

Reliability Trumps Truth

NANCY CARTWRIGHT, Department of Philosophy, Durham University, United Kingdom, Department of Philosophy, University of California at San Diego, USA and Research Associate, Department of Philosophy, University of Johannesburg South Africa

CARTWRIGHT, N.: Reliability Trumps Truth
FILOZOFIA, 79, 2024, No x, pp. xx – xx

Abstract

This paper argues that reliability is a better desideratum for scientific claims than is truth as soon as we focus on ‘helping to change the world (for the better)’ as a central aim of science.

Keywords: truth, reliability, warrant, aims of science

This paper argues that reliability trumps truth. What I mean by that in short is that *reliability is a better desideratum for science claims than is truth*.

Why? My late husband, Stuart Hampshire, pointed out that people are naturally curious: we want to *know* and we want to know *what is true*. That is an important part of human nature. So there is good reason to try to learn truths in science – because we naturally want to know. But it’s not good enough to act as if true scientific claims are an end in themselves. We also very much want claims in order to *use* them. That’s why I’m interested in science: because I have seen that it has proven so useful. We want to use the scientific claims we ‘establish’ to draw further inferences to change the world. This has two important knock-on consequences:

- If claims are to be useful, these claims need to be able to do what we expect of them.
- And if we’re going to be justified in using them, we need to be warranted that they are able to do what we expect.

That is, *these claims need to be reliable and we need warrant that they are reliable*.

My plan here is to lay out three advantages I think reliability has over the truth. First, reliability forces the question, “Reliable for what and on what

grounds can we defend that?" Second, reliability as opposed to truth does not invite *detaching* claims from the tangle of work that supports them. (I'll talk about what that means soon). Third, thinking in terms of the uses to which our claims are to be put and how reliable those uses are helps solve a standing problem in science: how to report study results. I start with detachment.

I. Detachment

Confirmation and truth invite what philosopher of social science Sharon Crasnow calls *detachment* of claims from the tangle of work that supports them, that constrains what you can expect them reliably to do and that warrants what you can expect them reliably to do. Detaching is a mistake. Focusing on the reliability of claims discourages making this mistake.

Here is how Crasnow explains detachment and the problems it can generate (which she discusses with respect to causal claims but her points are in no way restricted to causal claims):

Under the detachability ideal, the goal of research is to come up with a set of causal claims that once warranted can be separated from the context in which they received that warrant.

Thinking of causal inference in this way ... suppresses the relevance of background knowledge (beliefs and assumptions that reveal the relevance of particular information as evidence) and of aims (including practical implications) – the elements of context.

The ideal is intertwined with an understanding of knowledge as universal, general, and perhaps even lawlike. The transportability of the causal claim – where else it might be applicable – is thought to depend on these characteristics.

[My] alternative view calls for continued attention to context and argues for transportability to depend on detailed comparisons of contextual features (Crasnow 2024, 252).

With a number of other authors I recently published a book called *The Tangle of Science: Reliability beyond Method, Rigour and Objectivity*. (The other authors are Jeremy Hardie, Eleonora Montuschi, Matthew Soleiman and Ann Thresher.) The *Tangle* might seem in favour of detachment because it makes much of the fact that the sciences regularly put a myriad of products 'on the shelf' for others to take down and use for themselves. In the *Tangle* we are interested in the reliability of the vast panoply of kinds of products that science

produces: models, concepts, measure, devices, et cetera. This was meant to include scientific claims, though we did not discuss those in particular, as I do here. We say

Science's products [which includes claims] have been envisioned, developed, created, assembled and tested by one conglomerate of actors and then they're put on the shelf [i.e., they're deemed acceptable and propagated] to be taken down for use by different actors with different ideas and practices for employing them in their own endeavours (Cartwright et al. 2022, 3).

We use the term *products* of science deliberately because these *are* products and also because it resonates with other facts that matter. As we note,

Many products on shop shelves come with instructions for use: "You can bake good cakes with this soft flour but not good pasta"; "Take 2 tablets twice a day, best with food"; "Store in a cool dry place"; or "Use [Genie Crafts Cloth Plaster Wrap Rolls] to make keepsakes, pregnancy belly casting, masks, science projects, 3D sculptures, science projects, etc". But new users find [and share] new uses ... And there are general shared understandings: you don't expect your mobile phone to work in the middle of a cave or if you drop it in the bath (Cartwright et al. 2022, 3).

The same is true, we argue, for the products of science, including claims. There are only a few explicit users' manuals, but there is a very great deal of implicit and shared understanding. Scientists learn what their products can and cannot be used for in combination with others. Here's an example we quote in the book about the use of continuum versus particulate models for flow (quoting Gollub 2003, 10 (not on in the Bibliography):

Some researchers have used a continuum approximation to obtain a version of hydrodynamics for granular media. The strategy is sometimes useful – for example, in treating energy transport in highly excited granular matter. However, the continuum approach can be problematic. When the degree of excitation is low, gravity often produces persistent contacts between particles, so their interactions are not limited to the isolated collisions that are assumed in justifying a continuum approximation.

Also, if you shake a granular medium to excite it, the energy subsequently decays rapidly but unevenly due to huge numbers of inelastic collisions. This collision process leads to inhomogeneities in the local kinetic energy on scales only a few times larger than the particle size. The distribution of

internal stresses in both the static and dynamic states is highly nonuniform, and the stresses are transmitted along linear chains of particles, in contrast to the situation in ordinary solids (Cartwright et al. 2022, 4).

Returning to detachment. Detachment is removing the labels from these products, the labels that indicate

- Where and how the product was developed: e.g. derived from an idealized economic model.
- What to do with it, how to use it and how not to use it: e.g. don't use it to make inferences about settings where the central assumptions of the model in which it is derived are violated.

Think about how the process of 'establishing' scientific claims (in our terms, putting them on the shelf for other users to take away and apply elsewhere) and then using these scientific claims proceeds. Suppose you have one body of work $W1$ that leads you to take a claim φ to be established. But there might be another body of work $W2$ that does the same for φ and another ($W3$) and another ($W4$). When φ is 'established', given other premises, you then make inferences from it. If when φ is established – i.e. made publicly available for use outside the context in which it was developed and defended – it is detached, it is then treated as a true claim that you can use with other premises in any inference *regardless of which body of work originally led to deeming it true*. Imagine there are 6 or 7 obvious conclusions you can draw assuming the claim φ is true: $C1, C2, \dots, C7$. The nice thing about detachment is that you are entitled to each and all of these (supposing the other premises necessary are detached as well) no matter which body of work is taken to support the claim. Essentially you can go from ' $W1$ or $W2$ or $W3$ or $W4$ ' to ' $C1$ and $C2$...and $C7$ '.

An easy way to see this is to think about claims about measurement results. You can measure temperature in a variety of different ways that are taken to provide accurate readings. Having then measured the temperature of an object in one of these ways or another, you can go on and make all the standard inferences about things that have that temperature. If you have a body of warrant that looks good enough to detach a claim and treat it as well-formed and true, then you can draw any and all conclusions that follow from valid inferences using that claim as a premise. That's what happens under detachment.

Without detachment you have a lot more work to do, work which Sharon Crasnow and I both think needs to be done much of the time.¹ For any specific conclusion you have to argue that it is supported by some specific body of work. You cannot just derive the conclusion from some off-the-shelf claim used as a premise. You must show how some body of work that supports making that claim available for use elsewhere supports the specific use you put it to in drawing any further conclusion. You're not allowed just to say, "Now, this is true and I'm going to just use it to draw my conclusion." Without detachment, you can't do that.

II. Reliable for What and How Reliable?

'True' seems to be a one-place relation. 'Reliable' is clearly (at least) a two-place relation. *X* is reliable *for doing Y*. Moreover, reliability is not an all or nothing matter. You might judge it to come in degrees, like '*X* is highly reliable, fairly reliable, somewhat reliable for doing *Y*'. I don't know if you think you have some way to lay down a measure, but if so you might even judge something like '*X* is 70% reliable for doing *Y* in cases of type *Z*'. Or you might allow more modal possibilities, which I think is important since this is actually the way we think, even if we don't say things exactly this way in science: '*X* might be reliable, it may be reliable, it might well be reliable, or it could be reliable for doing *Y*'.

I claim that different ways of formulating a scientific result suggest different options with respect to either degree of reliability or the modal possibilities. But I'll come to that later. The immediate topic is that reliability is a two-place relation and truth is not. And, like truth, reliability for use should be warranted.

But how do we use scientific claims? What category of things are they useful for? Of course we use claims to affect the world. But for simplicity of presentation here, I will suppose that all uses are inferences. For example, you use the law of the lever to move your car out of the mud. But we might say instead that what you do is to infer the claim that if you put this board over the stone and wedge it under the car, the car will move. So just for simplicity, let's suppose we're not using claims directly to intervene in the world, but we're using them to make inferences that describe what happens if we intervene.

¹ It is no surprise that she and I have similar views on this since we have been discussing these issues together for a long time.

What I want to stress about making inferences from claims is that first, a claim is reliable for making some inferences and not for making others, and second that its reliability for so doing must be warranted if you are to be justified in drawing those inferences. Of course few claims support interesting inferences by themselves. We know that you need other premises. A more accurate way to put the question then is this: is this claim warranted as reliable for use in this argument by the body of work that leads us to accept this claim plus the bodies of work that support the other premises in the argument? Clearly a claim ϕ is not reliable for using in an argument if all the other premises in the argument are rubbish. But again for simplicity let's think about it like this: is this claim warranted as reliable for use in this argument by the body of work that leads us to accept it?

What we did in the *Tangle* book differs from what Sharon Crasnow says since we do in a sense allow, even endorse, detaching. After all, as I noted, we talk about creating the product in one area, then putting it on the shelf for others to take down and use in another area. *But*, then we make a big issue about the fact that the products that are put on the shelf come with labels. Though they are often implicit, the cautions on these labels are standardly recognized in the scientific of the communities that use them – and they need to be so recognized.

The labels discussed in the *Tangle* identify what I call 'bog-standard' uses that the claims are reliable or not reliable for: 'Just observe all these warnings and then it's pretty clear you can use it for this and not for that'. The labels mark out well-established facts about reliability that are generally known in the relevant scientific communities, facts that are backed up by the methods by which the claim was warranted in the first place. For instance here is a standard one: we all know that the claims in Newtonian mechanics come with a warning label 'Do not use where velocities are near that of light or where general relativistic effects matter'. My claim then is that once you take note of the labels, there's going to be a set of bog-standard uses that you are now warranted in taking the claim to be reliable for. But then other uses than the bog-standard ones can be *suggested* by the claim (given the other premises you've got around), and especially in new contexts or when needs must.

So the claim will suggest uses that it hasn't already been established as reliable for. You don't need to warrant the bog-standard uses case-by-case but the further suggested uses then need to be warranted case-by-case. And not doing so can us lead badly astray. Here's one example from political science: the democratic peace principle, which is the principle that democracies don't

go to war with other democracies. Studying it has been a huge industry in political science. It's something that five or six years ago I was told, "This is one of the few things almost all political scientists agree on." I don't know whether that's really the case, but the principle clearly has been an excuse for much American foreign policy.

I am going to make pastiche use of this example to illustrate the dangers of detaching. Here is Sharon Crasnow again:

Surveys indicate many Chinese perceive their form of government as democratic, and even more surprisingly, do not take elections to be a key element. The understanding of democracy underlying these views is that a regime is democratic but governs in a way that is consistent with the well-being of the people (Crasnow 2021, 1220).

Consider now the principle that democracies don't go to war with other democracies. Given the understanding of democracy Crasnow describes, which is held by a very great many people, can we now assume that the democratic peace principle allows us to infer that China will not go to war with Taiwan? I think that's a bad bet. That's detaching the principle from the body of work that supports it. Hardly anything in that body of work bears on that sense of democracy.

Shall we also conclude that the democratic peace principle tells you that France and the UK will not engage in the Scallop Wars? Again, that's a bad bet. As CNN on 20 August 2018 reports, 'French and British fishermen clashed in the English Channel Tuesday in the latest installment of the long-running "scallop war"' – and note that they are 'long-running'!

If you had detached the democratic peace principle and didn't pay any attention to the warning labels, you might feel entitled to bad inferences like those. But more positively, we think that there are clearly many uses of the democratic peace principle that are warranted by its huge tangle of support. We studied this work in political science for the *Tangle* book, looking at the big tangle of support for it and that tangle gets political scientists to so widely endorse it. We ended up with a slightly different interpretation of how to think about it than is usual in political science, partly because we come to it with the idea that some uses are pre-warranted, some uses are fairly clearly pre-discouraged and others are suggested – not ruled out but needing further warrant.

Bog-standard uses of that principle that I think are clearly warranted are ones that look like this: from the DPP you can infer if X and Y are democracies

by anybody's book, X and Y will not go to war-by-anybody'd book. There are a very great many contested different definitions of democracy and most of the conventional democracies that you'd like to label that way, like the UK and France, satisfy all the definitions. These are democracies by anybody's book. So on our reading of the background work in political science, it is reasonable to conclude that the UK and France will not engage in what counts as war-by-anybody's-book. As with 'democracy', there are a host of contested definitions of 'war'. But the Scallop War is not a war by anybody's book. It satisfies some definitions of war, but not many.

Example two comes from finance. It is a standard case we all know about: 'The mathematical equation that caused the banks to crash', the Black-Scholes formula. The quote is from a headline in the *Guardian* newspaper, which continues this way:

The Black-Scholes equation was the mathematical justification for the trading that plunged the world's banks into catastrophe...The equation provides a systematic way to calculate the value of an option before it matures, then the option can be sold at any time. The equation was so effective that it won ... the 1997 Nobel Prize in economics (*The Guardian*, 11 February, 2012).

The problem was that Black-Scholes got detached from the model that underwrites it and put to uses for which it was not reliable.

My colleague John Pemberton has prepared a description of these matters specifically for me to use here.² Pemberton is now a philosopher, but he used to be an auctions-pricer and criticized the use of Black-Scholes well before the crash. That's in fact how he started to become a philosopher. He came to the LSE and joined our modelling group to work on understanding how models are properly and improperly used. Here's what Pemberton says:

- [T]he Black-Scholes option pricing model: the value of the option is calculated using the Black-Scholes formula. Input parameters include ... Plugging the parameter values into the Black-Scholes formula then yields - hey presto! - the value of the option. No need for consideration of other context. What's not to like?
- Unfortunately the use of such models is highly unreliable. The Black-Scholes model implicitly assumes that the distribution of the price of the underlying at the end of each given future period is lognormal

² Personal communication.

Often, in many unexceptional prevailing contexts... this price distribution assumption may be roughly reasonable – so the formula as used may provide a value which is roughly correct.

- Sometimes, though, the shape of the future price distribution is not at all close to lognormal – now the Black-Scholes model offers no approximation at all.

Example three is from education, sent to me by my Durham colleague Adrian Simpson.³ Simpson talks about how results of a study of a reading programme were reported:

An RCT of the “abracadabra” reading intervention finds a statistically significant difference between groups... The effect is attributed to the “abracadabra” curriculum *rather than* the difference between the “abracadabra” curriculum and the curriculum used by the comparison group. [Italics added]

Suppose you detach the result claim as recorded. (I think indeed the report sounds detached.) Then if you think your pupils are much the same in reading ability and preparation as those in the RCT, and their basic socioeconomic backgrounds are also similar, then it seems you can infer

Probably if we introduce Abracadabra, we too will see a similar effect size. But that would be a big mistake if your pupils already have a far stronger reading program in place than the control in the RCT. I hope these three kinds of examples provide a good strong sense of why we consider not just the truth of our scientific claims but crucially what uses our scientific claims can reliably be put to.

III. Choosing among Result Reports

I will assume here that we are trying to figure out how to choose among true claims that you could write down as reports of the results of your study. Now, how to describe results is a problem we philosophers know really well. First, there’s what I call ‘the Donald Davidson problem’ that the event described column 4, paragraph 6, page 1 in the *New York Times* caused the event described in column 2, paragraph 2, page 6. Let’s suppose that’s true. Davidson notes, that sort of claim is not very interesting. What we really want is a claim that picks out the casually relevant features.

³ Personal communication.

There is also the so-called 'new problem of induction', which isn't so new because it's one scientists have struggled with since Aristotle. Nelson Goodman brought this problem to the fore by asking: do we say that all the emeralds we have observed are green, or do we say that all the emeralds we have observed are grue? They're both true. One suggests one kind of inference for the future, the other suggests a different one. This problem is just everyday business for philosophers. Again I shall illustrate with a couple, again oversimplified, real examples from science.

Here is a Covid case sent to me by Jonathan Fuller from the History and Philosophy of Science Department at Pittsburgh:⁴

Original mRNA COVID vaccine trials had short spacing interval between dose 1 and dose 2. The UK & Canada increased the spacing interval in their 2021 rollout. The US FDA/CDC explained why they would not do so for lack of evidence. Whether the RCTs provided direct evidence for the increased dosing interval depends on how the trial intervention is described:

As '2 doses given (approximately) 3-4 weeks apart' vs.

'2 doses given at least 3-4 weeks apart'.

You could describe the results either way. Both descriptions are true of what happened. What to do? Which of these is better to report? Well, if you favour detachment, you want to get the claim out there and not have too many strings attached, allowing people to just start making inferences from it without having to look back, then what you do is go small. That's because the more precise, detailed and unambiguous a claim is, the less the inferences licensed from it are hostage to the body of warrant supporting it, basically because you've stuck all the caveats into the claim itself.

The trouble with that is that precision and usefulness generally go in opposite directions. Claims suggest uses given their formulation and further often suggest how reliable those uses might be. After all, you should expect that. Claims are pieces of ordinary language, and they carry with them what you might think of as conversational implicatures. You put it this way, you suggest this. Put it that way, you don't. The more precise and unambiguous, the fewer uses are suggested. That's generally the trade-off.

Let's look at the COVID case again. We have these two descriptions. Of course as before we need some other premises to make inferences. Suppose it is justified to assume that the study population is representative of the target,

⁴ Personal communication.

maybe because it was a very, very large random sample from the target population. Well, with that assumption, consider the first way of expressing the result: '2 doses given (approximately) 3-4 weeks apart'. This implies you can reliably infer that giving doses three to four weeks apart in the target will be effective. The way it's put suggests not much else.

Now consider the second way of describing the result: '2 doses given at least 3-4 weeks apart'. This implies the same as the first and it *suggests* that it *may* be reliable to extend the dosage gap further still. In fact, it seems to me that it suggests that it's *probably* safe to extend it four to five weeks and *may well* be safe to extend it five to six weeks.

The suggestions are different between the two just from the choice of the way the claim is expressed. The first may be over-precise, discouraging possibly promising investigations. The second may be misleading, suggesting fruitless investigations. Note that I say suggests. The point is that if an inference is suggested, it's not what I called a 'bog-standard' use. You can't just take it to be reliable. You need to do a lot of new work to warrant it.

Why allow for suggestions? Because that a conclusion is suggested matters if you care about the conclusion. The formulation of the claims invites you to consider whether there is case-specific warrant for various further claims. It tells you, 'You might want to look here, or look there'. So the formulation of result claims and principles is important.

Also I want to remind you, you can never be totally unambiguous. Consider again, Adrian Simpson's Abracadabra RCT. He's thought about the case seriously in the context and thinks that the report that was written was misleading since it seems to allow inferences like the one I showed you. So he wants the result report to talk about Abracadabra *compared to* the particular other reading programme the control group had. He wants you to disambiguate. Perhaps that suggests that what you should do is just disambiguate, disambiguate, disambiguate. But to disambiguate there's a lot more you need to specify precisely, including

- The description of the treatment
- The description of the effect
- The description of the control
- The method of measuring the treatment
- The method of measuring the control
- The method of measuring the effect
- The kind of statistical analysis applied

• ...

Now, if you've done even only just that – which is not actually disambiguating it all the way since you can never get down all the details – what you get is a result claim that suggests almost nothing. That in fact might be the right thing to do. You might be in a position where your background information suggests we don't really know much at all. So you don't want to suggest anything except (again supposing the test population is a random sample from the target) that if you did it exactly this way on the target, you would get the same results.

Choice here is unavoidable. Just remember the Donald Davidson problem. Also bear in mind that you can never be entirely precise. You have to decide. And there's no rule for how to do so. Deciding is a matter of judgment, informed judgment of course, informed by the tangle of work that supports the claim. Are there any guidelines that can be offered? That's really a topic for another paper – or a serious research programme. Recall that I introduce these issues to show up some advantages that reliability has over truth as a desideratum for scientific claims. These concerns about how to express result reports – that we philosophers do not seem to have noticed and that are clearly of real practical importance – set issues of reliability right in Centre Court. So that is enough I think for my purposes here.

But, off the main topic, to get the ball rolling that I offer some very simple starting points that seem to make good sense and fit with the examples we've looked at.

1. Formulate claims so that *only* inferences that the tangle of work underlying the claim supports as reliable are clearly marked out as bog-standard. This is in aid of reliability. For instance, given Simpson's worry, you should include the description of the control program in the report about Abracadabra to foreclose the inference that effect sizes may be similar regardless of how effective the reading programme already in place is.
2. Formulate claims so that *all* inferences that the tangle of work underlying the claim supports as reliable are clearly marked out as bog-standard. This is in aid of making claims as useful as possible. For instance, it could be the case that background work supports the assumption that only if you measure reading scores in just the way they were measured in the experiment could you expect comparable outcomes in your target. But instead you might have a lot of background work that supports the assumption that actually it didn't

matter how you measured reading scores in the experiment. There's good reason to think any standard procedure would have yielded the same outcomes. If so, including measurement methods in the results statement will be overly prohibitive, stopping inferences that can be expected to be reliable.

3. Formulate claims so that promising inferences that are neither clearly admissible nor clearly to be prohibited are clearly marked out. This is in aid of more efficiency about what projects are pursuit worthy. We saw an example of this in the COVID report formulation 'doses given at least 3-4 weeks apart'.

It looks to me as if these issues are reflected in what we actually do in social science reports. There's generally a big methods section that says as precisely as possible what was done in the research. In the results section you see claims that are considerably less detailed. Presumably this is based on serious scientific consideration of what might matter. I hazard that what's happening is that a balance is being struck between being a little precise so that you're pretty sure that anything that looks to be licensed or prohibited actually is so and being too generous and suggesting a lot of red herrings.

Return now to truth and reliability. How did I begin to think about all this and come up with these issues of how to formulate results claims that I haven't heard philosophers talk about before? Recall, they're generated from the supposition that scientific result claims are meant not to be true or not *just* to be true. They are (at least equally!) meant to be useful and reliable for the uses they license.

That's where the argument starts. It's because of thinking of it this way that you can notice the problem and have some starting ideas about a helpful way to resolve it. Of course, in the end this will have to be case-by-case and of course what should be said depends entirely on the tangle of work available. For a little while I was troubled by the fact that in writing this I had no idea which of those results should have been reported in the COVID case. Then I realised, 'That's okay. Why should I know? They make different suggestions about what further inferences should and should not be made and what is promising and what is not. And, not surprisingly, I don't have anything like the background knowledge to have an idea about how to make that judgment'. I did though have views about the democratic peace principle because we went into it in some depth in our background work for the *Tangle* book.

Finally, let me summarise what I hope to have shown. Which is just what I started out with. True claims are not just an end in themselves. We want claims in order to use them. So they need to be able to do what we expect of them: they need to be reliable. But also, they need to be *warranted* to be able to do what we expect of them: they need to be warranted to be reliable. So reliability and its warrant matter.

Why though say that it *trumps* truth? Because there are several advantages of reliability over truth. First, reliability forces the question, reliable for what? What are we going to do with this? I talked about models earlier in the discussion. Models can be useful and help you derive results that are reliable for some purposes and give totally unreliable results for others. It's the same with claims. You need to know, what's it reliable for? Second, reliability does not invite detaching claims from the tangle that supports them. Third, thinking in terms of uses and reliability helps you recognise and deal with a standing problem that philosophy has had little to say about, how to formulate your study results.

So, I think this shows that reliability is a better desideratum for science claims than truth. Reliability turns truth.

Bibliography

CARTWRIGHT, N. – HARDIE, J. – MONTUSCHI, E. – SOLEIMAN, M. – Thresher, A. (2022): *The Tangle of Science: Reliability beyond Method, Rigour, and Objectivity*. Oxford: Oxford University Press.

CRASNOW, S. (2021): Coherence Objectivity and Measurement: The Example of Democracy. *Synthese*, 199, 1207 – 1229.

CRASNOW, S. (2024): Evidential Partnerships and Multi-method Research in Political Science: Methodological, Evidential, and Causal Pluralisms. In: *Philosophical Foundations of Mixed Methods Research: Dialogues Between Researchers and Philosophers*. Ed. by Yafeng Shan. London and New York: Routledge, 240 – 261.

GOLLUB, J.P. (2003): Discrete and Continuum Descriptions of Matter. *Physics Today*, 56/1.

Nancy Cartwright
Centre for Humanities Engaging Science and Society
Durham University
Stockton Road
Durham
DH1 3LE
United Kingdom

e-mail: nancy.cartwright@durham.ac.uk

Department of Philosophy
University of California, San Diego
La Jolla, California
United States
e-mail: ncartwright@ucsd.edu
ORCID ID: 0000-0002-0873-6966



Citation on deposit: Cartwright, N. (in press).
Reliability Trumps Truth. Filozofia

**For final citation and metadata, visit Durham
Research Online URL:** [https://durham-
repository.worktribe.com/output/3105252](https://durham-repository.worktribe.com/output/3105252)

Copyright statement: This accepted manuscript is licensed under the Creative
Commons Attribution 4.0 licence.

<https://creativecommons.org/licenses/by/4.0/>