Science, truth, and honesty

Andy J. Wills

September 25, 2017

1 Topics

This lecture series starts by considering two fundamental questions — "what is truth?" and "what is science?". We start with the nature of truth.

2 Truth as a property of claims

It's not possible to get anywhere if one considers truth as some kind of general ineffable property of the universe. Truth is, at heart, one property of a claim. Someone makes a claim and someone (the same person, or someone else) attributes the property "true" to it. Or the property "not true" (false). It is also possible to attribute some level of belief, perhaps using probability. Finally, it is possible to "opt out" and say you have insufficient information to make a sensible judgement.

Note that the last two statements on the slide are about the limits of our knowledge, not (necessarily) about the limits of truth. The last two statements are either true or false, it's just you don't know the answer.

3 Subjective and objective claims

Truth is a property of claims, but there are different types of claims — a recurring theme in this lecture. The first distinction we need to make is between subjective and objective claims.

Subjective claims are those whose truth differs for different people. So, for example, the truth of the statement "My favourite colour is blue" differs depending upon who says it. Objective claims are those whose truth value is not affected by who says it. "Charlotte Smith's favourite colour is blue". This is true, or false, irrespective of whether I say it, you say it, or Charlotte Smith says it.

Note that the term "subjective" is amongst the most widely mis-used, particularly when referring to psychology. For example, the claim "I think chocolate tastes better than cabbage" is subjective, but the claim "chocolate tastes better than cabbage" is objective, and is amenable to scientific test (for example, taste ratings taken over a representative sample of individuals).

4 Vague claims are subjective

Critically, a claim is not objective if it is too vague to be tested. For example, "smoking is wrong" - not objective because "wrong" has a number of different meanings.

5 What is truth?

When one asks "what is truth?", one generally is asking what it means to say that an objective claim is true. One possibility is to avoid the question by asserting that there are no objective claims – all claims are subjective and hence the truth value of all claims depends on who is speaking them.

This position about truth is called relativism. It's not an uncommon view, even amongst academics, but it's pretty problematic.

First, any definition of relativism is self-defeating. For example, relativism is the belief that "all truth is relative". This claim is, in itself, universal and so proves that not all truth is relative! The same problem exists for expressions such as "No one can say what is right or wrong".

Another real problem with relativism is that it seems to preclude the possibility of being wrong. If truth is simply what you believe, then your beliefs cannot be in error.

One can avoid these paradoxes by saying, for example "some truth is relative". This seems undeniably true, as we discussed under the issue of subjective claims. It still leaves us with the question - what does it mean to say an objective claim is true?

6 Truth, correspondence, and coherence (2 slides)

The answer is really quite prosaic. Truth can be defined both in terms of correspondence and coherence. Anyone who has ever watched a crime programme on TV knows this.

Correspondence Theory states that Truth is determined by the correspondence between what is claimed and what is observed. An murder suspect claims to have been at home between 10pm and midnight, but he was seen by a reliable witness some miles from his house at that time. The murder suspect's objective claim is judged to be false because it does not correspond to observations.

Coherence Theory states that Truth is determined by the coherence between new claims and other justified claims. The Defence claims that the apparent murder victim actually committed suicide. A number of Prosecution witnesses testify that the victim had no history of depression or other mental illness, and appeared happy in the days and weeks leading up the attack. The Defence's claim seems unlikely to be true, not on the basis of anyone having observed how the injuries were inflicted, but on the lack of coherence between this claim and other justified claims.

7 Scientific claims

Having established what truth is, we can turn to the question of what science is. The question of what science is, is important for two reasons. First, you are studying a science and hence you are expected to operate within the definition of what a scientist is. Second, science is not just a subject of study, it is also one approach to life in general that has much to commend it.

There are two parts to answering this question "what is science" - the nature of scientific claims, and the culture of modern science.

Scientific claims first. At its heart, science is the process of making scientific claims and then attempting to determine whether those claims are true or false. Scientific claims are objective rather than subjective, see earlier. Science also limits, or should limit, the claims that it makes to those whose truth or falsity can, at least in principle, be clearly determined.

In general, this means scientific claims are descriptive rather than prescriptive.

8 Descriptive versus prescriptive claims (3 slides)

Science's focus on descriptive claims does not exempt scientists from having moral or ethical principles, or from wanting to change the world in particular ways. But scientists choose to investigate the issues inherent in prescriptive claims by converting them to descriptive claims that are relevant to the prescriptive point.

For example, a scientist who wishes to make the claim "abortion should be made illegal" looks for an descriptive claim that would support the prescriptive claim. For example, if one wished to support the outlawing of abortion, one might make the claim "feotal pain receptors have developed by eight weeks gestation", in order to support your prescriptive claim. If one wished to argue against that prescriptive claims, you might then make the descriptive claim "foetal pain receptors are not connected to the brain until at least 20 weeks".

Similarly, in order to support the prescriptive claim about drink-driving, you would make some kind of descriptive claim. For example, "Risk of car accidents doubles at 80mg/100ml blood alcohol (UK drink-driving limit)"

So, science is not about avoiding societally difficult questions. It is about making claims that can be reliably examined, in order to help answer difficult questions.

9 Absolute versus contextual claims

To summarize, scientific claims are objective, and they are usually descriptive. However, scientific claims are seldom absolute.

Absolute claims are invariant. They hold always. Their truth value is not conditional on circumstances. They are not conditional on time or place.

Relative claims (perhaps better described as contextual claims) hold under a defined set of conditions. Scientists should define those conditions, but do not always do this. For example, consider the claim, "the leadership positions that women occupy are less promising than those of their male counterparts". This is not intended to be an absolute claim. It is, rather, a claim about a current state of affairs. If this claim is true now, but no longer true in 20 years, perhaps due to increased awareness of the issue, this does not undermine the truth value of the original claim – as long as it is seen in the appropriate context.

In general, scientists wish to investigate claims that are, as far as possible, context independent. In other words, as close to being absolute claims as is achievable. Indeed, it is hard to think of claims as scientific in the broader sense if they are too context-specific. For example:

"The leadership positions that women in 2003 in UK FTSE100 companies occupy are less promising than those of their male counterparts"

Such a claim is technically scientific, but the truth being claimed is so contextual that it is not possible to, for example, make any further predictions.

10 Scientific claims and truth

Scientific claims are objective, descriptive, and relatively context independent. They also have a number of other properties.First, some scientific claims are true and others are false. A claim that is false can still be scientific if it meets the other criteria.

11 Observable, measurable states

Second, scientific claims are based on observable, measurable states. Making sure your claim is descriptive rather than prescriptive helps to make it observable and measurable, but doesn't always get you all the way there. For example, consider the claim "impulsive people are more likely to be criminals". Being a criminal is a state that is observable and measurable in a number of relatively uncontroversial ways. For example, if you have been convicted of a crime then you are a criminal in an observable and measurable way.

Impulsivity is a fairly vague concept that has to be somehow translated into something directly measurable. One of the contributions of modern psychology has been to develop explicit measures that can be clearly assessed for reliability and validity. In the case of impulsivity, there is a 30-item standard questionnaire called the Barratt Impulsivity Scale (BIS). So, our claim becomes:

"People with one or more criminal convictions score higher on the BIS than people without a conviction"

12 Independent replication

Third, scientific claims must be expressed in such a way that they permit independent replication. For example, claiming that Willsian Therapy reduces depression is not a scientific claim if Wills is the only person who can perform Willsian Therapy.

13 Scientific claims

Finally, scientific claims have to be falsifiable. This basically means that it must be possible to imagine an outcome of an experiment or other study that would lead to the conlcusion the claim was false. If this is not possible, the claim is unfalsifiable and hence unscientific. Note the difference between falsifiable (one can imagine a set of circumstances in which it could be shown to be wrong), and falsified (one has demonstrated that it is wrong). Good, productive, science consists of the collection of claims that are falsifiable but not (yet) falsified.

The slide contains a summary of the properties of falsifiable claims.

14 Science and dishonesty (3 slides)

To summarize, science is about making scientific claims, and we've looked at what makes a claim scientific (or unscientific). But science, like all human endeavours, is about more than this. It is also about being part of a culture that holds certain behaviours and processes in high regard. Trangressing these rules leads to the normal cultural sanctions — isolation, exclusion and, sometimes, expulsion. In some cases, these cultural aspects represent A WAY of doing science, rather than THE ONLY WAY it could be done.

The first cultural norm of modern science is honesty. This is most easily defined by its opposite — dishonesty. It is dishonest to partially report your results if your intention is to obscure results that are inconsistent with your claim. It is dishonest to choose a form of data analysis because it gives you the best p-value, rather than because it is the method best suited to your data. It is dishonest (if you are using standard statistical methods) to keep testing until $p < .05$. It is dishonest to publish the same data more than once without acknowledging you have done this. And it is dishonest to say you collected some data when actually you just made it up!

The reason honesty is such an important cultural norm in science is that it gets in the way of reliably evaluating claims, and wastes peoples' time. Science is a highly social cooperative activity.

In the lecture, I considered the cases of Diderik Stapel, and Dirk Smeesters. These are pretty obvious cases of unacceptable dishonesty. I then went on to discuss Simmons et al.'s apparent demonstration that listening to music about old age makes you physically younger, as measured by your date of birth. This is of course impossible — and the authors realise this — but the conclusion comes from data that was really collected, and in which the analyses reported were really conducted, and conducted correctly. The problem is that, for the purposes of illustration, the authors engaged in a number of other practices I listed above as dishonest (first three items). Part of the problem is that some psychologists do not consider these practices dishonest — but this is changing.

15 Science and skepticism (4 slides)

Another important cultural norm in science is skepticism. Scientists should welcome it when others are skeptical about their claims, because such skepticism can lead to close examination of the evidence, which in turn advances our understanding. Essentially, being scientific is about soliciting and welcoming criticism. This is something that many of us, including me, still find emotionally quite difficult.

One recent development in skepticism is the increasing emphasis on not only scientific claims permitting independent replication, but the importance that such independent replication is actually conducted. Recent meta-analyses suggest that a half to two-thirds of results in psychology fail to replicate. In the lecture I discussed two case studies (Ap Dijksterhuis, John Bargh).

I deliberately picked social psychology examples, because it is here the results are worst. In social psychology, you should assume it won't replicate unless there is evidence that it does. In cognitive psychology, it's 50:50 — your first question should still be "has this been independently repilcated?" Ask this in all your lectures, all your tutorials, all your essays. If it hasn't been independently replicated, don't believe it!

16 Science and publication

The next cultural norm is publication. Science is highly social and cooperative. Collecting data, and evaluating claims, and then keeping the results of that evaluation to yourself is generally frowned upon. Some scientists would even go so far as to say it is unethical to not publish your work, although this is a minority view.

17 Science and Open Access publication

There are two important recent developments that have emerged from this belief in the importance of publication.

The first is a realization that the traditional method of academic publishing is getting in the way of the reasons we should be publishing in the first place. Generally, scientists publish in scientific journals. Subscriptions to these journals cost thousands of pounds a year, and are generally only held in university libraries, which are often not open to the public.

The solution is Open Access publishing. This means making your publications freely available at no charge. There are a number of ways to achieve this. For example, every scientific paper I have ever written is available for free download at `www.willslab.org.uk`.

18 Science and Open Access raw data

Another recent development related to publication is the realization that the internet now permits scientists to make not only their publications, but also their raw data, freely, publicly, and PERMANENTLY available. This is really incredibly important. It is one thing to run a study again because you want to independently replicate it. It is quite another to HAVE to run it again because the original author didn't analyze the data properly and the raw data no longer exists.

19 Science and peer review

A third, critical, aspect of the culture of modern science is the process of peer review. This refers to the fact that your publications should be scrutinized by your peers — people who are experts in the particular sub-field in which you are working. Peer review, despite some criticism, remains a critical part of the scientific process. Generally speaking, scientists (and others) would like to be able to assume that publications have passed some basic quality control checks. Authors would also like, in most cases, others to believe what they are saying, rather than discard it because the work contains some obvious flaw or is expressed unclearly. Peer review can help with this.

20 Peer-review process (2 slides)

Traditionally, peer review proceeds as follows. The scientist writes a paper and submits it to a journal. After some time (3-4 months in psychology), the paper is returned with reviews from 2-3 experts in the field. The journal editor tells the author either that the paper has been accepted as is (very unlikely), or that it might be accepted if the reviewers' comments can be addressed, or that they are not interested – please don't send it again. Resubmission of a revised paper is typically followed by another set of reviews. Eventually, if the editor agrees, the paper is published. The process can sometimes take more than 2 years.

Peer review is traditionally anonymous and closed. It's anonymous in the sense that you don't know the identity of the people reviewing your work. It's closed in the sense that the reviews you receive are confidential and cannot be published.

A hot topic at the moment is whether anonymous closed review tends to lead to unscientific behaviour. As a reviewer, your identity is a secret and your review will never be published. It is human nature to be more critical of work that is counter to your own views than work that is consistent with it. The cloak of anonymity and confidentiality seems likely to make this worse. If reviews were published, along with the reviewers' name, at the same time as the paper, reviewers would presumably be rather more careful to limit their review to objective unbiased comments.

The argument on the other side is that open named reviews would be largely useless because people would be unwilling to give them, or would not give their honest opinion for fear of retaliation. The presence of good journals that use open named reviews as part of their process suggests this fear is unfounded. For example, Behavioral and Brain Sciences

publishes papers alongside several invited open peer reviews, and an authors' response to those reviews. BBS also runs a closed review process before this as a quality-control check.

21 Culture of modern science (and future directions)

This is the final slide of the lecture, and summarizes the culture of modern science.