

# TWO FALLACIES

James CARGILE

ABSTRACT: In charging *argumentum ad hominem*, we accuse someone of attacking the source of a claim. In charging *argumentum ad verecundiam*, we attack the source of a claim. This is reason for attending to “attacking the source.” It is important to distinguish probabilistic reasons for doubting a claim and evidentiary reasons. Evidence that the source of a claim is likely to be wrong is not evidence against the claim. The tendency to overlook this is the essential feature of the *ad hominem* fallacy. This is relevant in assessing the view that someone who regards his thinking as made possible by Godless arrangements of matter largely determined by chance is, in taking this attitude, advancing a hypothesis which undermines his theorizing about the world or himself.

KEYWORDS: probabilistic reasons, evidentiary reasons,  
motivational reasons

There is considerable latitude in descriptions of the fallacy of *argumentum ad hominem*. One version makes it simply a matter of attacking the source of a disputed claim rather than addressing the claim itself. There is nothing generically wrong with this, and anyway, the attack may not be offered as a basis for a conclusion. Digressing to denounce an opponent or source of an opposing view is something which may be done in an argument, but is not a pattern of argument, any more than interrupting the argument to rest and read the newspaper is. A more specific kind of *ad hominem* would be to offer a criticism of your opponent as a basis for rejecting his view. You might advance the claim that your opponent is a dishonest politician in support of rejecting his claim that there are no dishonest politicians in his party. This case would not be one of a fallacy of relevance. In a case where the criticism is not relevant to your conclusion, that would be bad, and it is perhaps worth warning against the temptation to argue in this way.

A still more specific kind of *ad hominem* involves concluding that P is false or unjustified on the grounds that it was put forward by a defective source. To find that a case of this is fallacious is to conclude that the arguer has not justified his conclusion that P is false or unjustified. The fallacy of *argumentum ad verecundiam* involves concluding that P is true or justified on the grounds that it is endorsed by a reliable authority, where in fact, the authority is defective, or the

arguer does not have adequate basis for believing the authority to be reliable. (So this is a premise fallacy rather than an inferential one.) To find someone guilty of this fallacy is to conclude that he has not justified his conclusion that P because the source for P to which he appeals is defective or not justifiably believed reliable.

Accusing someone of a fallacious appeal to authority thus bears considerable resemblance to arguing *ad hominem*, in that it attacks the source of a claim as a basis for finding the claim unjustified. It is worth clarifying the relation between these two fallacies. It would be unsatisfactory if finding someone to have committed the fallacy of appeal to a bad authority was itself to commit the *ad hominem* fallacy. Consider these two patterns:

Pattern I: A claims that P. The question whether P is a kind with respect to which A is untrustworthy (unreliable, insane, etc.). Therefore, A's claim is unjustified.

Pattern II: A claims that P solely on the grounds that B testifies that P. A is not justified in regarding B as a reliable authority. Therefore, A's claim is unjustified.

These patterns have the same conclusion, and in the principal case, both support that conclusion by criticizing the reliability of a source of the claim. (II also includes a secondary case in which the objection is that the arguer is not justified in regarding B as a reliable authority. We will set that case aside in the following discussion.) In II, the source may be a different person B. But in the case where A=B, that distinction between I and II is eliminated. Still, even in that case, there must be a good distinction between I and II, since I is a characteristic pattern for the *ad hominem* fallacy, while II is delivering a verdict of fallacious appeal to authority. I is of course not generically fallacious and it is worth making clearer the conditions under which instances of it are fallacious, but those instances should not include pattern II. It should not be generically fallacious to find someone guilty of a fallacy. It is thus worth explaining why I is not a generically reliable style of reasoning and is such that to treat it as a reliable pattern would be to commit the *ad hominem* fallacy, while II is acceptable.

To begin with II, it is a good pattern only when understood with certain qualifications. If A *offers* only the consideration that P was endorsed by B, he may nonetheless have other adequate reason for claiming that P. Boyd may defend his claim that Bob did it by appeal to the fact that Bill testifies that he saw Bob do it. Bill may be a notorious liar, so that Boyd has failed to justify his claim. But Boyd may himself have seen Bob do it. He may have selected the appeal to Bill's authority out of a very unwise assessment of Bill's authority, thinking that his

hearers will be more impressed by hearing it comes from Bill. Boyd has failed to justify his claim, but he *has* adequate justification for making it.

This qualification cannot be adequately included by changing from “based solely on the grounds that...” to “having as grounds, only the consideration that...” This latter is too unclear. Boyd saw Bob do it. Boyd also knows that his eyes are in good working order, and that he is not prone to wild distortions of memory, etc. He does not mention these considerations in offering justification for his claim that Bob did it. He should have mentioned that he saw it. But to claim that his eyes are reliable is best left unsaid. Still it is a consideration. Just where such considerations leave off would be hard to say, and worthless to attempt to describe. The change would make pattern II much less practically applicable than it should be.

In a given context, the presentation of certain premises  $P_1, P_2, \dots, P_n$  may be adequate to justify concluding that  $C$ . One might think that the contextual features which make  $P_1 \dots P_n$  adequate could be summarized in an additional premise  $P_{n+1}$ , to yield a valid and context-independent justification. This is mistaken. A premise can always be found, the addition of which makes the argument valid. All we need is “If  $P_1 \dots P_n$ , then  $C$ .” Whenever anyone argues “ $P_1 \dots P_n$ ; therefore,  $C$ ,” if he is sincere, he believes that *if*  $P_1 \dots P_n$ , *then*  $C$ . If we just add that which he presumably accepts anyway, we always get formal validity. But there is no general guarantee that that will not beg the question.

If it is granted that the features of a situation which make a conclusion reasonable cannot always be stated to yield a context independent argument to justify the conclusion, then we should see that what a reasoner has as grounds or basis for the reasonableness of his conclusion cannot in general be stated completely. Thus we will not be in a position to say that all someone has to ground his claim is such-and-such authority. By contrast, it is easy to be clear about the fact that all someone has offered to ground his claim is such an appeal.

We could change the conclusion from “A's claim is unjustified” to “A has not succeeded in justifying his claim.” This latter is a possible interpretation of the original words, and it makes pattern II quite solid. If A has based his case for  $P$  solely on a bad authority, and has offered nothing else (such as good reason to justify his (nonetheless mistaken) assessment of the authority), then he has not succeeded in justifying his claim, whether or not he is in fact justified in making the claim.

This alteration (or clarification) of the conclusion would not help pattern I. It is just as bad to conclude from the premises of I that A has failed to justify his

claim as to conclude that his claim is unjustified. It is part of our present project to explain this.

We have two versions of the conclusion: (i) "A has not presented adequate justification for his claim" and (ii) "A does not have adequate justification for his claim." I propose to understand (i) in such a way that it follows from (ii), but not conversely. This interpretation is not the only possibility. We can understand saying that someone presented adequate justification for a claim without his being justified in that claim. Someone might present Q,R and S, and these be adequate justification for P, without his getting credited with *having* adequate justification for P, if he himself did not believe Q,R or S and was presenting them with the intent to deceive.

We could understand saying that someone had adequate justification for asserting that P but was not justified in asserting that P. He might know certain things which justify concluding that P and yet fail to recognize this and not believe and thus not assert that P. Thus he 'has' justification in one sense, but not in another. Furthermore, just adding belief that P would not be enough. It would have to get properly connected with the justifying considerations. I will take having adequate justification to entail being justified. More important here is the notion of successfully presenting adequate justification. As I use this notion, it entails having adequate justification. This is not at all strictly observed in ordinary usage, but hopefully it will be clear enough in what follows.

The possibility of an alternate understanding on which a presenter of adequate justification may not himself have adequate justification may explain why people tend to become more wary when presented with Pattern Ii than when given Iii. A thoroughly untrustworthy person may be paid to 'present' the justification of a reliable authority. We may not credit the presenter with being justified in making any of the claims involved while still crediting them to the person for whom he makes the presentation. It seems obvious that merely getting an unreliable person to recite an argument should not serve to make the argument inadequate.

If someone is sufficiently unreliable, it may be reasonable to conclude that he is not justified in claiming that P, for some suitably complex P. Thus if a difficult proposition of mathematics is put forward, and the question is asked whether it is true, if the village idiot pipes up from the back "It is true," we may conclude that this is at best a lucky guess and the man does not even understand the proposition. It could not be true that someone is an idiot (as opposed to, among other possibilities, an *idiot savant*) and yet understood a complex proposition of mathematics. We can reduce this to a matter of how we interpret

the relevant terms. There is a level of incompetence such that it is not possible to be at that level and accomplish certain intellectual successes, just as it is impossible for an absolutely failing student to pass a test. Passing proves conclusively that he is not 'absolutely' failing.

The argument scheme I': "A is incompetent on questions like P, therefore, A did not present a good argument for P such that he can be credited with thus justifying the conclusion that P," is a close relative of Pattern I. Pattern I would infer from the fact that A is incompetent and an argument for P is creditable to A, that the argument is not a good one. I' infers from the fact that A is incompetent that if he has offered a good argument for P, then it is not creditable to him. Another variation (I'') would be to infer that since A has presented a good argument for P, which is creditable to him, he is not incompetent on the topic. To show that I is fallacious involves showing that these variations are also wrong. That is not to show that all instances are bad arguments, but just that the pattern does not give any support to an instance, so that to rely on the pattern is fallacious.

Suppose that the putative village idiot submits a paper to a journal, and the journal politely declines, without reading the paper. When challenged by a peer, they explain that the 'author' is an idiot incapable of producing anything but garble, and they do not have time to read all the papers submitted. This could be a reasonable response. Suppose further, that the peer somehow manages to persuade the editors to read the paper, and they find it contains a brilliant argument for a novel conclusion. They may still quite reasonably object to publishing the article under the idiot's name, on the grounds that he could not have written it. Suppose still further, that the idiot comes in to the office and writes a new paper under close observation, which is also found to contain a brilliant and original argument. Then there is no course but to withdraw the claim that the man is an idiot, or totally incompetent. This still does not undermine the general correctness of the inference from "He's an idiot" to "His paper is worthless." It merely shows the premise of that inference to be false in this case.

Suppose that the putative idiot produces a paper arguing that he is mentally completely incompetent. Again, this is a brilliantly argued paper, appealing to a highly impressive study of the physiology of his brain and the intellectual limitations which follow by well established brain science, from having a brain in such a condition. We are ready to conclude that the production of this paper refutes its thesis. Even though the thesis is wrong, the paper is clever enough to be proof that the thesis is wrong. The paper, of course, does not *contain* proof it is wrong. Rather, the fact it is so clever disproves the thesis that its author is completely incompetent.

The purported idiot may have anticipated this line of criticism and attempted to answer it in his paper. He argues that it is not impossible and not of probability zero that he should produce a cogent argument. It is just fantastically unlikely. We assume the man's general behavior manifests thoroughgoing mental incompetence, so that it is hard to credit his present cogency to the same person. That he is himself amazed, on reviewing what he has written, to think that someone as stupid as he is could have produced such stuff, is a kind of perception, accurate or not, which is not in keeping with his usual manifest character. He says that even though he is too incompetent to be a reliable judge of the soundness of his paper, it looks good to him, and he would like to know what is wrong with it. As he is currently present to us, he is a partner in dialogue, with a claim on our attention and response he would not usually be capable of making.

One reply might be based on the principle that if a theory entails that it is fantastically unlikely that *e* will occur, then the occurrence of *e* is strong evidence against the theory. The idiot's theory implies that it is fantastically unlikely that he would produce a clever paper. But he did produce one, so this theory is refuted. This is a poor reply, since the idiot's theory includes the claim that *e* has happened in spite of being extremely unlikely. Unless you can give reason for thinking there is a better explanation than mere chance, it is question-begging just to complain about the low probability.

Consider this theory  $T:(1)$  Jones tossed an unbiased coin ten times on occasion *O* and got all heads. (2) (1) was due to chance and had a probability of 0.0009765625. Should we say that the probability of  $T(1)$  given  $T$  ( $\text{Prob}(T(1)/T)$ ) is 1, on the grounds that  $T$  entails  $T(1)$ , or should we say that it is 0.0009765625, on the grounds that  $T$  entails that is the unconditional probability of  $T(1)$ ? However this is answered, we should recognize that it is question-begging to reject  $T$  on the grounds that, since one part of it has a probability of 0.0009765625, it is very unlikely that it be correct.

An alternative to  $T$  is  $T'$ : that Jones was cheating and was not tossing in an unbiased fashion. This could of course be right, but this depends on further details. To say *a priori* that  $T'$  is a more likely theory than  $T$  is unwarranted. It is a feature of our human condition that we recognize that there is a nonzero probability that an unbiased coin should be tossed ten thousand times and land heads every time, and yet we must concede that we could not know this happened. If we observe ten thousand heads, we cannot reasonably ascribe this to chance, even if we can think of no other explanation. We cannot reasonably rule out the hypothesis that something is influencing that coin. If someone claims to have psychic power over a coin and offers to prove it by letting us toss the coin four times which he will

make straight heads, and we do get that result, then it would be reasonable to discount this performance as lucky. If we tossed a thousand in this situation, we would have to admit that there is some power at work. The difference between four and a thousand has to do with background information or inclination.

We do, I believe, have background information (or inclination) such that the production of a clever theory concluding that he is an idiot would refute the theorist. Inability to articulate this background information could lead to the illusion that we can properly reject the theory purely because it has a low probability of being true. But it is incorrect to argue "P is a very unlikely reason why e occurred; therefore it is not true that P is the reason why e occurred." The performance of a purported idiot who writes a brilliant argument for the conclusion that he is incompetent is not self-refuting in the sense of the performance entailing the falsity of its conclusion (though it does provide conclusive reason to reject its conclusion). By contrast, asserting you cannot make an assertion is self-refuting in that way. It is impossible for the premise to be true and the conclusion false.

If the conclusion of the argument is that the arguer is incapable of producing a good argument, then the goodness of the argument would refute the argument. Granting that a self-refuting argument could not be a good argument (which is not obvious, but conceded for purpose of argument), then that argument cannot be a good one. Thus the form: A is incompetent in the sense of being utterly incapable of producing a good argument; A argued that P; therefore, A's argument for P is not a good one; is valid. It is a relative of pattern I which is not an inferential fallacy, though the premise is likely to beg the question in a particular case. But the form: A is incompetent, that is, very unlikely to argue well or get anything right. A argued that P; therefore A's argument for P is not a good one, is not valid. To think otherwise would be fallacious.

Someone may of course, be unreliable on one sort of topic while being reliable on another. The interesting cases are those in which the arguer or assertor is unreliable on the topic about which he is arguing or asserting. Here it is important not to slip into a formulation of 'incompetence' such that the mere performance of presenting the argument proves the arguer is not at that level of incompetence. Such a formulation is worth noting, but is not the most interesting case. Once we are clear of those extreme readings of 'incompetent' it is fairly obvious that from the fact that the arguer is ever so unlikely to be right it does not follow logically that he is not right. What is not so clear is whether it is justified to conclude that he is not right.

Suppose that we are given that an urn contains the numbers of 10,001 propositions. 10,000 are proven falsehoods. One is not known to be false and not known to be true (but it is a definite proposition and thus either true or false). We are not given the content of any of these propositions. They may be about topics such as the doings of people in a foreign country, so that we know nothing about any of them just from hearing the content. We make a random drawing and get proposition number  $n$ . It is surely reasonable for us to infer that prop.  $n$  is false. It would be reasonable to infer also, that if a certain stranger believes prop.  $n$ , then he is mistaken. (We might add the assumption that for each of the 10,000 proven falsehoods, there is someone who believes it.)

Proceeding to the question whether it would be reasonable to infer that the stranger's belief in prop. $n$  is unjustified, we should become more cautious. We have not heard his argument, after all. Even if mistaken, he might still be justified. But we can bring this under probabilistic considerations by the same method. We assume an urn of 10,000 bad arguments and 1 unanalyzed argument, and we are given that argument  $X$  has been drawn randomly from that urn. Then it seems to be highly reasonable to conclude that  $X$  is a bad argument.

One difference here from the case of unargued (contingent) propositions is that we can evaluate some arguments even when we cannot evaluate the constituent propositions. We may have no idea whether it is true that Wong grows melons in Daigan or that all melon growers in Daigan are in debt, but we can readily see that it is correct to infer from these premises that Wong is in debt. To that extent we can determine the merit of the argument with less dependence on background information.

This still does not let us determine whether the argument is a good one. As an attempt to establish that Wong is in debt, the arguer may have appealed to two egregious falsehoods. Once familiar with the context of the argument, we might correctly assess it as a woefully inept one, not redeemed by its veneer of validity. Similarly, the argument "Wong has been giving lavish presents to Ho, therefore, Wong is in debt" might be a very shrewd basis for the claim that Wong is in debt, despite its not being valid. Assessing the merit of an argument is not in general an a priori enquiry. And even when it is, say, in mathematics, that does not mean that 'we' can do it. Most of us would not present ourselves as judges of mathematical arguments, and even a distinguished mathematician is apt to be modest about checking some candidates.

The principal upshot here is this: one may have good reason to reject an argument (or a conclusion) without having assessed the argument and without even being able to do so if allowed to study it. It would be a fallacy to mistake this



reasonableness for having an objection to the argument. And this is a distinction which is difficult to articulate and thus a dangerous confusion which can be hard to avoid. The “*ad hominem* fallacy” is a classification which can serve us in warning against the confusion. It is worth reinforcing this classification.

One way is this: that P is selected by a procedure highly likely to pick falsehoods is good reason to reject P, but it is not at all *evidence* against P. To be evident is to be obvious, in plain view. Evidence for P is something which tends to make it evident that P or is offered as tending to do this. Adequate evidence for P succeeds in making it evident that P. Contemporary use of ‘Evidence’ is not always in accordance with this standard, but it would be good to return to it. Not all cases of giving good reasons for accepting or rejecting P are cases of giving evidentiary reasons.

Consider this case: You are a prisoner of a terrible tyrant, King Mog, who is known always to keep his promises. He tells you that tomorrow you will be given a lie detector test to determine whether you believe M: Mog is a kindly, enlightened monarch. If you do, you will be given \$10,000 and released unharmed. If you don't, you will be skinned alive and salted. Furthermore, a harmless, non-addictive drug will be offered to you, and if you take it while listening to a recording praising King Mog's kindness, you will wake up believing M. Mog argues to you as follows: (i) If you believe M, which you can easily do, you get a reward. (ii) If you don't believe M, then you will suffer horribly. (iii) Therefore, M. Has Mog given you good reasons for believing M? He has given you very good motivation for accepting M, but wretched evidence. I would mark this by saying that he has given good motivational reasons but bad evidentiary reasons. This is characteristic of fallacious appeals to emotion – in this case, the fallacy of *argumentum ad baculum* or appeal to force. Note that if the conclusion M were replaced with “You ought to accept M” then the reasons offered are good both motivationally and evidentially.

The fact that P was selected by a process likely to pick falsehoods is not a merely motivating reason as in the above appeal to force. But it is not evidence of the falsity of P either. Given the going loose use of ‘evidence’ it would be counted, but by a proper standard it should be otherwise named, say, as ‘probabilistic reasons.’ Just as the *argumentum ad baculum* has (typically, when fallacious) motivating reasons posing as evidence, the *argumentum ad hominem*, in its most interesting and formidable form, typically has probabilistic reasons posing as evidence.

Of course a good reasoner may find good evidentiary reasons to be good motivating reasons and also good probabilistic reasons. And just as moving from P

to “You had better accept P” can make the reasons which were good motivation but bad evidence into adequate evidence, so too changing from P to “It is highly likely that P” can make reasons which were good probabilistically but bad evidence into adequate evidence.

Thus in pattern I, the fact that P has been advanced by someone unlikely to be right about such questions is good reason to doubt that P is true. It is just not evidence for the falsity of P. The *ad hominem* fallacy confuses these two, so that the considerations which make it unlikely that P are taken to be considerations against P. It is easier to see this mistake in the above cases because we do not even know what proposition P is. We just know it is some proposition picked by a method which is highly likely to pick falsehoods.

However, this can become blurred as we gradually get some understanding of P. “Knowing what proposition P is” is very vague. One knows it is the proposition whose number was drawn from the urn. Or one may know that it is a proposition about a certain person, to the effect that that person grows melons. Or one may know who that person is, etc. It is easy to get the illusion of understanding the proposition and having evidence it is false when in fact one has nothing but probabilistic considerations.

“To conclude that P” is vague, but one good reading is in terms of claiming to know on the basis of premises. To claim to know that P is inconsistent with allowing that there is a nonzero chance that not-P. If you conclude that someone's claim that P, is unjustified, then you are (in the relevant uses) claiming to know that his claim that P, is unjustified. This cannot be done merely on the basis of probability. If you admit that there is any nonzero chance that his claim is justified then you cannot consistently claim to know it is not.

If we understand ‘evidence’ as recommended above (and it must be stressed that this is not an appeal to correct ordinary usage but rather, a correction of ordinary usage) then we can clarify the fallaciousness of pattern I(i) by distinguishing further between the conclusion that (ia) ‘we have evidence showing that A has not produced argument sufficient to justify his claim that P’ and that (ib) ‘it is likely that A's claim is not justified’. Iia should be seen as a kind of *ad hominem* fallacy, the justification *ad hominem*.

This point can be extended. Not only is Iia fallacious. It is fallacious to conclude from the premises of I even that you have any evidence undermining the claim that P, even slightly. One might put this by saying that it is fallacious to think that considerations showing that it is unlikely that P constitute any kind of defeater for the claim that P, except that the term ‘defeater,’ which has in its natural use the proper suggestion of a counterargument or evidence to the

contrary, has been appropriated as a technical term which can be used on the basis of merely probabilistic considerations.

In this connection, one might be tempted to appeal to a distinction between asserting the negation of a proposition and asserting that the proposition is false. When you do not know what the proposition is and are unable to present it, you can still assert it is false (as in “What he said yesterday is false”) but you cannot assert the negation (as in “It is false that his program will reduce unemployment”). However, this is not helpful. You may ‘know’ that the proposition was that Wong grows melons in Daigan and have excellent probabilistic reasons for asserting the negation, that Wong does not grow melons in Daigan while still being utterly lacking in evidence for or against what you are asserting. Whatever logical difference may be found between asserting that a proposition is false and asserting its negation, or epistemological difference, will not be of use here. The right guide is our rule that evidence that the source of a claim is likely to be wrong is not evidence against the claim. The tendency to overlook this is the essential feature of the *ad hominem* fallacy.

It is a prominent feature of Descartes' philosophy to hold that we cannot know anything about the external world unless we know that (G) we are created by a benevolent God who has given us senses which are trustworthy if used with caution we are capable of exercising. I believe that G is true. But it would be a bad mistake to think that someone who does not believe that G and thus does not know that G could therefore be shown not to know anything about the world. Even worse would be the conclusion that someone who regards his thinking as made possible by Godless arrangements of matter largely determined by chance is, in taking this attitude, advancing a hypothesis which undermines his theorizing about the world or himself. Theories to the effect that you are a creature of a kind very unlikely to produce a good theory have been advanced both on theistic and atheistic grounds. In any case, the fact that such a theory would assign a low probability to the claim that its author has produced a good theory is not, on the proper understanding of evidence, evidence against the theory.

It is true that evidence that someone is utterly unreliable can be a reason for denying that he knows something. Suppose that an extremely unstable alcoholic mathematical genius, Jones, proves a brilliant theorem, say, that (H) all hypersimple manifolds are superousian. We might well have good reason for the verdict that Jones does not know that H. Jones might frequently deny H, or defend H by appeal to certain irrelevant facts about patterns in which his empty bottles break against the wall. We may find it nearly incredible that Jones got himself together enough to produce the proof. This sort of thing is indeed relevant to the

question whether Jones knows that H. But it is a bad confusion to think it counts as any kind of objection to Jones' argument for H. That Jones thinks that H could be a terrible reason for H. To advance it as a reason would be an *argumentum ad verecundiam*, an argument we can correctly reject. But that should not be confused with evidence against H, or we will slip into a fallacious *ad hominem*.

We can compare Jones' case with that of the earlier reputed 'idiot' who produces a brilliant argument. In either case, that argument is evidence that the arguer is capable of cleverness, is not altogether incompetent. The presentation of the argument tends to make this evident. When the 'idiot' was arguing that he is mentally incompetent, his reasoning was failing to make it evident that he is incompetent. The display of reasoning was in fact making it evident he has at least some competence.

Our model of merely probabilistic reasons was based on a random drawing model of probability – an argument drawn from an urn containing 10,000 proven bad arguments and one untested one. The argument that the evolution of human argumentative competence was highly unlikely would not plausibly be based on such a model. This is a limitation on charges of *ad hominem* reasoning in accordance with our account of such reasoning. That is only due to limitations on judgments of probability. Whether, for example, the development of human reasoning ability in a naturalistic evolutionary process has some probability, low or otherwise, is far less clear than the simplified probabilities based on urn drawings. Hopefully such unclarity in the notion of probability does not undermine the point that, insofar as we have a distinction between probabilistic and evidentiary reasons, it is fallacious to confuse them.