Any of several interpretations might be equally sound from a scientific point of view, and they may all require giving some results more weight than others. Frequently, the interpretation that upholds a tried and tested theory will be the correct one. Sometimes it will not. But any author's choice of interpretation is bound to be somewhat subjective.

- *The need to attract and maintain funding can influence authors' presentation*
Any grant-giving body wants to know that its money is being well spent and will expect to see publications coming out of the work that they have funded. There may be several angles to a scientist's work, and the one that is stressed may depend on the source of funding that is being approached. The source of funding may then go on to determine the way in which the work is written up and published in journal articles. There may be a rush to publish before the next grant application needs to go in, or a tendency to stress certain aspects of the work at the expense of others. Concerns about funding will affect the text that a reader of paper has to deal with.
- *Commercial publications raise their own questions*
A university scientist who discovers a result with no obvious commercial application is likely to want to publish it sooner rather than later. On the other hand, an industrial research scientist must consider how to protect

intellectual property if it is commercially valuable. This means lodging a patent application, which may take several months, or keeping the invention a secret. Publishing immediately and sharing unprotected results with competitors is ruled out. The commercial sector therefore has a built-in incentive not to publish results. When you find industry papers that have made their way into journals, it makes sense to ask 'why did they publish this now?' Here are some possible answers to that question:

- o The end of a project. If an avenue of research looks exciting and commercially promising, it's a good idea not to tell too many people about it. Once a project has stopped producing sales it starts producing column inches.
 - o The same is true of projects that have made good progress, but that the company has decided to drop for strategic reasons.
 - o Therefore a paper or series of papers on a similar topic from a commercial lab can provide a clue that the project is no longer active.
 - o A paper may also provide a clue that the research is going well and that the company wishes to pique interest in a forthcoming product.
- *The need to publish many papers*
Pressure from Universities and funding organisations, together with personal ambition tends to impel scientists towards publishing as many papers as possible/publishing in the most prestigious journals possible. The desire to get multiple papers out of a single set of results can tempt authors to 'sausage slice' their work into fragments jokingly referred to as MPUs – minimum publishable units. A piece of work that is quite extensive but not good enough for a paper in the most competitive journals may be written up in a second-tier journal – as two papers. Better still, the papers may cite each other, an advantage for authors since on rough-and-ready way of assessing the importance of a paper is to count the number of other papers that cite it as a reference.
- *Ambitious people may exaggerate the significance of their work*
It is important not to judge a paper on how important it would be if it were correct, but rather on how likely it is to be correct given the quality of the work.
- *Papers and editorials are not immune to journalistic rhetoric*
Sometimes editorials get excited by papers, making bold claims about their findings. Careful study of the paper may reveal that these claims were never actually made by the authors of the paper themselves. We

see this kind of thing happening in newspapers, but it may happen in the editorial section of journals too.

- *Publication bias and reference bias*

Negative results do not sell journals unless they contradict a well-established belief or relate to an area of particular interest or controversy. Negative results often seem unexciting to their discoverers too and may not even be submitted for publication.

If I studied the influence of mitochondrial dysfunction on insulin resistance and found no effect, I would not be surprised if my paper were turned down by a journal, and might quite likely not bother trying to get it published in the first place. Yet the consequent invisibility of my results could have profound repercussions. Another scientist – or several groups of scientists might waste time, effort and resources studying the same question and coming to the same negative conclusion. This is an especially serious problem when human subjects, lab animals or large investments are involved.

Alternatively, another scientist might perform a similar experiment and find that mitochondrial dysfunction *does* have a role in insulin resistance. If that experiment is published, then it will become the last word on the subject, unquestioned and unreplicated. What if it flies in the face of a hundred sets of discarded results suggesting the contrary? What if the published result is an error, or applies to only a small subset of patients?

Ben Goldacre, in his book *Bad Science* rails against this trend of non-publishing of negative results. He argues that such a practice frequently leads to the oversight of important information and the arrival at incorrect conclusions particularly in health and medicine.

Even where authors know of results contradicting their findings, they may on occasion conveniently forget to draw attention to them. It is far simpler to be able to say 'eating fish oil is good for you' than to have to say 'our experiment suggests a possible benefit, but the meticulous work of Jones and Doe (2005) failed to find any beneficial effect, so we might be wrong'. Omitting ifs and buts is simple and it increases the apparent impact of a paper, which is useful for authors. If a result appears too tentative, journals may be reluctant to publish it.

Therefore, treat bold claims with a healthy scepticism and keep an eye out for contrary claims.

- *The source of funding influences published research outcomes*

Of course, one of the most significant influences on the answer to a scientific enquiry is the question asked. This is clearly revealed in investigations into how the outcomes of medical research depend on the research funding. These investigations show that research funded by pharmaceutical companies is more likely to produce results that favour the sponsor's product than research with other sources of funding. This is most often about the question asked rather than the funding of shoddy research by the drug companies.

If you are testing the effectiveness of a new drug, there are a variety of questions that you might ask, such as; 'how does the drug compare to a placebo?', or 'how does the drug compare to the best available treatment at the moment?', or 'how does the drug compare to no treatment at all'. As you will already have concluded, testing a drug against a placebo or against no treatment at all is far more likely to give a favourable result than testing it against the best treatment currently available.

Even more dodgy practices may involve comparing the new drug to the best available treatment administered at lower than recommended levels. If you use your imagination, you could probably come up with a variety of tricks you could use to ensure that the test on your drug came out favourably!

So you should always ask who funded the research that you are reading about and what implications this could possibly have.

Summary

How to read a paper

1. Decide why you are reading the paper
2. Identify any possible sources of bias by considering the authors' affiliations and funding sources
3. Given your reason for reading the paper, decide on your reading strategy
4. Read the paper with scepticism, looking out for experimental flaws, alternative interpretations of the work and unsubstantiated claims in the discussion section

Some questions to ask while you are reading a paper

- Why am I reading this paper?
- What is the research question that this paper is trying to answer?
- How would *I* go about answering this research question?
- Are there some confusing findings that don't appear in the Abstract?
- How were the experiments done?
- Are the experiments direct and robust? What compromises have been made?
- Have suitable controls been put into place?
- Do the experiments actually answer the research question?
- Are the experiments representative?
- Is the paper in line with conventional thinking or does it represent a departure from the accepted view?
- Who are the authors? Do they have a particular axe to grind?
- Why has this paper been published now?