Robert E. Goodin

The Epistemic Benefits of Multiple Biased Observers

ABSTRACT

We know that we can learn much from the reports of multiple competent, independent, unbiased observers. There are also things we can learn from the reports of competent but biased observers. Specifically, when reports go against the grain of an agent's known biases, we can be relatively confident in the veracity of those reports. Triangulating on the truth via that mechanism requires a multiplicity of observers with distinct biases, each of whose reports might be one-way decisive in that fashion. It also presupposes that all observers share the same fundamental epistemic standards.

We know full well the epistemic benefits of receiving reports from a multiplicity of competent, independent, unbiased observers. Reports from competent and unbiased observers command credence. The more such independent reports we have, the more confident we are in the judgments we form on the basis of them.

Where competent, independent, unbiased observers are concerned, it is not so much diversity as independence that matters for epistemic confidence. If many genuinely independent observers speak with a single voice, all saying exactly the same thing, then we ought to be all the more confident in the veracity of what they say. And when it comes to competent, independent, unbiased observers, we may not need any great multiplicity of them. We can often be really very confident in the conclusions we form on the basis a few dozen reports, as standard sorts of Bayesian or Condorcet Jury Theorem calculations clearly demonstrate.

In this article I want to discuss one way in which analogous epistemic benefits might accrue from reports from biased observers. For those benefits to accrue, those biases must take the very particular form that I shall discuss. In addition to that, there must be suitably many such observers and there must be a suitable diversity of them. Specifically, there must be at least as many (and ideally several times as many) different biased observers as there are distinct biases afoot in the community.

The argument I develop about the epistemic benefits of biased observers thus departs from the standard story about the epistemic benefits of unbiased observers in two crucial respects. First, ordinary Bayesian/Condorcetian accounts of the epistemic benefits of unbiased observers set no particular store by the 'diversity' (as distinct from 'independence') of observers. In contrast, my account of the epistemic benefits of biased observers sets great store by the 'diversity' among them (as well

https://doi.org/10.3366/epi.2006.3.3.166 Published online by Cambridge University Press

as the 'independence' of each set of biases from one another). Second, the ordinary stories about the epistemic benefits of unbiased observers do not require any very large number of them. My account of the epistemic benefits of biased observers, in contrast, attaches real importance to the number of such reports, which must be at least as great as the number of distinct 'biases' represented among our biased observers.

HOW TO TAKE ADVANTAGE OF BIASES

We are ordinarily biased against biased observers, for perfectly good reasons. When we learn that some scientist's research has been funded by the tobacco industry, we rightly discount any findings emerging out of that scientist's laboratory that tobacco smoke is not carcinogenic. Or when we learn that some scientist's father is president of Monsanto, we rightly discount any findings emerging out of that scientist's laboratory that genetically-modified crops will not cross-pollinate wild species.

What we rightly discount, notice, are reports that 'go *with* the grain' of the biases (or, less pejoratively, 'analytic predispositions') that we expect to dominate the work of those particular scientists. When scientists who we have independent grounds for supposing are biased in favour of the tobacco industry report some finding favourable to the interests of the tobacco industry, we rightly reckon that that is basically a worthless piece of additional information. The conditional probability we assign to the proposition being true, given that *they* say it is true, is basically just the same as the probability we would assign to the proposition being true without hearing their views on the matter.

In other words, when we believe that someone's reports are likely to be biased, we rightly engage in one-way discounting of their reports. If they report 'with the grain' of their presumptive bias/predisposition, we set no great epistemic store by that report.

But of course there is another side to the story. Suppose that someone who was in the pay of the tobacco industry reported, quite against the grain of ordinary expectations, that tobacco smoke was indeed carcinogenic. Or suppose that someone who was in the pay of a firm with heavy investments in genetically-modified foods reported, quite against the grain of our ordinary expectations, that the genetic modifications do indeed cross-pollinate wild species, I think that we would – and should – be mightily impressed by *those* sorts of reports, and all the more so because of the presumptive biases of the people making them.

Why so? The basic thought is just this. We all know that there are a great many more-or-less equally legitimate ways of setting up our experiments, and a great many experimental results we might more-or-less equally legitimately report deriving from those different set-ups.[1] By the same token, there are a great many auxiliary hypotheses that we might reject in order to protect our core hypotheses from empirical disconfirmation.[2] Combine all that with the hypothesized disposition of the putatively 'biased' observer to choose among them to report the result among those that is most consonant with her bias. If someone who is looking for any way she can to report a result that favours the tobacco industry, and finds none (reporting instead

https://doi.org/10.3366/epi.2006.3.3.166 Published online by Cambridge University Press

that 'tobacco smoke is carcinogenic'), then that surely counts as strong evidence that it is.

Given what we know (or confidently suppose we know) about *that* researcher's biases, if she tells us that tobacco smoke is carcinogenic, we can be pretty sure that she has explored all the alternative ways of setting up and interpreting the experiments, and she has been forced kicking and screaming to that conclusion. Coming from her, it really *must* be true. Or anyway, in light of what we know about her predispositions and the 'extra steps' that those lead her to take to check out the results that she is now reporting 'against her grain', we attach a very much higher conditional probability to that report being true, given she was the one reporting it, than we would to the same sort of report coming from someone from whom that sort of conclusion is 'only to be expected'.[3]

SOURCES OF BIAS

'Bias' of the particular sort I am discussing here crucially involves a differential propensity toward 'false positives' and aversion to 'false negatives'. A biased person's reports are expected to contain a great many 'false positives' – false reports in support of that person's bias. But the biased person's reports are expected to contain very few 'false negatives' – false reports that go against that person's bias. That expectation of '(almost) no false negatives' is what leads us to attach a particularly high credence to any 'negative' reports (reports contradicting what we expect to be her bias) coming from biased observers – however little credence we attach to any 'positive' reports coming from them that purport to support their bias.

Why should we expect some – perhaps many, perhaps most – observers to be 'biased' in this way? Well, the answer lies in facts about the sociology of science. That is a subject on which I can claim no special expertise. Still, casual empiricism suggests that, both in terms of the trajectory of individual careers and in terms of the reputational dynamics of research teams and laboratories, there are strong incentives that lead to some such phenomena.

If Latour and Woolgar's (1979) description is to be trusted, laboratories and research teams are organized around certain 'articles of faith', scientifically speaking. That is to say, they take as their mission making into 'received wisdom' certain scientific speculations that are not yet generally accepted. The work of any particular research team or any particular laboratory simply proceeds 'as if' its preferred truth were true, and goes on from there. In the greater game of science, some teams win and many lose. But it is the very rare team/lab that wins by hedging its bets. The way to play the game is, on this account of the sociology of science, to put all your chips on some speculation or on some particular small sets of cognate speculations.

What is true of scientific research teams and labs writ large is as true of individual careers, writ small. In choosing to specialize in one subfield rather than some other, in choosing to associate one's name with one 'thesis' rather than another within one's chosen subfield, one is inevitably making a more-or-less informed bet that that proves

https://doi.org/10.3366/epi.2006.3.3.166 Published online by Cambridge University Press

to be a fruitful vein of enquiry that will ultimately be seen to be vindicated. A few individual scholars, just as a few research teams, 'hedge their bets' by associating themselves with several different hypotheses at once. Insofar as they do, outside observers have correspondingly reduced expectations about any differential biases being built into the results that they report. But most individuals, like most research labs/teams, seem to reckon that the path to professional fame and glory lies in the riskier strategy of putting all their chips on the ultimate truth of a strictly limited set of propositions.

Certainly having no distinctive theses of one's own is deeply deleterious to one's reputation. Exercises in pure 'replication' – merely seeing if you can confirm findings that are already on record and that are already generally accepted – are hardly the stuff of which great reputations are made. That fact is much bemoaned, and rightly so. Many of the findings that are 'already on record' and that are 'generally accepted' are often based on very slender evidence indeed. There is every reason to suppose that attempts to replicate them would indeed fail (King 1995). But look at it from the perspective of individual scholars or research teams trying to build a reputation. What is in it for them in revisiting some generally accepted hypothesis in this way?

Occasionally, major professional reputations are made by being 'giant-killers'.[4] But rarely does a failure to replicate decisively dislodge a generally accepted hypothesis. More typically, the upshot of an unsuccessful replication is much more commonly merely 'case not proven, more investigation required'.

At worst, from the point of view of the scientist trying to make a reputation, replication replicates – confirming what everyone thought they already knew. There are no professional kudos in that. At best, from the reputationally hungry scientist's point of view, replication fails to replicate – failing to confirm what we thought we already knew. But insofar as that is merely that 'more work is needed', there are no real professional kudos in that, either. That is the stuff of scientific underlabourers, not of scientific superstars.

The only way to hit a genuine 'home run', in terms of professional reputation-building in science, is to stake out a positive claim which is genuinely 'surprising' when you first make it but which eventually turns out to be generally corroborated and generally accepted.

If that is a broadly accurate account of the sociology of scientific reputations, it is unsurprising that ambitious scholars and research teams would generally have a 'favourite hypothesis' to which they hitch their professional star and which they fervently hope turns out to be true. That hope, in turn, explains why most scientists and research teams would be expected to be 'biased' in the way I have described, being highly reluctant to report evidence that 'goes against' the hypothesis in which their professional aspirations are invested.

That is not to say that such people's biases lead them literally to 'suppress' embarrassing facts. Sometimes of course they do. But it is not merely, nor probably even mainly, a matter of that.

As every scientist knows, there are lots of ways an experiment can get 'contaminated'

https://doi.org/10.3366/epi.2006.3.3.166 Published online by Cambridge University Press

— lots of things that can 'go wrong'. It would genuinely be a fool's errand to write up for scientific publication the results of every experiment that probably involved a failure of the experimental apparatus – a contaminated Petrie dish or an imperfect vacuum or an electronic surge or whatever. There are dozens, hundreds or perhaps even thousands of results that we regard as 'uninteresting' in these ways, for every gem that we reckon is truly worth reporting.

POLITICIZED EXAMPLES

So far I have been telling this tale in terms of 'scientific incentive structures' and 'laboratory life'. But notice that the same phenomena are found in other realms as well.

Consider politics, most especially. Politicians, too, are 'biased' in just the ways I have been describing. They too have invested heavily in certain propositions being 'true'. Their reputations, and career prospects, are dependent upon the propositions with which they have associated themselves coming to be seen as true propositions by the public at large.

Politicians labouring under those incentives thus have every interest in disseminating information they believe will support their favourite propositions and in being reluctant to disseminate information they believe will undermine their favourite propositions. Biased politicians thus share the biased scientist's differential propensity toward 'false positives' and aversion toward 'false negatives'. Favourable facts they will report with gay abandon, while unfavourable ones they will check many times.

That in turn allows us to draw strong inferences from the reports of a politician that go 'against the grain' of that politician's bias. Just as with the cancer researcher in the employ of the tobacco industry reporting that tobacco is indeed carcinogenic, so too with the old anti-communist witch hunter Richard Nixon saying it is in the American interest to open relations with communist China: in both cases, if *that* person says that, so much against the grain of his or her bias, then we have unusually strong grounds for thinking that that simply *must* be true.

Politicians do not act dramatically 'against the grain' all that often, of course. (If they did too often, they would not have a 'grain' – a perceived bias – to act against.) But it does happen, if not often, nonetheless importantly. Consider these instances:

- Begin, a famously fierce Zionist, signed a peace treaty with Egypt returning the occupied lands of the Sinai.
- Johnson, who won the 1964 US presidential election by characterizing his opponent as a mad war monger, committed the US to the massive troop buildup in Vietnam immediately thereafter.
- De Gaulle, the very embodiment of French imperial grandeur, withdrew from Algeria.
- Nixon, fierce anti-communist, not only opened up US relations with Red China but also signed the first Strategic Arms Limitation Treaty with the USSR.

https://doi.org/10.3366/epi.2006.3.3.166 Published online by Cambridge University Press

Furthermore, when politicians conspicuously act 'against the grain' in these cases, the general public reacted in just the way I suggest they should. Opening relations with China would have been highly controversial in 1970s America, had it been proposed by anyone except someone with such impeccable anti-communist credentials as Nixon. Making peace with Egypt would have been highly controversial in 1970s Israel, had it been proposed by anyone except someone with such impeccable Zionist credentials as Begin. The general public accepted those policies, coming from those politicians, precisely on the grounds that 'if *that* person says so, it must be true'.[5]

INFERRING FROM BIASED REPORTS

Suppose, then, that people – or at least certain, identifiable individuals within that community – are reliably 'biased' in the ways just discussed. What follows from those facts?

Principally, just that such reports from such informants can be one-way decisive. We place no great faith in reports from them that are in line with what we expect from them – what we perceive to be their 'biases'. But we would place great faith in reports from them that go against what we expect from them – contrary to their 'biases'. If they say their favourite hypothesis is not true, then it's awfully likely that it is not true.

Such informants prove invaluable. Information from them, if it goes against their favourite hypothesis, can allow us to decisively reject that hypothesis. Notice, however, that such biased observers provide decisive information in this way purely one proposition at a time (or one cluster of connected propositions that constitute their 'bias' at a time).

If we are trying to choose among several distinct propositions, we will need input from lots of different informants, each biased in favour of different ones of those propositions. Negative reports from each of them can, in effect, 'knock out' propositions that are in line with that informant's own 'bias'. But to assess all the independent propositions that are on the table with the aid of informants who are biased in these ways, we will need (at least) as many biased informants as there are different kinds of bias.

Furthermore, if we are trying to assess that larger set of propositions with the aid of informants who are biased in these ways, we need for those informants to be importantly different from one another – different in their biases, different in which of the propositions they are differentially disinclined to report evidence against.

In short, if we are trusting to biased informants reporting 'against the grain' to help us with the winnowing process that Popper called *Conjecture and Refutation* (1965), we require epistemic diversity of a very particular sort. Specifically, we require a substantial number of informants, who are operating with genuinely diverse biases.

In pursuing this strategy, we will also need lots of reports from all those reporters. While it is true that we will not get many 'false negatives' from biased observers – so when we get a report of not-ø from someone biased in favour of ø we ought take that particularly seriously – it is also true that, by virtue of that same bias, we will not

https://doi.org/10.3366/epi.2006.3.3.166 Published online by Cambridge University Press

get many 'negative' (not-ø) reports from such observers at all. The nature of the bias instead inclines them toward disproportionate numbers of 'false positives' (reports that ø is true when in fact it is not). Hence, we need to get lots of reports from such observers in order to stand any chance of getting the rare, telling 'not-ø' report from them.[6]

DIFFERING BIASES, BUT SHARED EPISTEMIC STANDARDS

So far I have been proceeding implicitly as if everyone in the community, however different their 'biases', nonetheless operated according to the same 'epistemic standards', at least insofar as what they regard as good evidence for the falsity of their 'preferred' proposition I would regard as good evidence for its falsity. In practice, of course, epsitemic standards often differ – and worse, they often vary in ways that are related to the propositions that one would prefer to believe to be true.

That is precisely the way in which 'epistemic diversity' of that other, related sort – diversity of epistemic standards, rather than of preferred propositions – creates political problems. Diversity in that other epistemic respect gives rise to 'culture wars', with those having 'secular' standards of what counts as good epistemic warrant coming into conflict with those who find their epistemic warrants in some sacred text (Kitcher 2001).

The possibility of differing epistemic standards complicates the story I have been telling in the following way:

- Suppose my informant is 'biased' toward proposition ø, in the sense I have been discussing: she would like for ø to be true; she will be unsurprised to find evidence suggesting ø is true, and she will be rather cavalier in reporting such favourable evidence; but she will be surprised to find evidence suggesting ø is false, and will check that evidence very carefully indeed before reporting such unfavourable evidence.
- Suppose however my informant has different epistemic standards from my own. Then evidence that may lead her to reject ø would not necessarily lead me to reject it, and vice versa.
- From the sheer fact that she rejects her favourite proposition ø I therefore cannot infer anything about what attitude I should take toward it. I know she is reluctant to reject ø, but if she rejects it for reasons I would not then the fact she rejects it provides me with no great reason to reject it. What is a negative for her is not for me.
- The converse is also true. Suppose my informant finds evidence that leads her to accept ø but would lead me to reject ø. When she reports evidence that (in her view) 'supports ø', that provides me with no (very strong) reason to reject it: since the evidence 'goes with the grain of her biases', I have no particular reason for thinking she will necessarily have subjected it to any particularly special scrutiny.

https://doi.org/10.3366/epi.2006.3.3.166 Published online by Cambridge University Press

Those sorts of thoughts tend toward the following conclusion. The epistemic advantages of one-way-biased observers accrue, in circumstances of diverging epistemic standards, only where my informant's epistemic standards and my own converge on how to interpret the particular piece of evidence in question, and that evidence 'goes against the grain' of my informant's bias. Then and only then will I have good reason both (a) to make any strong inference from the fact that her conclusion 'goes against her grain' and (b) to suppose that evidence does indeed tell strongly against the proposition in question.

CONCLUSION

As with Popper's *Conjectures and Refutations*, so too with one-way-biased observers: nothing is ever positively proven true; crucial experiments (here, reports from 'biased' observers) simply falsify various conjectures, leaving the rest 'eligible'. Even if the upshot of my account takes that familiar shape, it is nonetheless interesting that this process can be driven even (indeed, especially) by 'biased' observers – just so long as we have a sufficient number and a sufficient diversity of them.

ACKNOWLEDGEMENTS

I am grateful to Mark Brown, Lina Eriksson and Philip Kitcher for discussion of these issues and to Christian List for comments on an earlier draft. Like many of my best ideas, this grows out of a joke originally shared with Diane Gibson.

REFERENCES

Freeman, Derek. (1983). *Margaret Mead and Samoa: The Making and Unmaking of an Anthropological Myth*. Canberra: ANU Press.

Freeman, Derek. (1999). *The Fateful Hoaxing of Margaret Mead*. Boulder, Colorado: Westview.

Goodin, Robert E. (1983). Voting through the looking glass. *American Political Science Review* 77: 420–34.

King, Gary. (1995). Replication, replication. *PS: Political Science & Politics* 28: 444–52.

Kitcher, Philip. (2001). *Science, Truth and Democracy*. New York: Oxford University Press.

Latour, Bruno and Steve Woolgar. (1979). *Laboratory Life*. Beverly Hills, California: Sage.

List, Christian and Philip Pettit. (2007). *Agency Incorporated*. Manuscript, Princeton University and London School of Economics.

Marx, Karl. (1852). Eighteenth brumaire of Louis Bonaparte. In *Complete Works*, vol. 11, pp. 99–197. London: Lawrence & Wishart, 1979.

Popper, Karl. (1965). *Conjectures and Refutations*. New York: Harper & Row.

Quine, W. V. O. (1961). Two dogmas of empiricism. In *From a Logical Point of View*, 2nd edn., pp. 20–46. Cambrdige, Mass.: Harvard University Press.

Williamson, David. (1996). *Heretic: Based on the Life of Derek Freeman*. Ringwood, Vic.: Penguin/Currency Press.

https://doi.org/10.3366/epi.2006.3.3.166 Published online by Cambridge University Press

NOTES

[1] Michelson and Morley had dozens of 'erroneous' oil drop experiment results that they did not report, for every one they did, for example.

[2] Quine 1961.

[3] The 'joke' (which I take to tap an important truth, all jest aside) which prompted this article was just this. Empirical hypothesis testing in the social sciences ought be regarded as a loose (and only one-way-decisive) 'reality check': given all the alternative ways of specifying your test, it ought be easy to find at least one (respectable) one among them that confirms your favourite hypothesis; and if you fail to find any way of supporting your favourite hypothesis, that is therefore truly a damning finding (although finding positive support for it is, pretty much, neither here nor there...).

[4] My ANU colleague, Derek Freeman (1983, 1999), is a case in point: his devastating critique of Margaret Mead established him in popular culture (including an excellent play by David Williamson 1996) in a way that his other more positive academic work never could have done.

[5] I have elaborated this logic in Goodin (1983). Marx (1852, pp. 193-4) had noticed the phenomenon before me, writing in 'The Eighteenth Brumaire of Louis Bonaparte' that, 'Manifestly, the bourgeoisie had now no choice but to elect Bonaparte.... After the coup d'etat, the French bourgeoisie cried: Only the ... [the lumpenproletariat] can still save bourgeois society! Only theft can save property; only perjury, religion; bastardry, the family; disorder, order'.

[6] In the terminology of List and Pettit (2007, ch. 4), this strategy is good at 'truth-indicating' (if an observer biased in favour of ø reports not-ø, that is telling), but it is not particularly good at 'truth-tracking' (such a biased observer does not reliably say ø only when ø is true).

Robert E. Goodin is Distinguished Professor of Philosophy and of Social and Political Theory in the Research School of Social Sciences at the Australian National University. He is the founding editor of the *Journal of Political Philosophy* and general editor of the ten-volume series of *Oxford Handbooks of Political Science*; in that series, he is editor of the volumes on *Public Policy* (with Michael Moran and Martin Rein) and C*ontextual Political Analysis* (with Charles Tilly). His recent own books include *Reflective Democracy* (Oxford University Press, 2003) and *What's Wrong with Terrorism?* (Polity, 2006).

https://doi.org/10.3366/epi.2006.3.3.166 Published online by Cambridge University Press