Cause and Effect, Conditionals, Explanations

Essays on Logic as *The Art of Reasoning Well*

Richard L. Epstein

Cause and Effect, Conditionals, Explanations

*Essays on Logic as
the Art of Reasoning Well*

Richard L. Epstein

Illustrations by Alex Raffi

ARF

Advanced Reasoning Forum

Cause and Effect, Conditionals, Explanations

Richard L. Epstein

Preface

Essays on Logic as The Art of Reasoning Well

Contents of Other Volumes in the Series

Preface

This series of books presents the fundamentals of reasoning well, in a style accessible to both students and scholars. The text of each essay presents a story, the main line of development of the ideas, while the footnotes and appendices place the research within a larger scholarly context. The essays overlap, forming a unified analysis of reasoning, yet each essay is designed so that it may be read independently of the others. The topic of this volume is the evaluation of reasoning about cause and effect, reasoning using conditionals, and reasoning that involves explanations.

The first essay summarizes material that can be useful as background from *The Fundamentals of Argument Analysis* in this series.

The essay "Reasoning about Cause and Effect" sets out a way to analyze whether there is cause and effect in terms of whether an inference from a claim describing the purported cause to a claim describing the purported effect satisfies specific conditions. Different notions of cause and effect correspond to placing different conditions on what counts as a good causal inference. An application of that method in "The Directedness of Emotions" leads to a clearer understanding of the issue whether every emotion need be directed at something.

In the essay "Conditionals" various ways of analyzing reasoning with claims of the form "if . . . then . . ." are surveyed. Some of those uses are meant to be judged as inferences that are not necessarily valid, and conditions are given for when we can consider such an inference to be good.

In "Explanations" verbal answers to a question why a claim is true are evaluated in terms of conditions placed on inferences from the explaining claims to the claim being explained. Recognizing that the direction of inference of such an explanation is the reverse of that for an argument with the very same claims is crucial in their evaluation. Explanations in terms of functions and goals are also investigated.

In a companion volume *Reasoning in Science and Mathematics* causal reasoning and explanations are connected to the use of models and theories in science. In particular the nature of causal laws is

discussed in the context of theories. The study of laws and confirmation is also analyzed in the essay "Generalizing" in *The Fundamentals of Argument Analysis* in this series.

Reasoning well about cause and effect, understanding how to use conditionals, employing and evaluating explanations—these are skills that can benefit us not only in our daily lives, but in science and our search for fundamentals of our knowledge and experience. Come, let us reason together.

> For never yet has any one attained
> To such perfection, but that time, and place,
> And use, have brought addition to his knowledge;
> Or made correction, or admonished him,
> That he was ignorant of much which he
> Had thought he knew; or led him to reject
> What he had once esteemed of highest price.
>> Attributed to the old man in the comedy by
>> William Harvey, *De generatione animalum*

Acknowledgments

Publishing history of the essays in this volume

The essay "The Directedness of Emotions" is new and grew out of an e-mail dialogue with Fred Kroon. The other essays are revisions, often quite substantial, of material in the corresponding sections of my *Five Ways of Saying* "*Therefore*".

The first presentation of the relation of arguments and explanations was at a talk at the Second Conference on Logic and Reasoning of the Advanced Reasoning Forum that was held in Bucharest, Romania, sponsored by New Europe College; it was published as "Arguments and Explanations" in the *Bulletin of Advanced Reasoning and Knowledge*, vol. 1, 2001, pp. 3–17.

The first exposition of the method of evaluating causal claims as inferences appeared in the first edition of my *Critical Thinking*.

BLANK PAGE

Background:
Claims, Inferences, Arguments

Claims

Claims   A *claim* is some written or uttered declarative sentence that we agree to view as true or false, but not both.

The word "uttered" is meant to include silent uttering to oneself. From now on I'll use "utterance" to include writing, too.

   We do not need to make a judgment about whether a sentence is true or whether it is false in order to classify it as a claim. A claim need not be an *assertion*: an sentence put forward as true by someone.

   Some say that claims only represent things that are true or false, namely, abstract propositions or thoughts. But it's utterances we use in reasoning together, and we can focus on those, as representatives, if you like, of abstract propositions or thoughts.

   The word "agree" in the definition of "claim" suggests that it is a matter of convention whether we take a sentence to be a claim. But almost all our conventions, agreements, and assumptions are implicit. Our agreements may be due to many different reasons or causes, including perhaps that there are abstract propositions.

   Often when we reason we identify one utterance with another, as when Dick says "Ralph is not a dog" and later, when Suzy thinks about it, she says "I agree. Ralph is not a dog." We do so when we believe the utterances are equivalent for all our purposes in reasoning.

 Equivalent claims   Two claims are equivalent for our purposes in reasoning if no matter how the world could be, the one is true if and only if the other is true.

I will often assume such equivalences without explicitly saying so.

   Often what people say is *too vague* to take as a claim: there's no single obvious way to understand the words, as when someone says

"This is a free country." Yet, since everything we say is somewhat vague, it isn't whether a sentence is vague, but whether it's too vague, given the context, for us to agree to view it as true or false. In an auditorium lit by a single candle some parts are clearly lit and some are clearly dark, even if we can't draw a precise line where it stops being light and starts being dark. The *drawing the line fallacy* is to argue that if you can't make the difference precise, there's no difference.

Inferences

We reason in order to discern whether certain claims are true. But we also reason to discern whether a particular claim follows from one or more other claims. We might not know whether those other claims are true. But were they true, would the truth of this other claim follow? At the basis of all reasoning is the notion of one claim following from one or more other claims.

Inferences   An *inference* is a collection of claims, one of which is designated the ***conclusion*** and the others the ***premises***, which is intended by the person who sets it out either to show that the conclusion follows from the premises or to investigate whether that is the case.

Whether some claims constitute an inference depends on the intent of the person uttering them. Sometimes people indicate that intention by using certain words to indicate that a claim is meant as a premise, or as a conclusion, or to indicate whether he or she believes the claim.

Conclusion indicators
    hence; therefore; so; thus; consequently;
    we can then show that; it follows that

Premise indicators
    since; because; for; in as much as; given that;
    suppose that; it follows from; on account of; due to

Indicators of speaker's belief
    probably; certainly; most likely; I think

These and many other indicator words are not part of a claim but show our intent in using the claim in a particular way.

In order to investigate the idea of a conclusion following from premises we make some definitions.

Valid, strong, and weak inferences   An inference is *valid* if it is impossible for the premises to be true and conclusion false at the same time and in the same way.

An inference is *strong* if it is possible but unlikely for the premises to be true and conclusion false at the same time and in the same way. An inference is *weak* if it is neither valid nor strong.

The classification of invalid inferences is on a scale from strongest to weakest as we deem more or less likely the possibilities we consider in which the premises are true and conclusion false.

For example, the following is a valid inference:

Maria is a widow.
So Maria was married.

We do not know if the premise is true, but if it is, then the conclusion is not false. In this case the conclusion surely follows from the premise.

The following is valid, too:

All dogs bark.
Spot is a dog.
So Spot barks.

Here we know that the first premise is indeed false: Basenjis can't bark, and some dogs have had their vocal cords cut. But it's not the truth or falsity of the premises and conclusion that determines whether an inference is valid, strong, or weak; rather, it is the ways in which the premises and conclusion could be true or false. In any way that the premises of this inference might be true, the conclusion would be also.

In contrast, the following inference is strong:

Almost all dogs bark.
Ralph is a dog.
So Ralph barks.

If we know no more about Ralph than that he is a dog, then any way in which the premises could be true and conclusion false is unlikely, for we know how rare those are. In this case, too, we say that the conclusion follows from the premises, though there is no certainty, no "must" in that. It is only that relative to what we know, it seems to us very unlikely that the premises could be true and conclusion false.

The following, however, is weak:

Louise is a student.
So Louise isn't married.

There lots of ways the premise could be true and conclusion false: for all we know, Louise might be forty years old with a husband and child.

Our evaluation of the strength of an inference is relative to what we believe. "Likely" means "likely to us." But typically the scale from strong to weak is not so completely relative to a particular person that there is no hope we can agree on the strength of inferences. Suppose we disagree. I find a particular inference strong, and you find it weak. If we wish to reason together, you should describe to me a way the premises could be true and the conclusion false that you think is not unlikely. That may depend on knowledge you have of how the premises could be true which I do not have, but once you've made that explicit we can agree or disagree that there is such a possibility. The only issue, then, would be whether we agree that the possibility is likely. Sometimes we can't come to a clear determination, but further examination will leave us with a clearer understanding of what our differences in evaluation are, based on more than just whim. When the beliefs involved in determining the strength of an inference are made explicit, determining the inference to be strong or weak is far more likely to be a shared judgment.

In sum, we say that the conclusion of an inference **follows from** the premises if the inference is valid or strong.

Arguments

The paradigmatic use of inferences is in attempts to convince someone that a claim is true.

Arguments   An *argument* is an inference that is intended by the person who sets it out to convince someone, possibly himself or herself, 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 that the criterion for whether an argument is good is whether the argument actually does convince. If I'm 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 her, but though you may be convinced, that does not mean the argument is good. Perhaps other ways to convince, such as entreaties, exhortations, sermons, or advertisements, can be judged by how well they convince, but that is not a criterion for judging attempts to establish the truth of a claim. A *good argument* is one that gives us good reason to believe the conclusion. But what does "good reason" mean?

If an argument is to give us good reason to believe its conclusion, we should have good reason to believe its premises, for from a false claim we can reason as easily to a false conclusion as a true one.

The Prime Minister of England is a dog. All dogs have fur.
So the Prime Minister of England has fur. (false conclusion)

The Prime Minister of England is a dog. All dogs have a liver.
So the Prime Minister of England has a liver. (true conclusion)

It seems, then, that a good argument should have true premises. But 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 well. Rather, for an argument to be good we must have good reason to believe its premises. We might, though, have good reason to believe the premises and not be aware of that. For an argument to be good, we need to recognize that we have good reason to believe the premises and actually believe them, for otherwise what convincing is done has no basis in our beliefs.

Plausible claims A *claim is plausible* to a particular person at a particular time if:

- The person has good reason to believe it.
- The person recognizes that he or she has good reason to believe it.
- The person believes it.

A claim that is not plausible is *implausible* or *dubious*.

The classification of claims as plausible or implausible is on a scale from the most plausible, ones we recognize as true, to the least plausible, those we recognize as false. Though we do not have precise measures of plausibility, we can often compare the plausibility of claims and by being explicit about our background we can usually agree on whether we will take any particular claim to be plausible. If we did not think that we can share our judgments of what is plausible, we would have no motive for trying to reason together. So if I say a claim is *plausible* without specifying a particular person, I mean it's plausible to most of us now, as I'm writing.

Good reason to believe a claim needn't always be established by reasoning, for then we would have no place to start, no plausible claims that would not require further justification, continuing forever. Some claims we take as plausible because of our personal experience, or our trust in authority, or our beliefs about the nature of the world.

But it's not just that the premises of a good argument have to be plausible. They have to be more plausible than the conclusion, for otherwise they would give us no more reason to believe the conclusion than we had without the argument.

Begging the question   An argument *begs the question* if it has a premise that is not more plausible than its conclusion.

Further, for an argument to be good, the conclusion must follow from the premises. For example, consider:

Richard L. Epstein is the author of this essay.
So Richard L. Epstein is bald.

This argument is weak: there are lots of likely ways the premise could be true and conclusion false. Though you know that the premise is true, it gives no reason to believe the conclusion.

Arguments, being inferences, are classified as valid, strong, or weak, and as with inferences it is valid or strong ones we take as establishing that the conclusion follows from the premises. But do strong arguments give good reason to believe the conclusion?

As an example, consider that last week Dick heard there are parakeets for sale at the mall. He knows that his neighbor has a birdcage in her garage, and he wonders whether the cage will be big enough for one of those parakeets. He reasons:

(‡)   Every parakeet I or anyone I know has seen, or read, or heard
about is less than 50 cm tall. So the parakeets on sale at the
mall are less than 50 cm tall.

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

Strong arguments, as in this example, give us good reason to
believe the conclusion, or at least good enough reason for our daily
lives and, as we'll see in the following essays, for science. Moreover,
a strong argument is often better than a valid one with the same conclu-
sion. Replacing the premise of (‡) with "All parakeets are less than 50
cm tall" would yield a valid but worse argument, for that claim is less
plausible than the premise of (‡). 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.

Arguments that lie in the broad center of the scale and the clearly
weak ones are certainly not good. We needn't bother to classify them
as bad versus very bad, since any bad argument tells us nothing about
the conclusion we didn't already know.

We now have three tests an argument must pass for it to be good.

Necessary conditions for an argument to be good

- The premises are plausible.
- The argument does not beg the question.
- The argument is valid or strong.

These conditions are relative to a particular person, though we can have
confidence that they establish an intersubjective standard for the eval-
uation of arguments.

Whether these conditions are also sufficient is a large topic which
is examined in the companion volume in this series *The Fundamentals
of Argument Analysis*. In what follows, though, I will generallly treat
them as both necessary and sufficient.

Repairing arguments

In our daily lives we often treat many arguments as good that do not seem to satisfy these conditions. For example, consider:

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. 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 play with 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 should revise many arguments by adding an unstated premise or even an unstated conclusion.

When are we justified in doing so? How do we know whether we've revised an argument well or just added our own assumptions? To repair arguments that are apparently defective, we must have some standards, for otherwise we will end up putting words in other folks' mouths. 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 reasoning we are evaluating:

- Knows about the subject under discussion.

- Is able and willing to reason well.

- Is not lying.

Often someone with whom we wish to reason does not satisfy these conditions. But when we discover that, then it makes no sense to continue to reason with him or her. We should be educating, or consoling, or pointing out errors.

The Principle of Rational Discussion justifies adopting the following guide.

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

- The argument becomes valid or strong.
- 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 implausible, we may *delete* that premise if the argument becomes no worse. In that case we say the premise is ***irrelevant***.

Given only this Guide, we might try to repair every argument into a good one. That would be wrong, for there are 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 that there's nothing obvious to add.
- A premise is implausible or several premises together are contradictory and cannot be deleted.
- The obvious premise to add would make the argument weak.
- Any obvious premise to make the argument strong or valid is implausible.
- The conclusion is clearly false.

It's 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, for it would be putting words in his or her 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.

Consideration of two particular kinds of arguments is important for the essays that follow: generalizations and analogies.

Generalizations

Generalizations A *generalization* is an argument in which we conclude a claim about a group, called the ***population***, from a claim about some part of it, the ***sample***. Sometimes we call the the conclusion the *generalization*. Plausible premises about the sample are called the ***inductive evidence*** for the generalization.

The following are generalizations:

(‡) Every dog I've seen barks.
    So all dogs bark.
        sample: The dogs the speaker has seen.
        population: All dogs.

Every dog I ever met except one can bark.
So almost all dogs bark.
        sample: The dogs the speaker has met.
        population: All dogs.

Of dog owners who were surveyed, 98.2% said their dogs bark.
So about 98% of all dogs bark.
        sample: The dogs of the pet owners surveyed.
        population: All dogs.

The last is called a *statistical generalization* because its conclusion is a statistical claim about the population.

If we have no reason to think that the sample is similar to the population, then the generalization is worthless, a bad argument. What we want for a good generalization is for the sample to be representative.

Representative sample A sample is *representative* if no one subgroup of the whole population is represented more than its proportion in the population. A sample is ***biased*** if it is not representative.

The first and second examples at (‡) are bad because we have no reason to think that the dogs the speaker has seen are representative of

all dogs. We don't know enough about the sample in the third generalization to make a judgment about whether it's representative.

Random sampling is an important method for getting a representative sample.

Random sampling   A sample is *chosen randomly* if at every choice there is an equal chance that any of the remaining members of the population will be picked.

Random sampling does not guarantee that the sample will be representative. Choosing two students randomly from the 716 at McEpstein High School to interview about their views on gay marriage is not going to give a representative sample. The sample has to be large enough for us to have good reason to think it is representative.

But even if we have confidence that the sample is representative, if it's not studied well then it's no use for concluding anything about the population. Maria asked all but three of the thirty-six people in her class whether they've ever used cocaine, and only two said yes. So she concluded that almost no one in the class has used cocaine. But there's no reason to think that people will answer truthfully to such a question, so her generalization is not good.

Necessary conditions for a generalization to be good

- The sample is representative.
- The sample is big enough.
- The sample is studied well.

These conditions do not establish a different standard from 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.

Analogies

Analogies   A comparison becomes *reasoning by analogy* when it is part of an argument: On one side of the comparison we draw a conclusion, so on the other side we say we should conclude the same.

For example, consider:

> We should legalize marijuana. After all, marijuana is just like
> alcohol and tobacco, and those are legal.

The comparison here is between marijuana on the one hand and alcohol and tobacco on the other. The latter are legal. So we should make marijuana legal, too. But no connection between the premise and the conclusion has been supplied besides saying that marijuana "is just like" alcohol and tobacco.

The difficulty in reasoning by analogy is to make clear what we mean by "is just like" or "is the same as" in order to justify the inference in terms of the comparison. Such a justification calls for some general claim under which the two sides of the comparison fall. Often analogies are sketchy, with only the comparison offered, so that their main value is to stimulate us to search for such a general claim. It must be one that relies on the similarities and for which the differences between the two sides of the comparison don't matter. Though that procedure is somewhat more involved than in analyzing many other arguments, it is does not require further necessary conditions for an argument to be good.

Reasoning backwards

One particular mistake in reasoning is important to note for some of the discussions that follow. For example, Suzy said to Lee:

> All CEOs of computer companies are rich. Bill Gates
> is a CEO of a computer company. So Bill Gates is rich.

Lee sees that Suzy's argument is valid, and he knows that Bill Gates is rich and that he's a CEO of a computer company. So he reckons that the other premise, "All CEOs of computer companies are rich" is true, too. But he's wrong: there are lots of CEOs of small, struggling computer companies who are not rich. Lee is arguing backwards.

Arguing backwards   *Arguing backwards* is the mistake of concluding that the premises of an inference are true because the inference is valid or strong and its conclusion is plausible.

This concludes the very brief summary of the basics of inference and argument analysis needed for the succeeding essays.