

May 17th, 9:00 AM - May 19th, 5:00 PM

Bias, Bayes, and Group Psychology

Guillermo Barron
Red Deer College

Follow this and additional works at: <https://scholar.uwindsor.ca/ossaarchive>

 Part of the [Philosophy Commons](#)

Barron, Guillermo, "Bias, Bayes, and Group Psychology" (2001). *OSSA Conference Archive*. 3.
<https://scholar.uwindsor.ca/ossaarchive/OSSA4/papersandcommentaries/3>

This Paper is brought to you for free and open access by the Department of Philosophy at Scholarship at UWindsor. It has been accepted for inclusion in OSSA Conference Archive by an authorized conference organizer of Scholarship at UWindsor. For more information, please contact scholarship@uwindsor.ca.

Title: Bias, Bayes, and Group Psychology

Author: Guillermo Barron

Response to this paper by: David Hitchcock

© 2001 Guillermo Barron

ABSTRACT. An agent who desires, *ceteris paribus*, to maximise her true beliefs and minimise her false beliefs will be rational if she pays due attention to the question of bias when appraising the testimony of others, since bias may cause a testifier to say things which are not true. But how much is “due”? According to the genetic fallacy, it does not follow deductively that testimony is false because it is biased. However, Alvin Goldman and Wesley Salmon have argued that where an testifier has historically made false claims due to bias, this is good probabilistic evidence to think her future (biased) claims false as well. And Elliot Sober argues that if the cause of a belief is *independent* from the grounds of its truth, this is a good probabilistic reason to count the belief false. I contend to the contrary that genetic arguments (with very rare exceptions) do not offer even probabilistic reasons for disbelief or even agnosticism about claims supported by biased testimony, and I offer a Bayesian analysis of Sober’s arguments as well as intuitively plausible arguments to support my claim. Moreover, “independence,” as Sober defines it, is neither a necessary or sufficient condition to support inductively compelling genetic arguments.

Annette Baier, on the other hand, has suggested that recognizing genetic arguments as fallacious requires us to “ignore” the origins of ideas. I show that this interpretation is also mistaken. I offer some suggestions on when and where attributions of bias will be helpful, and I rely on robust, well-confirmed, and widely-observed psychological phenomena in group psychology to suggest that suggestions of bias, even where true, will very frequently reduce, rather than increase, an agent’s epistemic success.

One reason why we might aspire to a value-free science is that some values - call them “biases” - will unacceptably influence the way scientists frame scientific questions, select, interpret, and reject data, devise methodologies, and formulate explanations. Since theories are always underdetermined by observation, observer bias will fill the epistemic gap and create observer-relative theories. In short, bias can taint an entire scientific enterprise with non-objectivity. Following Robert Nozick, let us call this the Contamination Thesis (Nozick 1997, 34). It is now tempting to suggest that that if some theory Q is tainted by bias, that fact will count as a reason to think that Q is false. If we combine the Contamination Thesis with the further claims that (1) bias is endemic to some research tradition *and* (2) that this bias is the *chief* causal force behind theory construction, promulgation, and acceptance within that tradition, we now have a powerful conceptual tool with which to reject entire research traditions.

However, the Contamination Thesis is not sufficient to establish that a theory is false. Practical logic warns us that genetic arguments such as these are invalid and fallacious since even apparently disreputable origins can yield true claims (Copi 1996, 121-2). But there are at least three suggestions that this is not the whole story.

- 1) Lorraine Code points out that there is a curious asymmetry between the way we consider appeals to authority and the way we treat genetic arguments (Code 1993, 27). To paraphrase Code slightly, if an unbiased and credible authority utters some claim P, we typically count this as a good probabilistic reason to think that P is true. By parity of reason then, if a biased and therefore *non-credible* source utters Q, this ought to be a good probabilistic reason to think that Q is false. In other words, genetic arguments are (or ought to be) the epistemic

mirrors of appeals to authority. Of course, in saying this, Code is in no way arguing that all appeal to authority or genetic arguments are legitimate.

- 2) Alvin Goldman, following Wesley Salmon, also argues for parallelism between appeals to authority and genetic arguments if those arguments are based on an agent's epistemic history. Suppose an agent makes some set of claims about subject S and the great majority of these claims prove false. If the agent makes some further claim Q about S, Goldman claims it is then simply a matter of "proper induction" to conclude that Q is false as well (Goldman 1999, 152-3).
- 3) Elliot Sober has argued that insofar as an agent's belief is independent of the truth of the belief, it is for that reason likely to be false. Given these considerations, it seems that a probabilistic version of the genetic argument may be acceptable where a deductive version is not.

I contend to the contrary that both the "hard" (deductive) and "soft" (probabilistic) variants of the genetic argument are flawed. Except for some exceedingly rare cases that I'll discuss later, *evidence of human bias is never a good reason to think a claim is false*. I shall use Bayes' Rule and intuitively acceptable arguments to show that Sober's argument is mistaken and shall offer some normative suggestions about the epistemic value of genetic arguments.

Sober's Probabilistic Genetic Argument

Sober argues that even though deductive forms of the genetic argument are indeed invalid, this point has been over-interpreted (Sober 1993, 206; 1994, 105). There are, he argues, perfectly respectable probabilistic versions of the genetic argument. Sober offers this thought experiment, which he directs against a variant of subjectivist ethics, but which is more applicable:

Suppose I walk into my introduction to philosophy class one day with the idea that I will decide how many people are in the room by drawing a slip of paper from an urn. In the urn are a hundred such slips, each with a different number written on it. I reach in the urn, draw a slip that says "78," and announce that I believe that exactly 78 people are present. (206)

Since Sober's belief is almost certainly incorrect, Sober thinks we can construct the following genetic argument:

- (1) Sober decided that there were 78 people in the room by drawing the number 78 at random from an urn.

p =====

It isn't true that there are 78 people in the room.

The " p " and double line indicate that the argument is non-deductive and that the premise confers probability p on the conclusion. Sober contends that p is high and that this is "a perfectly sensible genetic argument" in which "the conclusion is justified *because* of the process that led me to this belief (Sober 1993, 206, emphasis added)" even though "*what caused me to reach the belief had nothing whatever to do with whether the belief is true* (Sober 1993, 207,

emphasis in original).” By way of contrast, if Sober’s alter ego Rebo carefully counts all the people in the class and consequently believes there are 104 people present, we have good probabilistic grounds to think that Rebo is right, because she arrived at her belief in a respectable way.

Sober’s moral is this: where an independence relation holds between a belief’s cause and the truth of the belief, the belief is likely false. Contrariwise, if there is a dependence relation between the belief’s cause and its truth, the belief is likely to be true. So this example proves that genetic arguments can offer probabilistic grounds to think some claims are false (Sober 1993, 207). This interpretation looks intuitively compelling, but I think it’s incorrect, and Bayes’ Rule demonstrates this conclusively. According to the simplified version of Bayes’ Rule:

$$(2) P(Q|R) = (P(Q) \times P(R|Q))/P(R)$$

where:

Q = There are exactly 78 people in the class.

P(Q) = The probability that (Q)

R = Sober randomly draws the number 78.

P(R) = The probability that (R) = 1/100

P (Q|R) = The probability that there are exactly 78 people in the class conditional on the fact that Sober randomly draws the number 78.

P (R|Q) = The likelihood that Sober draws 78 conditional on the fact that there are exactly 78 people in the class. R and Q are independent, so P (R) = 1/100.

Substituting these values in (2) yields

$$(3) P(Q|R) = (P(Q) \times 1/100)/(1/100)$$

which simplifies to:

$$(4) P(Q|R) = P(Q)$$

In other words, the probability of there being 78 people in the room given that Sober drew 78 from the urn is exactly equal to the prior probability that there are 78 people in the room. Therefore Sober’s conclusion that there are exactly 78 people in the room is probably false just in case we think P (Q) is small. But Sober’s argument (1) offers no evidence whatsoever that P (Q) is small. This suggests that the argument relies crucially on a suppressed premise:

(5) Sober decided that there were 78 people in the room by drawing the number 78 at random from an urn.

[University classes rarely contain exactly 78 people.]

p =====

It isn’t true that there are 78 people in the room.

Making this suppressed premise explicit shows that it, and not Sober’s independence thesis, is in

fact doing all the evidential work. If you doubt this, consider these two variants of Sober’s argument:

(6) Sober decided that there were 78 people in the room by drawing the number 77 at random from an urn.

[University classes rarely contain exactly 78 people.]

p =====
It isn't true that there are 78 people in the room.

(7) Sober decided that there were 78 people in the room by drawing the number 78 at random from an urn.

[Universities rigidly enforce rules requiring there to be exactly 78 people in every class.]

p =====
It is *true* that there are 78 people in the room.

Now let me make the disagreement between Sober and me as explicit as possible. Sober thinks that:

(8) “drawing 78 at random” and believing that “there are 78 people in the class” are two independent states of affairs.

_____ (1) is a “convincing” argument (Sober 1993, 207).

While I contend that

(9) “drawing 78 at random” and believing that “there are 78 people in the class” are two independent states of affairs.

_____ (1) is a weak argument.

(2)

I want now to diagnose just why we disagree. Consider for a moment the following matrix:

	INDEPENDENT BELIEF FORMATION: Sober draws the number 78 from an urn and consequently ...	DEPENDENT BELIEF FORMATION: Rebo carefully counts the people in the room and consequently ...
--	---	--

YIELDS TRUE BELIEF (PROBABLY)	POINTLESS RANDOM CHOICE: Sober believes there are NOT 78 people in the room.	EMPIRICISM: Rebo believes there are 104 people in the room.
YIELDS FALSE BELIEF (PROBABLY)	RANDOM CHOICE: Sober believes there are 78 people in the room.	PERVERSE EMPIRICISM: Rebo believes there are NOT 104 people in the room.

In the first column, the cause of both of Sober's beliefs is independent of the facts. But the belief generated by Pointless Random Choice in the upper square is almost certainly true. Hence it is false to say that the independence relation cannot reliably produce true beliefs. The lower right hand corner represents beliefs formed by Perverse Empiricism. A Perverse Empiricist believes some claim Q iff she has carefully investigated Q and found compelling evidence that Q is false. Hence her beliefs are dependent on the truth, but, so to speak, inversely so. This is the only case in which a genetic argument has any force. Precisely because we know the Perverse Empiricist's epistemic practices reliably create false beliefs, appeals to those origins count legitimately as a reason not to think her claims false. So, properly considered, the argument from authority and the genetic argument are indeed analogous. One can, for example, infer that a given measurement is inaccurate if one knows that the measuring instrument that produced it reliably measures too high or too low. But because human Perverse Empiricism is pathological, rare, and largely irrelevant to scientific study, I will henceforth disregard it.

Sober only considers the beliefs represented by the lower left hand and upper right hand boxes. This, I suggest, is why he thinks that a belief's plausibility is linked to its dependence. But the examples of Pointless Random Choice and Perverse Empiricism prove that dependence is neither necessary nor sufficient for true belief. We also need to recognise that Sober's example also fails (no doubt for the sake of lucid exposition) to model real world belief formation procedures where degrees of epistemic in/dependence are far more difficult to ascertain and where prior probabilities may not be so obviously minuscule.

If my line of reasoning is correct, Sober has given us no reason to think that the origins of a belief will count as reasons to think that it is false - even where those beliefs are arrived on the basis of random choice. Goldman's contention that one can make a persuasive inductive argument from an agent's past false claims fares no better. Suppose Jones makes a set of false claims {A ...P} about subject S. If Jones further asserts claim Q about S, isn't her espousal of the earlier set of false claims a good reason to think that Q is also false? No. If there is some direct logical or evidential link between {A ...P} and Q such that the falsity of the former is evidence for the falsity of Q, then you have good reason to think Q is false independently of Jones's actually asserting Q. For example, if Q entails {A ... P}, then Q must be false. On the other hand, if Q is logically independent of the former discredited claims, then its probability cannot

be less than its probability prior to Jones asserting it, and this follows for precisely the same reasons I adumbrated above against Sober's example. In neither case is the shared origin of claims $\{A...P\}$ and Q relevant to our evaluation of Q 's truth-value.

Bias and Probability

Evidence of bias is thus never, all by itself, a reason to disbelieve any claim Q . And if the prior probability of Q is high, it may not even be reason to be agnostic about Q . For my purposes here, I define "agnostic" as having no reason to prefer one truth-value over another for some claim. Formally, the agent might say she assigns a probability of 0.5 to that claim. But it is probably more common that she is unwilling to be so precise. In these cases, it is sufficient for my purposes that the agent be much more willing to accept a probability close to 0.5 than she is to accept a probability of 1 or 0. Now consider the case in which exactly one of four theories $[A, B, C, D]$ is true, and one has no reason to prefer any of them (i.e., $P(A) = P(B) = P(C) = P(D)$). The prior probability of not- C is therefore 0.75. Later evidence might increase one's belief in not- C , but if this evidence turns out to be biased, one is still justified in assigning not- C a relatively high level of probability (0.75).

Nonetheless, it is straightforwardly false to assert, as Annette Baier does, that arguments against genetic arguments require us to "ignore" Q 's origins as irrelevant (Baier 1995, 325). After all, in many cases, knowing that Q comes from unsavoury origins will be good reason not to increase one's degree of assent. But this will only be the case where a testifier is strongly towards Q . Where an agent is only weakly biased towards Q , it may be rational to increase one's belief in Q even though one knows that the testimony for Q is biased. Consider, for example the testimony of two chicken sexers. One is strongly male-biased- she identifies all chickens as male. The other is only weakly pro-male in her assessments - she correctly identifies all male chicks as male, and misidentifies only 1% of female chicks as female. It is therefore rational to believe the testimony of the latter (and to believe it to a high degree) - even though we know her testimony to be biased.

So a rational agent who desires to maximise her quotient of true (or justified) beliefs over false beliefs will be acting in an epistemically responsible manner if she takes credible accusations of bias seriously. I offer some modest suggestions about how she should do this.

First, it is tempting to believe that if a scientific claim Q has a low prior probability, one can therefore infer the likely existence and influence of bias in its creation. Given this, one can next deduce the nature of that bias from the content of Q itself, and this inference can then be used to erect a genetic argument against the veracity of Q . (Lorraine Code's dialogical epistemology apparently licences this methodology. Code employs this form of argument against Philippe Rushton, who has argued (notoriously) that there is an interracial correlation between penis size and intelligence. Code contends that since Rushton could not have found his data "by coincidence," he must therefore have been driven by some right-wing agenda. Code proceeds to lay out in detail what Rushton's politics must be, and then suggests that the existence of these politics constitutes a probabilistic reason to reject Rushton's work (Code 1993, 28-9). This

tempting conclusion, however, rests on a false dilemma, since there is any number of motivations that could have informed Rushton's research.)

This practice is, however, deeply flawed for several reasons. First, the existence of bias or error will not be revealed in the prior probability of Q itself, but in the conditional prior probability of an agent's asserting Q *given that* Q is in fact true (Nozick 1993, 101). Let me explain this a bit. If, for example, one *knows* that Q is false, one might *guess* which biases might have led a researcher to espouse it. Contrariwise, if one *knows* that a researcher has a given set of non-scientific commitments, one might *guess* how those commitments would affect her work. But both of these approaches are very dodgy enterprises, since there are any number of non-scientific considerations that might have motivated a researcher, and since a researcher's political and moral commitments do *not* exert a deterministic influence on her scientific claims. After all, people frequently do arrive at counterattitudinal beliefs (Goldman 1999, 236). Now consider the case in which one knows neither that Q is false nor that bias played a role in Q's construction. In this case, to infer the existence of bias from the content of Q and to then argue that that bias now counts against the truth of Q is surely to build epistemic castles in the air. And, as the Bayesian argument above shows, erecting probabilistic arguments against Q that are based on Q's own low prior probability will lead to double discounting. So, to avoid these evils, claims that a researcher is biased should be based on independent evidence about the researcher's non-scientific commitments.

Some science critics, (Rose 1984, 8) for example, have assumed that this measure is sufficient all by itself: prove that a scientist has a given political commitment and you've proved that it also adversely affects her research. But whether this is so is an empirical question, and must be settled by empirical means. Alvin Goldman argues that several flaws hamper many case studies on scientific bias. First, studies that show that political interests are coincident with claim Q cannot, by their very nature, establish the counterfactual condition that had those political facts not obtained, the claim Q would not have been made. Such case studies therefore cannot establish the causal efficacy of politics on the development of Q. And even where they do, they are less persuasive in explaining Q's continued acceptance. Finally, Goldman suspects, many case studies are not undertaken on a random or representative set of scientific episodes, but are handpicked to prove the very points which science critics wish to make (Goldman 1999, 37-40). Consequently, these case studies cannot be used to licence generalisations about the effect of bias across all science. Given all this, it seems that the best way to conclusively prove that bias has led a researcher astray is to show first that she did go astray, and then to show that bias was the cause. But where this can be done, one of course no longer needs a genetic argument.

Furthermore, Robert Nozick has argued that no factor is intrinsically biasing and that whether or not a factor biases epistemic products depends crucially on the process in which it occurs. For example, although jurors are supposed to be unbiased, it may well be that a jury with two biased and opposed jurors will more frequently arrive at the truth - and this, again, is an empirical question (Nozick 1993, 33-4). And an example drawn from Donald Brown suggests that political bias can in some cases increase a researcher's credibility. If, for example, feminist

anthropologists have political interests in discovering evidence of ancestral matriarchies, then when feminist anthropologists such as Pam Bamberger and Sherry Ortner fail to find any such evidence, this is particularly persuasive in showing that matriarchies did not in fact exist (Brown 1991, 52).

All this aside, one might still think that an awareness of bias cannot help but improve one's critical objectivity, especially when one cannot assign any prior probability to Q. My final point suggests that this may not be so. To see why, notice first that attributions of bias are frequently made in the third person. You and I, gentle listener, have our commitments. They have biases. *We* have intuitions. *They* have prejudices. Notice also that accusations of bias are commonly made on the basis of some difference between us and them. That is, if I warn you about Jones' dualist, anti-feminist, or reductionist bias, I typically do so in the belief that you and I are not dualists, anti-feminists, or reductionists. So when I ascribe bias to some third party, this accusation will frequently elicit in my listener what social psychologists call *ingroup/outgroup bias* (Tyler 1999, 2-3). This well-known and pronounced bias displays three relevant features:

The Minimal Group Paradigm: The listener will display bias against the outgroup even when she knows the differences between groups are minimal. Investigators have found that members of one group will discriminate against another group even when they know that the groups have been divided on the basis of such irrelevant criteria as a coin toss or differing preferences in modern art.

Outgroup Derogation: The listener will tend to favor the ingroup over the outgroup and will attribute more negative attributes to the outgroup. In this vein, Sandra Harding attributes the numerous failures of mainstream science and the many epistemic successes of marginalised knowers to the fact that marginalised knowers can somehow throw off their "covers and blinders" and thereby understand the world "how in fact it is" while scientists are "destined" to study not nature itself, but only "socially constituted objects" (Harding 1993, 54, 64)."

Stereotyping: The listener will tend to believe that there is far less intragroup diversity within the outgroup than within the ingroup. Val Dusek, for example, flatly asserts that "Certainly none of the evolutionary psychologists support major egalitarian change in social or gender arrangements (Dusek 1999)."

Ingroup/outgroup bias is a robust, widespread, and highly confirmed phenomenon and it is almost impossible to overestimate its effect on our epistemic practices. If this model is correct, the mere suggestion that some theory Q serves some pernicious social function, or is situated within some noxious political nexus, or is the product of an oppressive power structure, et cetera, will thereby condemn Q as the doctrine of outgroup members whose beliefs have irredeemably contaminated their scientific understanding. Notice that the very rhetoric of "contamination" supports the gratuitous assumption that it is only ideologies that we find repugnant that could contaminate science. If, for example, one were a liberal, one would hardly say that Jones' liberal commitments had "contaminated" her scientific views. And the claim of "contamination" further implies - again gratuitously - that the degree of bias is absolute and total, where in fact only a slight bias may be present. (It is incoherent to suggest that a research program is

“contaminated” with political bias and then to suggest that its testimony is still somewhat credible.) And the assumption that all bias contaminates completely invites the further gratuitous conclusion that if there is a political cause for believing Q, there is therefore no epistemic reason to believe it. Even worse, such accusations will prevent us from seriously considering Q, because serious consideration of Q implies the possibility of conversion, and if the group bias theory is correct, many of us will fear conversion to outgroup beliefs more than we fear error.

If group bias effect is universal, powerful, and anti-veritistic, nostrums counseling open-mindedness are simply not sufficient. Rather, we should take care to construct arguments in ways that do not trigger well-known epistemic failings (such as group bias) in our listeners. Paul Viminiz suggests (p.c., 2000) that the naturalised epistemologist can reasonably object that group bias is an epistemic pattern that must be doing some useful work for us. And my concerns about the argument from bias apply with equal force against my own appeal to group bias. That said, I am not certain that arguments which rely, even implicitly, on group bias are, on the whole, epistemically advantageous within modern science. My intent here has been neither to trivialise the role of bias nor to counsel quietism. Rather, my modest suggestion is that accusations of bias may be incapable of bearing the entire epistemic load that they are sometimes asked to support.

References

- Baier, Annette C. 1995. “A Naturalist View of Persons.” *Moral Prejudices: Essays on Ethics*. Cambridge, Mass.: Harvard UP.
- Brown, Donald E. 1991. *Human Universals*. Philadelphia: Temple UP.
- Code, Lorraine. 1993. “Taking Subjectivity into Account.” *Feminist Epistemologies*. Linda Alcoff and Elizabeth Potter (eds.). New York: Routledge. 15-48.
- Copi, Irving M., and Keith Burgess-Jackson. 1996. *Practical Logic*. New York: Macmillan.
- Dusek, Val. 1999. “Sociobiology Sanitized: The Evolutionary Psychology and Genic Selectionism Debates.” *Science as Culture*. <<http://www.shef.ac.uk/~psych/rmy/dusek.html>>
- Goldman, Alvin I. 1999. *Knowledge in a Social World*. Oxford: OUP.
- Harding, Sandra. 1993. “Rethinking Standpoint Epistemology.” *Feminist Epistemologies*. Ed. Linda Alcoff and Elizabeth Potter. New York: Routledge. 49-81.
- Lemyre, L. and Smith, P.M. 1997. “Intergroup Discrimination and Self-esteem in the Minimal Group Paradigm.” *The Blackwell Reader in Social Psychology*. Miles Hewstone, Anthony S.R. Manstead, and Wolfgang Stroebe (eds.). 558-575. Malden, MA.: Blackwell.
- Nozick, Robert. 1993. *The Nature of Rationality*. Princeton, N.J.: Princeton UP.
- . 1997. “Invariance and Objectivity” in *Proceedings and Addresses of the APA*. 72:2. 21-48.
- Rose, Steven, Leon J. Kamin, and R.,C. Lewontin.1984. *Not in Our Genes: Biology, Ideology, and Human Nature*. Harmondsworth: Penguin.
- Sober, Elliot. 1993. *Philosophy of Biology*. Boulder: Westview.
- . 1994. *From a Biological Point of View: Essays in Evolutionary Philosophy*. Cambridge: CUP.
- Tyler, Tom R., Kramer, Roderick K., and John, Oliver P. , eds. 1999. *The Psychology of the Social Self*. Mahwah, N. J.: Erlbaum