Synthese
DOI 10.1007/s11229-015-0995-7

The no miracles argument and the base rate fallacy

Leah Henderson[1]

Received: 26 May 2015 / Accepted: 14 December 2015

Abstract The no miracles argument is one of the main arguments for scientific realism. Recently it has been alleged that the no miracles argument is fundamentally flawed because it commits the base rate fallacy. The allegation is based on the idea that the appeal of the no miracles argument arises from inappropriate neglect of the base rate of approximate truth among the relevant population of theories. However, the base rate fallacy allegation relies on an assumption of random sampling of individuals from the population which cannot be made in the case of the no miracles argument. Therefore the base rate fallacy objection to the no miracles argument fails. I distinguish between a "local" and a "global" form of the no miracles argument. The base rate fallacy objection has been leveled at the local version. I argue that the global argument plays a key role in supporting a base-rate-fallacy-free formulation of the local version of the argument.

Keywords No miracles argument · Scientific realism · Base rate fallacy · Base rate neglect

1 Introduction

The no miracles argument (NMA) has been called the 'ultimate argument' for scientific realism (Fraassen 1980; Musgrave 1988). Yet it has also been criticised for begging the question, one of the worst faults that a philosophical argument could have (Laudan 1981; Fine 1986, 1991). Proponents of the argument have made various attempts to mitigate the seriousness of this circularity charge (Boyd 1983; Psillos 1999). Recently,

✉ Leah Henderson
l.henderson@rug.nl

1 University of Groningen, Groningen, The Netherlands

Published online: 28 December 2015

however, a new and potentially equally devastating objection to the argument has been raised. Colin Howson has alleged that when the argument is formulated probabilistically, it becomes clear that it is an instance of the base-rate fallacy (Howson 2000, 2013). This allegation has been endorsed by a number of others (Lipton 2004; Magnus and Callender 2004; Worrall 2005). In this paper, I will argue that this new allegation is mistaken, and the NMA does not in fact commit the base rate fallacy.

2 The no miracles argument (NMA)

The NMA is an argument for scientific realism that takes scientific success as its starting point. The idea is basically the following. If things were not roughly as our best confirmed theories say they are, it would be a miracle that our theories work so well and predict so successfully. We should believe our best scientific theories are approximately true, because, for example, we would not be landing on comets and taking pictures of their surface, if they were not.

In fact the NMA has been formulated in two different versions, a 'global' version, and a 'local' version. The global version concerns the success of science as a whole. The classic formulation of this argument was first given by Hilary Putnam, and developed by Richard Boyd and Stathis Psillos (Putnam 1975; Boyd 1983; Psillos 1999). The idea is that science as a whole is highly successful, and scientific realism provides the best explanation of that success. Scientific realism, then, is taken to be a high level hypothesis about science, that can be supported by evidence of scientific success from the history of science. The argument has been reconstructed along the following lines:

Global NMA

1. If realism is true (i.e if the best confirmed successful theories in mature science are typically approximately true), then the fact that they are successful in new domains is just what we expect.
2. If realism were not true, then the success would be very surprising—a 'miracle'.
3. Theories are empirically successful.

C. Therefore, probably realism is true. i.e. our best confirmed mature and successful theories are probably approximately true.

Some authors have preferred to formulate the NMA as a 'local' argument (Musgrave 1988; Worrall 2005). In this formulation, the explanandum is the success of a particular theory and the explanans is the approximate truth of that theory. The argument then says that the success of the theory is best explained by the approximate truth of that theory.

Local NMA

For a mature theory T which is well confirmed and successful,

1. If T is approximately true, the fact that it is successful in new domains is unsurprising.
2. But if T is not approximately true, its success would be very unexpected—a miracle!
3. T is successful in new domains.

C. Therefore, probably T is approximately true.

3 The base rate fallacy objection

Howson has argued that the NMA commits the base rate fallacy and should therefore be discarded (Howson 2000, 2013). As we shall see, Howson's argument targets the local version of the NMA. The base rate fallacy is a psychological tendency that people have to analyse problems presented in terms of probabilistic information in a certain way (Kahneman and Tversky 1973). It has been revealed by giving experimental subjects a problem such as the following:

> Suppose that doctors conduct a routine test for a rare disease that occurs in 1 in a thousand patients. There is virtually no chance that a person with the disease will test negative (i.e. the false negative rate can be regarded as zero). On the other hand, the chance of a person without the disease testing positive is about 5 % (i.e. the false positive rate is 5 %). A patient takes the test and his result is positive. What is the probability that the patient has the disease?[1]

Most subjects, when confronted with this problem, judge that the probability that the patient has the disease, given a positive test, is very high. However, the correct answer is just under 2 %. This can be seen as follows. Suppose we have 1000 people undergoing tests. One of those people has the disease and tests positive. But there will also be approximately 50 false positives (5 % of 1000), who do not have the disease but still test positive. Therefore the fraction of those with a positive result who actually have the disease is $\frac{1}{51}$ or just under 2 %.

This can also be seen using a Bayesian calculation. Let P represent 'the patient tests positive for the disease' and D represent 'the patient has the disease'. The information that was given in the problem is:

(i) $p(P|D) = 1$ (the false negative rate is zero)
(ii) $p(P|\neg D) = 0.05$ (the false positive rate is 5 %)
(iii) $p(D) = \frac{1}{1000}$

The question is what is $p(D|P)$, that is, the probability that the patient has the disease given that he tests positive for it. The correct answer can be calculated using Bayes' rule

$$p(D|P) = \frac{p(P|D)p(D)}{p(P|D)p(D) + p(P|\neg D)p(\neg D)} = \frac{p(D)}{p(D) + BF.(1 - p(D))}$$

where BF is the 'Bayes factor' $BF = \frac{p(P|\neg D)}{p(P|D)}$ which in this case is 0.05. Combining this with $p(D) = \frac{1}{1000}$ gives the Bayesian conclusion that $p(D|P) = \frac{1}{51}$.

[1] This example is cited in (Howson 2000, p. 52), as a problem which was actually given to students and staff at Harvard Medical school.

One influential explanation of why people give the wrong answer in this, and similar experiments, is that people have a tendency to neglect the base rate. The problems all require a subject to integrate two types of probabilistic information: generic information about base rates in a population—such as the rate of the disease—with specific information about which scenario would make the particular result found more likely—such as the probability that someone with the disease will test positive. One reason that people may overestimate the probabilities is because they neglect the generic base-rates in favour of the more specific information—this inappropriate neglect of base rates is called the 'base-rate fallacy'. So, in the disease case, for instance, people ignore the given information that the overall rate of disease in the population is low, and focus on the information that a patient with the disease is very likely to test positive. The base rate fallacy is by now established as quite a robust phenomenon, which has been revealed in a number of different psychological experiments. This suggests that base-rate neglect may be a common feature of people's thinking. Therefore, suggests Howson, people's thinking about philosophical arguments may be also vulnerable to it.

Howson's argument is directed at the local NMA. Suppose that for a given theory h, T = 'h is approximately true', and S = 'h is predictively successful'. Then if we translate the local NMA as formulated in the previous section into probabilities, we have

(i) $p(S|T)$ is high.
(ii) $p(S|\neg T)$ is very low.

Therefore $p(T|S)$ is high.

According to Howson, someone who draws this conclusion is committing a similar fallacy to someone who concludes that there is a high probability of a patient having the disease, given a positive test result in the medical case. They are failing to take into account that the base rate that the theory is approximately true is very low. The appeal of the NMA, it is alleged, relies on our psychological tendency to neglect base rates. It is not in itself a valid argument. In order to be a valid argument, we would need to assume an additional premise that $p(T)$ is not very low.

4 Response to the base rate fallacy objection

One response to the base rate fallacy objection is to attempt to minimise the importance of priors in the NMA. For example, one might appeal to the convergence of priors. According to this line of thought, if it is repeated and varied successes that drive the NMA conclusion, not just one-off successful predictions, then it doesn't matter much whether the prior probability is high or low. If we update our prior probability for the realist claim T in the light of success of the theory, the updated probability will, as long as the theory continues to predict successfully, get higher and eventually we will end up with a high value for $p(T|S)$. Thus, even someone who starts with a very low prior probability for the realist claim, will potentially end up a realist if enough success is gathered.

However, the realist will not necessarily want to rely on this convergence of priors, since we don't know how long we will have to wait before the realist claims accrues substantial posterior probability. The strategy weakens the realist position since we cannot say confidently at any given moment that the theory is probably approximately true. We can only say that it is getting more likely that it is. Thus most realists will still want to make a stronger argument.[2]

I claim that even if we do not deny that priors play a role in the argument, still the base rate fallacy allegation should be rejected. This is not because the base rate of approximately true theories is not very low. Indeed approximately true theories do make up a tiny subset of the skeptical alternatives considered by the anti-realist skeptic. However, even if the base rate is very low, this does not mean we should necessarily use that base rate as our prior probability. To see this, let us return to the diseases case. So far we have assumed that the individual in question was randomly drawn from the population as a whole. But consider what would happen if this were not true. Suppose instead that the individual were sent for the medical test by the doctor after presenting various symptoms of the disease. If the sampling process is such that those being sent for test are more likely to have the disease, then it is no longer appropriate to take the prior to be the base rate in the overall population. And in this case, it need not be so unlikely that the person does have the disease, if he tests positive.

Similarly, if the scientific method somehow selected a group from the general population which was not just a random sample of all theories that fit the data observed so far, then again it would no longer be appropriate to take the prior to be the base rate of approximately true theories in the overall population of skeptical alternatives. And again, in this case, it is not necessarily unlikely that the theory is approximately true, if it is successful.

The moral here is that one should only take the base rate in the population as one's prior if one knows that the individual was randomly drawn from the population. In the cases which have been presented to subjects in psychological experiments, this is usually stipulated, or at least implied.[3] However, in the case of the NMA, we cannot assume that the theory we have in hand was drawn randomly from the population of skeptical alternatives. In fact, the nature of the sampling process is exactly the point at issue in the realism debate. In the global NMA, realism is treated as a high-level hypothesis—and it is a hypothesis about the nature of the sampling process that scientific method produces. The realist hypothesis R says that the theories which the scientific method produces as best-confirmed in their given fields, are mostly approximately true. The anti-realist makes an alternative hypothesis AR about the sampling process. The exact nature of this alternative depends on the type of anti-realist in question, but it generally asserts that whilst the best confirmed theories have empirical merits (they may be empirically adequate, or well supported by the empirical evidence so far observed), they are not necessarily approximately true.

[2] Psillos' attempt to reconstruct the NMA in likelihoodist terms (Psillos 2006) similarly yields too weak a conclusion, as pointed out in (Howson 2013). In the context of the positivist versus realist dispute, Dorling discusses conditions under which the accumulation of ordinary observational evidence may move someone with a low prior for realism by Bayesian conditionalisation to the realist position, though he is not specifically concerned with the NMA (Dorling 1992).

[3] Notably, however, in variants of the experiments, where subjects do the random sampling themselves, base rate neglect is substantially diminished (Gigerenzer 1991).

Let us return for a moment to the diseases case. Suppose you do not know what kind of sampling was in fact used. Is there a way to tell, by looking at the outcomes of the test? It seems there is. If the sampling process is random, we expect relatively few positive test results out of the total number of test results, whereas, if the sampling process does produce a group in which more people have the disease, there will be a relatively large proportion of positive test results. Thus, the overall proportion of people testing positive could give us some information about whether the tested individuals were selected randomly from the population, or non-randomly by their doctors. If we find a high proportion of positive tests, it is more likely that the individuals were selected for test by a non-random selection process. This in turn would give us reason to think that the prior probability that a tested individual has the disease is not very low.

Similarly, according to the global NMA, a high overall proportion of success among theories serves as evidence that the scientific 'sampling' procedure is not random, but rather biased in favour of approximately true theories. That is, the fact that scientists so often succeed in coming up with successful theories is a reason to think that the scientific method generally produces approximately true theories. Again, in turn this gives us a reason to think that the prior probability that a particular best confirmed scientific theory is approximately true is not very low.[4]

The local NMA then does not stand alone. It is supported by the global NMA. In the local NMA the realist does not neglect the low base rate of approximately true theories among the skeptical possibilities. She would only be guilty of ignoring that base rate if she knew that the theory was randomly sampled from the population. But in fact, the realist *argues* that the theory is *not* randomly sampled by appeal to the success of science more broadly, as in the global NMA.

One might wonder whether the issue of base rates has just been pushed up to the level of the global argument. Certainly, in the global NMA, there is still a question of what priors to assign to R and AR respectively. However, whatever issue there may be over this, there should be no suggestion that a base rate fallacy is committed at this level. At the local level, the allegation was based on the idea that the base rate of approximately true theories amongst all scientific theories is very low. However, at the global level, there is no analogous base rate—we just have alternative hypotheses about the nature of the sampling process. As long as we regard the realism debate as

[4] Magnus and Callender also consider the possibility that in order to assess the prior probability for T the realist could take the set of theories $\mathcal{H}$ 'actually professed by our mature sciences', then 'sample the overt declarations of mature sciences and check their success' (Magnus and Callender 2004, p. 326). However, they reject this possibility, saying 'this would be a biased sample, since theories in mature sciences were chosen (in no small part) because they were successful. Any theory [in $\mathcal{H}$] would probably be successful, just on account of its membership in $\mathcal{H}$' (Magnus and Callender 2004, p. 326). Magnus and Callender then say that this route to avoiding base rate neglect is only accomplished 'at the cost of sample selection bias'. However, Magnus and Callender's objection rests on not clearly distinguishing between success at different times. The set of theories under consideration is the set of those which are best-confirmed in their fields. It is true that all those theories have been successful in the past. But the question at issue in the global NMA is whether they continue to succeed when tested in novel ways. There is no guarantee that a theory which succeeded before will go on succeeding, rather than needing to be replaced in substantive ways. Thus Magnus and Callender's statement that any best confirmed theory would trivially be successful by virtue of its membership in the set $\mathcal{H}$ is not true.

a live issue, which is not already effectively decided in favour of the anti-realist, the prior probability for the realist hypothesis in the global argument will not be vastly smaller than that for the anti-realist hypothesis.

5 Conclusion

In this paper, I have argued against the allegation that the NMA commits the base rate fallacy. The allegation is based on the idea that the NMA relies on inappropriate neglect of the very low base rate of approximately true theories among the skeptical alternatives. However, I have argued that it is only appropriate to use that base rate as your prior if the theory in question is randomly sampled from the overall population. In the case of the NMA, we cannot assume that there is random sampling. Rather the realist's argument that a theory is probably approximately true, given its success relies not only on support from the particular success of that theory, but also implicitly involves a more global form of the NMA that takes scientific success in general as evidence that the sampling provided by the scientific method is not random.

This conclusion does have some broader implications. Howson clearly means his allegation to do in the NMA. He describes the argument as 'dead in the water' (Howson 2013). This has led some authors to suggest that perhaps the realist needs a whole new approach. In particular, Magnus and Callender describe the NMA as a 'wholesale' argument which 'seeks to explain the success of science in general' (Magnus and Callender 2004, p. 321). They suggest that, in the light of the base rate fallacy objection, wholesale arguments for (and against) realism should be abandoned in favour of a 'retail' approach of assessing the scientific evidence for each disputed element of a theory on a case-by-case basis. [5] This is actually an odd moral to draw, given that the local version of the NMA is not a 'wholesale' argument, and, as we have seen, the base rate fallacy objection targets the local version. Be this as it may, what I hope to have shown is that the base rate fallacy objection fails, and hence it provides no grounds for a call for the realist to adopt a new strategy.

Acknowledgments I would like to thank Agustín Rayo, Robert Stalnaker and Roger White for discussion of early versions of this paper.

References

Boyd, R. N. (1983). On the current status of the issue of scientific realism. *Erkenntnis*, *19*, 45–90.

Dorling, J. (1992). Bayesian conditionalization resolves positivist/realist disputes. *The Journal of Philosophy*, *89*, 362–382.

Fine, A. (1984). The natural ontological attitude. In J. Leplin (Ed.), *Scientific realism* (pp. 83–107). Berkeley: University of California Press.

Fine, A. (1986). Unnatural attitudes: Realist and instrumentalist attachments to science. *Mind*, *95*, 149–179.

Fine, A. (1991). Piecemeal realism. *Philosophical Studies*, *61*, 79–96.

[5] Interestingly, similar moves were already made by Fine (1991). Fine suggested that realists might take refuge in a 'piecemeal' defence of their position, given that, as he saw it, the NMA was fundamentally flawed because of circularity. Fine actually concluded that the piecemeal approach was equally unpromising, famously leading to a general recommendation to dissolve the realism debate altogether (Fine 1984).

Gigerenzer, G., Eells, E., & Maruszewski, T. (1991). On cognitive illusions and rationality. *Probability and rationality: Studies on L. Jonathan Cohen's philosophy of science* (pp. 225–250). Amsterdam: Rodopi.

Howson, C. (2000). *Hume's problem: Induction and the justification of belief*. Oxford: Oxford University Press.

Howson, C. (2013). Exhuming the no-miracles argument. *Analysis*, *73*(2), 205–211.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, *80*(4), 237–251.

Laudan, L. (1981). A confutation of convergent realism. *Philosophy of Science*, *49*, 19–49.

Lipton, P. (2004). *Inference to the best explanation* (2nd ed.). London: Routledge.

Magnus, P. D., & Callender, C. (2004). Realist ennui and the base-rate fallacy. *Philosophy of Science*, *71*, 320–338.

Musgrave, A. (1988). The ultimate argument for scientific realism. In R. Nola (Ed.), *Relativism and realism in science* (pp. 229–252). Dordrecht: Kluwer Academic Publishers.

Psillos, S. (1999). *Scientific realism: How science tracks truth*. New York: Routledge.

Stathis, P. (2006). Thinking about the ultimate argument for scientific realism. In C. Cheyne & J. Worrall (Eds.), *Rationality and reality: Conversations with Alan Musgrave* (pp. 133–156). Berlin: Springer.

Putnam, H. (1975). *Mathematics, matter and method*. Cambridge: Cambridge University Press.

van Fraassen, B. C. (1980). *The scientific image*. Oxford: Oxford University Press.

Worrall, J. (2005). Miracles, pessimism and scientific realism. LSE webpage.