Advanced Reasoning Forum

Proceeding of the 2nd Annual Conference on Logic and Reasoning

New Europe College
Bucharest, Romania
July 2000

Bulletin of Advanced Reasoning and Knowledge
Volume 1, 2001
Edited by Richard L. Epstein

New Europe College, Bucharest

Arguments and Explanations

Richard L. Epstein[1]

Reasoning is the art of saying "therefore": How to determine whether the truth of one claim follows from the truth of some others. There are many different ways to say "therefore." In this paper the relation between two of those, arguments and explanations, will be discussed.

To say the truth of one claim follows from the truth of others, we first need to be clear what claims are.

Claim A *claim* is a declarative sentence that is used in such a way that it is either true or false, but not both.

Claims are the smallest parts of spoken or written language with which we reason that we can call "true" or "false." This is not to deny that there are abstract objects—propositions—that are the real things that are true or false. If there are, then perhaps claims are just our imperfect representations of those. But claims are what we use in reasoning.

In trying to understand what we mean by saying one claim follows from others, it's best to consider first the archetype of saying "therefore."

Argument An *argument* is a collection of claims; one is called the *conclusion,* whose truth the argument is intended to establish. The others, called the *premises,* are meant to lead to, or support, or convince that the conclusion is true.

Arguments are attempts to convince, whether someone tries to convince you, or you try to convince someone else, or you try to convince yourself. But that does not mean the criterion for whether an argument is good is whether the argument actually does convince. If I am drunk, you may give me an excellent argument that my driving home is dangerous; though I remain unconvinced, the argument is no worse. A politician may make a bad argument that you should vote for him, but just because you are convinced does not make the argument good. Perhaps other ways to convince, such as entreaties, exhortations, sermons, and advertisements, can be judged by how well they convince, but attempts to establish the truth of a claim cannot be so judged.

[1] This is a revised version of a talk given at the second Advanced Reasoning Forum meeting sponsored by New Europe College in Bucharest, Romania. Many of the ideas and the exposition were refined over the last two years through discussions with the members of the Advanced Reasoning Forum. A much fuller account of this work can be found in my *Five Ways of Saying 'Therefore'*. A textbook version of these ideas can be found in my *Critical Thinking,* and in the *Science Workbook* for that text. I am grateful to Alex Raffi for providing the illustrations.

Bulletin of Advanced Reasoning and Knowledge 1 (2001) 3–17.

Good argument An argument is *good* if its premises give good reason to believe the conclusion is true.

This seems a very imprecise standard, if any standard at all. What does "good reason" mean?

First, note that from a false premise we can derive a true claim or a false claim. For example,

(1) All authors of books on arguments are women.
Richard L. Epstein is an author of a book on arguments.
Therefore, Richard L. Epstein is a woman.

(2) All authors of books on arguments are women.
Richard L. Epstein is an author of a book on arguments.
Therefore, Richard L. Epstein is a human being.

The first of these has a false conclusion, the second a true one.

So it would seem that a good argument should have true premises. But that is too strong a condition. Consider:

There are an even number of stars in the sky.
So the number of stars in the sky can be divided by 2.

There are an odd number of stars in the sky.
So the number of stars in the sky cannot be divided by 2.

One of these has a true premise, but we cannot tell which. A standard that gives us no way to evaluate arguments is not part of the art of reasoning. Rather, for an argument to be good, we must have good reason to believe its premises; that is, the premises must be ***plausible***.

What counts as good reason to believe a premise? That is not for the logician to say. It may depend in part on the subject matter: The biologist, the car mechanic, the professional football player, the physicist all have their own standards. It may depend in part on the metaphysics we adopt: Some say we never have good reason to believe claims whose truth cannot be ascertained through empirical tests. At best, we can give some rough standards for when to accept unsupported claims, which I present in *Critical Thinking*.

So a good argument should have plausible premises. But more is needed. Consider:

(3) Richard L. Epstein speaks English.
Richard L. Epstein wrote *Critical Thinking*.
Therefore, Richard L. Epstein lives in the United States.

Each of these claims is highly plausible. But the conclusion does not follow from the premises. What, though, do we mean by saying that the conclusion follows from the premises? There are two different standards.

Valid argument An argument is *valid* if it is impossible for the premises to be true and conclusion false (at the same time).

For example,

(4) Maria is a widow.
So Maria was married.

It is not possible for the premise of this argument to be true and conclusion false.

In our daily lives, however, we often cannot employ valid arguments. For example, Dick hears that there are parakeets for sale at the mall. He knows that his neighbor has an old birdcage in her garage. He wonders whether the cage will be big enough for a parakeet if he buys one. He reasons:

(5) Every parakeet I or anyone I know has seen, or read, or heard about is less than 60 cm tall.
Therefore, the parakeets on sale at the mall are less than 60 cm tall.

This argument is not valid. There could be a new kind of parakeet discovered in the Amazon that is 1 meter tall; or a new supergrow bird food has been developed that makes parakeets grow very tall; or aliens have captured some parakeets and shot them with rays to make them very large; or . . . But any possibility that we can think of for the premise to be true and conclusion false is very unlikely—so unlikely that Dick has good reason to believe the conclusion.

The Scale from Strong to Weak We classify invalid arguments on a scale from *strong* to *weak*. An argument is *strong* if it is very unlikely that the premises could be true and conclusion false (at the same time). An argument is *weak* if it is not unlikely that the premises could be true and conclusion false (at the same time).

Here "very unlikely" means relative to what we know.

In reasoning in our lives, and in almost every area of science, we rely on strong arguments that we cannot replace with valid arguments. For example, replacing the premise of (5) with "All parakeets are less than 60 cm tall" would yield a worse argument, for that claim is less plausible than the premise of (5). Indeed, there is often a trade-off between how plausible the premises of an argument are and how strong the argument is: the less plausible the premises, the stronger the argument.

So for an argument to be good, it should be either valid or strong. Argument (3), for example, is bad because it is weak: It is not unlikely that I could speak English and write *Critical Thinking* while living in Barbados.

The standard of whether an argument is strong is subjective. But it is not unusable, any more than the judgment of whether an auditorium lit at one end by a candle is dark at one end and light at the other. In practice, we almost always agree on the evaluation of the strength of an argument once we exchange the background assumptions on which we base the evaluation.

In any case, it is not just the criterion of whether an argument is strong that introduces a subjective element into judging whether an argument is good. The question of whether the premises of an argument are plausible is also subjective. Even the classification of a collection of claims as an argument is subjective. It may be that there are collections of abstract propositions that satisfy objective criteria that our notion of strong argument imperfectly attempts to describe. But those abstract objects will be of little use to us in our reasoning.

In order for an argument to be good, it is not enough for it to be valid or strong and have plausible premises. Consider:

> Every dog has a soul.
> Therefore, dogs should be treated humanely.

Even if you find the premise plausible, it is not more plausible than the conclusion. For an argument to give us good reason to believe the conclusion, its premises should be more plausible than its conclusion.

Begging the question Any argument that uses a premise that is not more plausible than the conclusion is said to *beg the question.*

Many valid arguments beg the question, for example (4) above. But not every valid argument begs the question. For example, "George is a duck. If George is a duck, then George quacks. So George quacks."

We now have three tests that an argument must pass in order to be good.

Necessary Conditions for an Argument to be Good
- The premises are plausible.
- The argument is valid or strong.
- The argument does not beg the question.

In our daily lives there are many arguments that we classify as good that do not seem to satisfy these conditions. For example,

> Lee: Tom wants to get a dog.
> Maria: What kind?
> Lee: A dachshund. And that's really stupid, since he wants one that will chase a frisbee.

Lee has made an argument, if we interpret rightly what he said: Tom wants a dog that will chase a frisbee, so Tom shouldn't get a dachshund. But on the face of it that argument is not strong or valid. Still, Maria knows very well, as do we, that a dachshund would be a bad choice for someone who wants a dog to chase a frisbee. Dachshunds are too low to the ground, they can't run fast, they can't jump, and the frisbee is bigger than they are, so they couldn't bring it back. Any dog like that is a bad choice for a frisbee partner. Lee just left out these obvious claims, but why should he bother to say them?

We normally leave out so much that if we look only at what is said, we will be missing too much. We can and must rewrite many arguments by adding an unstated premise or even an unstated conclusion.

When are we justified in adding an unstated premise? How do we know whether we've rewritten an argument well or just added our own prejudices? To "repair" arguments that are apparently defective, we must have some standards. Otherwise we will end up putting words in someone's mouth. Such standards depend on what we can assume about the person with whom we are reasoning or whose work we are reading.

The Principle of Rational Discussion We assume that the other person with whom we are deliberating or whose arguments we are reading:

- Knows about the subject under discussion.
- Is able to reason well.
- Is willing to reason well.
- Is not lying.

Often people with whom we wish to reason do not satisfy these conditions. But then it makes no sense to reason with that person. We should be educating, or consoling, or pointing out the errors in what he or she says.

The Principle of Rational Discussion justifies adopting the following guide for when we can add or delete claims in an argument.

The Guide to Repairing Arguments Given an (implicit) argument that is apparently defective, we are justified in *adding* a premise or conclusion if and only if all the following hold:

- The argument becomes strong or valid.
- The premise is plausible and would seem plausible to the other person.
- The premise is more plausible than the conclusion.

If the argument is valid or strong, yet one of the original premises is false or dubious, we may *delete* that premise if the argument becomes no worse.

Given only this Guide, we might try to repair every argument into a good one. That would be wrong, for there are clear standards for when an argument is unrepairable.

Unrepairable Arguments We cannot repair a (purported) argument if any of the following hold:

- There is no argument there.
- The argument is so lacking in coherence there's nothing obvious to add.
- A premise is false or dubious or several premises are contradictory and cannot be deleted.
- The obvious premise to add would make the argument weak.
- Any obvious premise to add to make the argument strong or valid is false.
- The conclusion is clearly false.

It is not that when we encounter one of these conditions we can be sure the speaker had no good argument in mind. Rather, we are not justified in making that argument for him or her—it would be putting words in the speaker's mouth.

In addition to these conditions for an argument to be unrepairable, a list of other kinds of arguments, called fallacies, have been deemed to be typically so bad that they, too, are rejected as unrepairable when we encounter them.[2]

In *Critical Thinking* I set out how the Principle of Rational Discussion justifies adopting these guides, and I give many examples how these standards are useful in evaluating arguments, from arguments in our daily lives to those we encounter in science journals. But one particular kind of argument that is related to explanations seems to require different standards of evaluation.

Generalization An argument in which we conclude a claim about a group, called the *population,* from a claim or claims about some part of it, called the *sample.* Sometimes the conclusion of the argument is called a "generalization."

For example, the following is a generalization: "All dogs that I have ever met except for one can bark; so almost all dogs bark." Some standards for a generalization to be good are the following.

Necessary Conditions for a Generalization to be Good A good generalization requires as premises the following three claims (whether stated not):

1. The sample is representative.
2. The sample is big enough.
3. The sample is studied well.

But these are not different standards than the necessary conditions for an argument to be good. They only spell out in more detail what is required for the argument to be strong and have plausible premises. Still, it remains open whether the necessary conditions for an argument to be good are also sufficient. That is a difficult issue to resolve, and takes up much of *Five Ways of Saying "Therefore."*

Finally, we need to note one particular mistake in reasoning with arguments, for it is surprisingly common and is related to the use of explanations. Some people, when encountering an argument that is valid or strong which has a clearly true conclusion, conclude that the premises are plausible. But that's wrong, as argument (2) shows.

Arguing backwards To conclude that the premises of an argument are true because the argument is valid or strong and its conclusion is highly plausible.

What in the study of arguments can we carry over to all other reasoning? It cannot be the requirement that the premises be plausible. Proofs in mathematics are often acceptable even when that is not the case.

It cannot be that the premises are more plausible than the conclusion. A good explanation of why the sky is blue will certainly use claims more dubious than "The sky is blue."

What must be analyzed in any reasoning is the relationship of the premises to the conclusion.

---

[2] See *Five Ways of Saying "Therefore"* or *Critical Thinking.*

Inference A collection of claims, one of which is designated the *conclusion* and the others the *premises,* that is intended to be judged as valid/invalid or on the scale from strong to weak.

An inference is *valid* if it is impossible for the premises to be true and conclusion false (at the same time). An inference is *strong* if it is very unlikely that the premises could be true and conclusion false (at the same time). An inference is *weak* if it is not unlikely that the premises could be true and conclusion false (at the same time).

An argument is an inference that is meant to convince that the conclusion is true. This is "therefore$_{\text{argument}}$".

Many explanations are inferences, too. Consider:

> Why is the sky blue? Because sunlight is refracted through the atmosphere so as to absorb other wavelengths of light.

"The sky is blue" is explained in terms of why it is true—what it follows from, the reasons for its truth.

Not every explanation can be understood as an inference. For example, when a traveler asks a policeman to explain how to get to the Post Office, she's not asking him to show why some claim is true.

For an explanation that can be judged as an inference, the conclusion, what's being explained, should be highly plausible: We can't explain why the sun rises in the west. The claims that do the explaining should be plausible, too. We don't accept "The sky is blue, because there are blue globules high in the atmosphere" as a good explanation, because "There are blue globules high in the atmosphere" is known to be false. But for an explanation the premises shouldn't be more plausible than the conclusion, for otherwise we'd have an argument.

Moreover, the inference should be valid or strong. We don't accept "Dogs lick their owners because they aren't cats" because the inference is neither valid nor strong and there is no obvious way to repair it. As with arguments, we allow that the inference might be repaired: We understand that an explanation "E because of A" may require further premises to supplement A. And as with arguments, we can invoke the Principle of Rational Discussion to motivate the Guide to Repairing Arguments, except for the condition that the premises be more plausible than the conclusion.

Inferential explanations An explanation that can be judged as an inference, "E because of A, B, C, . . .". For it to be good, all the following must hold:

1. E is highly plausible.
2. Each of A, B, C, . . . is plausible, but at least one of them is not more plausible than E.
3. "A, B, C, . . . therefore E" is a valid or strong inference, possibly with respect to some plausible unstated claims.
4. The explanation is not "E because of D" where D is E itself or a simple rewriting of E.

We call E the *explanandum* and collectively A, B, C, . . . the *explanans.* Sometimes the explanans alone is called "the explanation."

It is often said that an explanation must answer the right question. For example,

Mother: There were two pieces of cake in the cupboard. Why is there only one now?
Flo: Because it was dark and I didn't see the other piece.

Flo thinks she has given a good explanation: Her answer makes it clear why the claim "There is only one piece of cake in the cupboard now" is true (assuming some other fairly obvious claims). But her mother won't accept it. Flo answered "Why is there only one piece of cake in the cupboard, instead of none?" but her mother meant, "Why is there only one piece of cake in the cupboard, instead of two?"

Questions are often ambiguous, and a good explanation to one reading of a question can often be a bad explanation to another. If the explanandum is ambiguous, then that is a fault of the questioner; we should not be expected to guess correctly what's meant. Still, when it is clear that a different reading of the explanandum is meant, we can say that an explanation is bad: It's answered the wrong question.

How are explanations and arguments related? A good explanation is not a good argument. Consider what Zoe said to Dick Sunday morning:

> You drank three cocktails before dinner, a bottle and a half of wine with dinner, and then a couple of glasses of brandy. Anyone who drinks that much is going to get a headache. So you have a headache.

Zoe has given a good explanation of why Dick has a headache. But it is a bad argument, because it begs the question: It is much more plausible to Dick that he has a headache than that anyone who drinks that much is going to get a headache. With a little rewriting, this shows that aristotelian syllogisms that are faulted as begging the question are often perfectly good attempts to codify or explain.

For the relation of particular explanations to arguments, consider the following. Dick, Zoe, and Spot are out for a walk in the countryside. Spot runs off and returns after five minutes. Dick and Zoe notice that Spot has blood on his muzzle. And they both really notice that Spot stinks like a skunk. Dick turns to Zoe and says, "Spot must have killed a skunk. Look at the blood on his muzzle. And he smells like a skunk." Dick has made a good argument:

> Spot has blood on his muzzle. Spot smells like a skunk.
> Therefore$_{\text{argument}}$ Spot killed a skunk.

He's left out some premises that he knows are as obvious to Zoe as to him:

> Spot isn't bleeding.
> Skunks aren't able to fight back very well.

Normally when Spot draws a lot of blood from an animal that is smaller than him, he kills it.
Only skunks give off a characteristic odor, an odor that drenches whoever or whatever is near if they are attacked.
Dogs kill animals by biting them and typically drawing blood.

Zoe replies, "Oh, that explains why he's got blood on his muzzle and smells so bad." That is, Zoe takes the same claims and uses them to make a good explanation, relative to the same unstated premises:

Spot killed a skunk.
Therefore$_{\text{explanation}}$ Spot has blood on his muzzle and smells like a skunk.

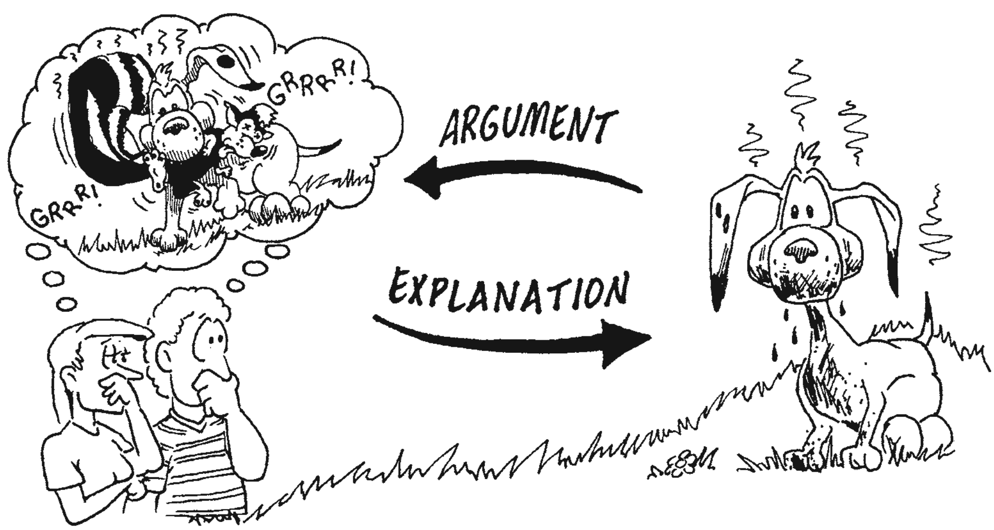

Explanations and associated arguments Given an explanation:

A therefore$_{\text{explanation}}$ E (relative to some other premises P, Q, R, . . .)

the *associated argument* is:

E therefore$_{\text{argument}}$ A (again relative to P, Q, R, . . .)

For an explanation with many claims in the explanans, Σ therefore$_{\text{explanation}}$ E, reversing the role of E with any one of the claims in Σ is an associated argument.

An *independent* explanation is one in which each premise that is less plausible than the explanandum can be established by an associated argument. An explanation is *dependent* if the reason to believe at least one of the premises must be established by claims outside the explanation.

Zoe's explanation of why Spot has blood on his muzzle and smells bad is independent, because the associated argument that Dick makes is good: It establishes that Spot killed a skunk.

But many explanations are dependent. Consider what Dick told his neighbor's little girl:

Spot chases cats because he sees cats as something good to eat and because cats are smaller than him.

Are the claims in the explanans plausible? Certainly "Cats are smaller than Spot" is plausible. But how about "Spot sees cats as something good to eat"? What reason do we have to accept this? The associated argument for it is:

> Spot chases cats, and cats are smaller than Spot.
> Therefore, Spot sees cats as something good to eat.

This is weak. So the explanation is dependent. Without more evidence for "Spot sees cats as something good to eat," we shouldn't accept the explanation.

Any explanation with a generalization in the explanans is likely to be dependent, too. For example,

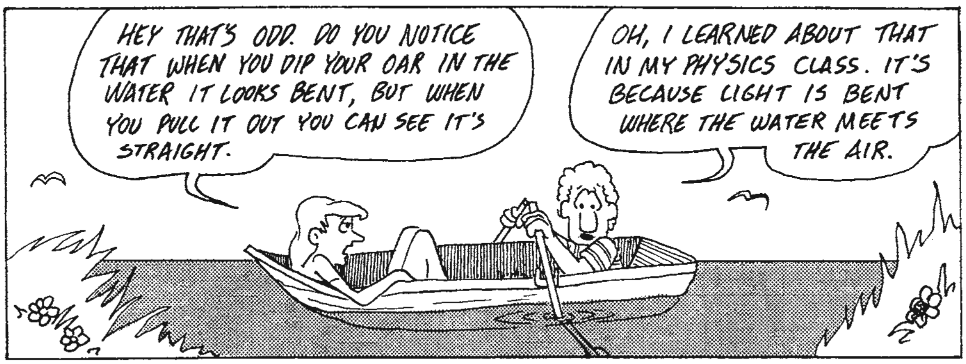

Dick offers an explanation:

> The oar appears bent because light is bent where the water meets the air.

The explanans here is a generalization: "Light is bent where the water meets the air." Zoe hasn't taken a physics course, so the only reason she has to believe that claim is the associated argument:

> The oar appears bent, therefore$_{\text{argument}}$ light is bent where the water meets the air.

But this is a weak generalization, needing more examples to convince. So Dick has given a dependent explanation. Dick has other reasons to believe the claim from his studies in physics, which he can offer to Zoe to make this explanation good.

Testing explanations is often how we establish a generalization. For example, consider what the children Flo and Becky were saying last week:

> Flo: Spot barks. And Wanda's dog Ralph barks. And Dr. E's dogs Anubis and Juney bark. So all dogs bark.
>
> Becky: Yeah. Let's go over to Maple Street and see if all the dogs there bark, too.

Flo is generalizing. Relative to her experience it's a pretty good generalization. Becky wants to test the generalization.

Suppose that A, B, C, D are given as inductive evidence for a generalization G. (Some other highly plausible unstated premises may also be needed, but we'll keep those in the background.) Then we have that G explains A, B, C, D.

But if G is true, we can see that some other claims must be true, instances of the generalization G, say L, M, N. If those are true, then G would explain them, too. For example, Rodolfo barks, Lady barks, Fido barks, . . .

That is, G *explains* A, B, C, D and *predicts* L, M, N, where the difference in this case between the explanation and the prediction is that in the explanation we know the conclusion is true, whereas we don't know if the predictions are true.

Suppose we find that L, M, N are indeed true. Then the argument "A, B, C, D + L, M, N therefore G" is a better argument for G than we had before. At the very least it has more instances of the generalization as premises.

But how do more instances of a generalization prove the generalization better? They can if (*i*) they are from different kinds of situations, that is, A, B, C, D + L, M, N cover a more representative sample of possible instances of G than do just A, B, C, D. This is typically what we do: We deduce claims from G for situations that we had not previously considered.

And (*ii*) because we had not previously considered the kind of instances L, M, N of the generalization G, we have some confidence that we haven't got G by manipulating the data, selecting situations that would establish just this hypothesis.

The best way to test an hypothesis-generalization, it's often said, is to try to falsify it. Trying to falsify the generalization just means that we are consciously trying to come up with instances of the generalization to test that are as different as we can imagine from A, B, C, D. Trying to falsify is a good way to ensure (*i*) and (*ii*). So we say that an experiment ***confirms*** the explanans if it shows that a prediction is true. *Confirmation amounts to strengthening the associated argument.*

Here is an example of this relation between explanation and prediction:

> Consider the explanation offered by Torricelli for a fact that had intrigued his teacher Galileo; namely, that a lift pump drawing water from a well will not raise the water more than about 34 feet above the surface of a well. To account for this, Torricelli advanced the idea that the air above the water has weight and thus exerts pressure on the water in the well, forcing it up the pump barrel when the piston is raised, for there is no air inside to balance the outside pressure. On this assumption the water can rise only to the point where its pressure on the surface of the well equals the pressure of the outside air on that surface, and the latter will therefore equal that of a water column about 34 feet high.[3]

That is, Torricelli offered an explanation, but the only evidence he had for the explanans, which was a generalization, was the explanandum.

> The explanatory force of this account hinges on the conception that the earth is surrounded by a "sea of air" that conforms to the basic laws governing the equilibrium of liquids in communicating vessels. And because Torricelli's explanation presupposed such general laws it yielded predictions concerning as yet unexamined phenomena. One of these was that if the

[3] Carl G. Hempel, *Aspects of Scientific Explanation,* p. 365.

> water were replaced by mercury, whose specific gravity is about 14 times that of water, the air should counterbalance a column about 34/14 feet, or somewhat less than $2^1/_2$ feet, in length. This prediction was confirmed by Torricelli in the classic experiment that bears his name. In addition, the proposed explanation implies that at increasing altitudes above sea level, the length of the mercury column supported by air pressure should decrease because the weight of the counterbalancing air decreases. A careful test of this prediction was performed at the suggestion of Pascal only a few years after Torricelli had offered his explanation: Pascal's brother-in-law carried a mercury barometer (i.e., essentially a mercury column counterbalanced by the air pressure) to the top of the Puy-de-Dôme, measuring the length of the column at various elevations during the ascent and again during the descent; the readings were in splendid accord with the prediction. [4]

Predictions are made of further instances of the generalization or of consequences of the claim in the explanans; those are shown to be true; the explanans thus becomes more plausible because the associated argument for it (adding as premises all the instances of the generalization that have been tested and found to be true) is strengthened. The story is much the same for explanans that aren't generalizations, too.

This relation of arguments and explanations is often misunderstood and used badly. Some scientists believe that if you have an explanans that could explain a lot, it must be true. For example, Charles Darwin said:

> It can hardly be supposed that a false theory would explain, in so satisfactory a manner as does the theory of natural selection, the several large classes of facts above specified [the geographical distribution of species, the existence of vestigial organs in animals, etc.]. It has recently been objected that this is an unsafe method of arguing; but it is a method used in judging of the common events of life, and has often been used by the greatest natural philosophers.[5]

Darwin was arguing backwards: From the truth of the conclusion(s), we can infer the truth of the premises. The direction of inference is incorrect. Rather, the conclusions that are drawn from the explanans together serve as evidence for the explanans, not because the explanans gives the best explanation of them, but because they are premises for an argument concluding with the explanans.

Gilbert Harman, however, thinks Darwin's method of arguing is right if the explanation is the best:

> In making this inference one infers, from the fact that a certain hypothesis would explain the evidence, to the truth of that hypothesis. In general, there will be several hypotheses which might explain the evidence, so one must be able to reject all such alternative hypotheses before one is warranted in making an inference. Thus one infers, from the premise that a given

---

[4] Ibid.

[5] *On the Origin of Species,* p. 476. Paul R. Thagard presents further examples where scientists reasoned in this way in "The best explanation: criteria for theory choice"

hypothesis would provide a "better" explanation for the evidence than would any other hypothesis, to the conclusion that the given hypothesis is true.[6]

This method of argument he calls *inference to the best explanation.*

But even if we had clear criteria for what we mean by "best explanation," which we certainly don't, it wouldn't rescue inference to the best explanation from the charge of arguing backwards. Consider what happened to me a few years ago:

(In the hospital emergency room at 2 a.m.)
Me: Doctor, doctor, why do I have such pain in my back? It doesn't feel like a muscle cramp or a pinched nerve.
Doctor: (after examining me) A kidney stone would explain the pain. Kidney stones give that kind of pain, and it's in the right place for that. I can't think of anything else that would give you that much pain there.

The doctor offered me the best explanation he had:

Your back hurts this way because you have a kidney stone.

This would have been a good explanation if we'd had good reason to believe the explanans. But at that point the only evidence for the explanans was the associated argument, and that was not strong.

So the doctor made predictions from the explanans: "A kidney stone would show up on an X-ray," "You would have an elevated white-blood cell count," "You would have blood in your urine," "A kidney stone will show up on a CAT-scan." He tested each of these and found them false. He reasoned by *reductio ad absurdum* that if the explanans were true, one or more of these would be true; they are false; therefore, the explanans is very likely false. Hence his original explanation turned out to be bad. Nothing else was found, so by process of elimination it was conjectured that I had a severe sprain or strain, for which exercise and education were the only remedy.

If inference to the best explanation were a good method of argument, there would have been no point in doing tests. The doctor and I should have believed "You have a kidney stone." But I'm glad we didn't, or I'd have undergone needless surgery.

Still, you might say that in this example the explanation the doctor offered was not good, since "You have a kidney stone" was not plausible. So we didn't have reason to believe it. Inference to the best explanation doesn't fail here.

But if we require in an inference to the best explanation that the explanation be not only better than all the others but good as well, then we don't need inference to the best explanation. The explanans of a good explanation is plausible, perhaps because of the associated argument, or another argument, or just by inspection. We have good reason to believe it. But the explanans of the best explanation need only be more plausible than any (or all) of the explanans of

6 "The inference to the best explanation", p. 89.

the other explanations we can think of. Ruling out all but one explanation does not by itself show that that one is plausible. It could equally show how bad our imagination is. We need further evidence.

Finding an explanation that is better than all others does not justify belief in the explanans. It only provides motive for us to investigate whether the proposed explanans is true, which is exactly how the doctor saw it. Used that way, inference to the best explanation is called *abduction.*[7]

Fallacy of inference to the best explanation An argument "These claims give the best explanation, so they are true."

Instead of "That's the best explanation we have, so it's true," think "That's the best explanation we have, so let's investigate it." Instead of "That's the only explanation of . . ." say "All the evidence points to . . .".

Inference to the best explanation is what stands behind many claims that there are numbers and sets, abstract objects of that sort. Mathematics and science "need" them, in the sense that they best explain why our mathematical and scientific theories are true. Certainly if those objects exist, they explain that. But, as in any use of inference to the best explanation, we must ask what other evidence we have to believe the claim "Numbers, as abstract objects, exist," since the inference from these scientific theories being "true" to numbers and sets existing is weak.[8]

Inference to the best explanation is no better in mathematics than in daily life. The difference, it seems, is that in mathematics there is no other evidence we can cite for "Numbers, as abstract objects, exist." Mathematics, for the platonist, is built on faith; and the necessity of numbers for mathematics—all numbers, natural, rational, real, in their abstract plentitude—is a guide, a sign towards that faith.

---

[7] C. S. Peirce saw this clearly (Vol. 5.189 and Vol. 5.171):

> The surprising fact, C, is observed.
> But if A were true, C would be a matter of course.
> Hence, there is reason to suspect that A is true.
>
> Deduction proves that something *must* be; Induction shows that something *actually is* operative; Abduction merely suggests that something *may be.*

[8] W. D. Hart, Introduction to *The Philosophy of Mathematics,* p. 6:

> Sophisticated natural science as it comes is always to be formalized in an extension, in the logician's sense, of some mathematics, often number theory and analysis. Equations are obvious to anyone reading serious science. So, by abduction [inference to the best explanation], we are justified in believing true at least as much mathematics as we need for the best scientific explanations of what we observe. Since the truth of that much mathematics requires very abstract objects, Quine thereby began an empiricist justification for belief in the abstract objects required for mathematical truth.

Conclusion

Arguments are different from explanations. By carefully distinguishing the criteria for what counts as a good argument from what counts as a good explanation, we can see that arguments and explanations are related. It is not inference to the best explanation that turns explanations into arguments, but simply reversing the roles of premise and conclusion, which accounts for the notion of confirmation of an hypothesis. To argue well, we must be able to distinguish different ways of saying "therefore."

Bibliography

Darwin, Charles

1872 *On the Origin of Species,* Sixth edition. Collier, 1962.

Epstein, Richard L.

2001 *Five Ways of Saying "Therefore"*, Wadsworth.

2001 *Critical Thinking,* 2nd edition, Wadsworth.

Harman, Gilbert

1965 The inference to the best explanation, *Philosophical Review,* vol. 74, pp. 88–95.

Hart, W. D. , ed.

1996 *The Philosophy of Mathematics,* Oxford University Press.

Hempel, Carl G.

1965 *Aspects of Scientific Explanation,* The Free Press.

Peirce, Charles

1960 *Collected Papers,* vol. 5 and vol. 6, eds. Charles Hartshorne and Paul Weiss, The Belknap Press, Harvard.

Thagard, Paul R.

1978 The best explanation: criteria for theory choice, *The Journal of Philosophy,* vol. 75, pp. 76–92.