Chapter Five

INFERENCE TO THE BEST EXPLANATION

> 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 the inference. Thus, one infers, from the premise that a given hypothesis would provide a "better" explanation for the evidence than would any other hypothesis, to the conclusion that the given hypothesis is true.
>
> Gilbert Harman[1]

INFERENCE TO THE BEST EXPLANATION

We have been treating the expression ***inference to the best explanation*** as technical jargon. It is a way of looking at evidence, or at least purported evidence, in an inductive argument. If we look at the component words in this expression, we will discover quite a lot. First of all, we are dealing with an ***inference***. For most purposes, we can consider that just another way of saying that we have an argument to be considered. This inference is to an ***explanation***. But we are not dealing with just an inference

to an explanation, but to the ***best*** explanation. This implies two very important things. First, in order for there to be a comparison, there must be other possible explanations of the data in the argument, ***rival explanations***. And, the argument is also committed to this original explanation being better than all of these rivals. Therefore, there seems to be some ***rank ordering*** of the explanatory candidates, even if this is not explicitly stated.

I will use all of this as a way of articulating at test of the quality of evidence within an argument. This test will be most straightforward when you are what I have called a consumer of an argument. Connie thought she had evidence that her boyfriend was smooching Mary Jane. Leslie laid out her evidence that she had lost Johnny. Holmes had evidence about Watson's decision about the investment, and about what happened at Ridling Thorp Manor. We must decide whehter these arguments are any good. Was the evidence for these hypotheses strong? What I am going to call the IBE "recipe" is a procedure for answering these kinds of evaluative questions.

Inference to the Best Explanation Recipe

1. Schematize the argument
2. List some serious (hopefully challenging) rival explanations
3. Rank order all of the explanations – the original along with the rivals
4. Based on the rank order see if the original is the best explanation. If it is, the evidence has passed the test and looks pretty good. If it isn't, it's failed the test, and the evidence is weak, maybe nonexistent.

Let's apply the test or recipe to you cousin, and the Times', argument.

SCHEMATIZING THE ARGUMENT

Maybe as you tried your hand at schematizing the argument from the last chapter you had a glimmer of this, but let me state it explicitly. ***Often the single hardest part of argument analysis or the IBE recipe is simply identifying what the argument is in the first place***. There are a number of reasons for this. First and foremost, people aren't always as clear as the might be when they state their arguments. But there are other complicating factors as well. My guess is that if you contacted the reporter, Mr. Liptak, he would tell you that he wasn't presenting an argument at all, but simply "reporting" on a controversy that's brewing in Washington regarding Mr. O'Neill. Still, the fact that there is this controversy, that law professors like Professor Rhode say that the evidence is "disqualifying" for his nomination, and that other legal experts see at least potential evidence of professional misconduct, indicates that your cousin hasn't simply gone off the deep end. Add to all of this the fact that judicial nominations are clearly political, and that few of us can read articles like this and set our personal politics to the side. These unavoidable biases that we all carry with us will often tempt us to simply misread what the argument is. Finally, as we get a hint of in this rather long newspaper article, but becomes daunting when an argument is developed over the course of a whole book, the sheer number of words, thoughts, and sentences makes it extremely challenging to keep the structure of the argument clearly in mind.

Granted all of this, the first step in the IBE recipe is not only the most difficult, it is the most important. If we misrepresent what the argument is, then all of our work in analyzing it will be a waste of time. Who cares if you show "the argument" to be spectacularly successful, or a dismal failure, if it wasn't you cousin's argument in the first place?

Useful schematization requires three virtues, all of which defy simple characterization. First and foremost, as we have just emphasized, you should strive for ***copy fidelity***. Your task is to represent "the other person's argument," a representation of his or her evidence. You may think of better ways to make the argument, or you may even think that the evidence points in a different direction. That's all fine, and will be useful in later steps. Right now, however, your job is to faithfully represent the argument as is was stated. You want to also strive for ***brevity***. We just saw how an argument might take several columns of newspaper, but just imagine when we look in a later chapter at Darwin's "abstract" of his theory in ***Origin of Species***,[2] and try to keep straight all of the evidence presented in over four hundred pages. In order for your schematization to useful to you, you will need to keep your representation of the evidence down to, say, no more than a page. Finally, and most difficult of all, you should strive for ***charity*** in your schematized arguments. You want to present the argument in the strongest form you can. This is not because you are being nice, or discounting the above virtue of copy fidelity. It is because you want to avoid at all costs weakening the evidence in the way you choose to schematize it. This is particularly important when you are dealing with arguments with which you disagree. If you come to the judgment that the evidence is weak, you need to make darn sure that you've given the evidence its best shot.

START AT THE BOTTOM [FIND THE CONCLUSION]

Since Step One in the recipe is often the most difficult, and most important, we will indulge ourselves with a couple more subsections on how to do it, or more modestly, some hints for doing it better.

We've already discussed the fact that ***conclusions*** may come anywhere in a statement of an argument. Still, in the ***schematic form*** I am urging on you, they always come at

the bottom, they are always are identified as t_0 ("0" to start a sequence of explanations, and "o" as a reminder that this explanation is the original one), and they are always ***explanations*** of the data, not simply statements of the data. I strongly suggest that you begin your schematizing of the argument by trying to identify its conclusion.

Often times you will find hints in the statement that will guide you to the argument's conclusion. There are many words and phrases that are commonly used to alert readers or listeners that an inference is being drawn. Some of the classics you will find in any introductory logic book are: "therefore," "hence," "so," "it follows that," and many others. But at other times you are simply expected to pick out what the theory is that is supposed to be supported by the evidence. The best advice in these latter cases is simply to ask yourself something very general and vague like, "what's the point of all this?," or as suggested above, "what the heck is going on?" Once you have a candidate, now see whether it explains some of the data in the argument. If it doesn't seem to, you might want to look for another candidate as the argument's conclusion.

Two other general comments are appropriate here. First, don't get discouraged. This is hard stuff. It will get easier and more natural as you get more experience using the recipe. And secondly, there will be times when you fail to discover a conclusion to begin your schematization because the passage of prose in front of you is not an argument in the first place. We obviously use language do lots of things – make simple assertions, push people's buttons, or simply vent – stating an argument is only one use of language.

There is probably data in the *Times* article that might be used to support a number of conclusions -- that Mr. O'Neill will not be confirmed, or that he will, or that the administration will vigorously defend him. But your cousin

focused on the charge of plagiarism. So, for our purposes the conclusion is:

> t_0. Mr. O'Neill plagiarized from the book review in his *Supreme Court Economic Review* article.

FIND THE RELEVANT EVIDENCE

In the New York Times article we learned a lot of stuff. Mr. O'Neill is a boyish looking 46 year old man, and was wearing blue jeans at the time of the interview. That's probably good newspaper style. It makes the subject of the story human and real. It is certainly not part of your cousin's evidence, though. Also, there is a great deal of data about the political implications of the scandal, but that data is only tangential to the charge of plagiarism. Your cousin is going to be concerned with four or five key facts.

> e_1. The *Supreme Court Economic Review* issued a retraction of Mr. O'Neill's article.
> e_2. "Substantial portions" of the article, the editors wrote, were "appropriated without attribution."
> e_3. Details of the similar passages.
> e_4. Other articles by O'Neill have similar problems.
> e_5. Mr. O'Neill voluntarily stepped away from tenure.

We should feel free to exercise some good judgment about what can safely be omitted from our schematized representation of the argument, but we must always bend over backwards to include everything that is relevant, even those facts that might point in the away from the argument's conclusion. We need to consider all of the evidence, not just the stuff that suits our (or our cousin's) purposes. Simple

fairness requires this. We learn other data that seems to help Mr. O'Neill's case. At the very least I think we would want to include the following.

e_6. Mr. O'Neill attributes the similar passages to "poor work method," and not "keep[ing] appropriate track of things."
e_7. One of the authors of a source article admires Mr. O'Neill and does not believe he would copy such banal points.
e_8. Republican congressional leaders and the President knew of the charges and have accepted Mr. O'Neill's explanation.

When all is said and done, I think we get something like the following as an accurate schematization of your cousin's evidence against Mr. O'Neill.

e_1. The *Supreme Court Economic Review* issued a retraction of Mr. O'Neill's article.
e_2. "Substantial portions" of the article, the editors wrote, were "appropriated without attribution."
e_3. Details of the similar passages.
e_4. Other articles by O'Neill have similar problems.
e_5. Mr. O'Neill voluntarily stepped away from tenure.
e_6. Mr. O'Neill attributes the similar passages to "poor work method," and not "keep[ing] appropriate track of things."
e_7. One of the authors of a source article admires Mr. O'Neill and does not believe he would copy such banal points.
e_8. Republican congressional leaders and the President knew of the charges and have accepted Mr. O'Neill's explanation.

---

t_0. Mr. O'Neill plagiarized from the book review in his *Supreme Court Economic Review* article.

RIVAL EXPLANATIONS

For our purposes, ***rival explanations*** will be accounts of the data that flat-out ***deny*** the original explanation and substitute a completely different story of the data offered as evidence. It will be useful to imagine each truly rival account of Mr. O'Neill's article as starting out with a lengthy preliminary phrase -- "***no, no, no,*** O'Neil did not plagiarize from the book review in his *Supreme Court Economic Review* article, rather …" This is important because the original explanation might be phrased in very different language.

t_0'. The article in question was copied from another law journal.

Or an account might offer a more detailed (or less) account of what happened.

t_0". Mr. O'Neill discovered Ms. Daily's review on-line. He cut and pasted the passages into his manuscript, then changed a few words here and there.

Neither t_0' nor t_0" will count as rival explanations. If you were to challenge your cousin with them, I think she would not say, "oh yeah, maybe I was wrong," but rather she'd exclaim, "exactly!"

So what else might have happened? Mr. O'Neill provides one very different account in the New York Times article.

t_1. The similar wording was inadvertent. Notes from his sources, Ms. Daily, Mr. Katyal,

> and Mr. Caplan, got mixed up with Mr. O'Neill's thoughts and observations.

Other things might have been going on as well. Maybe Ms. Daily and Mr. O'Neil, independently, and without each other's knowledge, plagiarized from some third source that simply has not surfaced yet.

> t_2. Mr. O'Neill and Ms. Daily independently plagiarized from a third source.

Perhaps there is some really deep conspiracy going on.

> t_3. Mr. O'Neill's original manuscript to the *Supreme Court Economic Review*, which was good, honest scholarship, was altered by liberal law student editors seeking to discredit him for his work on behalf of conservatives.

Once we allow our imaginations to start going, it is remarkable how many different accounts of data can be produced. I am constantly amazed, and quite delighted, at the number of creative, and often funny, rival explanations my students have been able to create for me on thirty-plus years of quizzes and examinations. A word of caution is relevant here, however. Don't let your desire to be creative and amusing trump the need for listing the most plausible and challenging rivals. Martians or political conspiracies may bring a smile to our faces, but simple bad record keeping is much more challenging to your cousin's account.

RANK ORDERING EXPLANATIONS

We now have on the table four competing accounts of the key data in the New York Times article.

t_0. Mr. O'Neill plagiarized from the book review in his *Supreme Court Economic Review* article.

t_1. The similar wording was inadvertent. Notes from his sources, Ms. Daily, Mr. Katyal, and Mr. Caplan, got mixed up with Mr. O'Neill's thoughts and observations.

t_2. Mr. O'Neill and Ms. Daily independently plagiarized from a third source.

t_3. Mr. O'Neill's original manuscript to the *Supreme Court Economic Review*, which was good, honest scholarship, was altered by liberal law student editors seeking to discredit him for his work on behalf of conservatives.

Inference to the best explanation asks us to judge one of these explanations as better than all the rest. How in the world do we start the process of judging one explanation as superior to another? What counts, and what doesn't count, in such a comparison? We will address this in some detail in Chapter Six, but for now let's simply phrase the question as, "***which account makes the best sense of what we know?***"

I assume that both t_3 and t_4 would rank way down on your list, compared to t_0 and t_1. Isn't part of the reason for this, the fact that both of them introduce something "out of the blue" to explain the article? Where did this mysterious third source come from? Or those devious law student editors? Why in the world, if Professor Daily is a plagiarizer, would she be so reckless to call attention to her article? And if the law students substituted the damning passages, why didn't Mr. O'Neill catch it when he approved the "proof copy" of his article?

What about t_0 and t_1, however? They both seem reasonable enough. Let me simply assert some factors that do not count in rank ordering explanations. The ***best explanation*** is not necessarily the one we like the best, nor the one that best accords with our politics, religion, or moral perspectives. ***It is the one that is most plausible***.

Here comes a scary fact! ***You*** have to make the judgment about which explanation is best. There is no "objective," "reliable" test or formula you can utilize that automatically identifies the best explanation. The whole recipe, therefore, rests on a step that is candidly, unavoidably, ***subjective***. When it comes to flavors of ice cream, or styles of beer, being subjective means that people's preferences are relative to who they are, and, consequently, all over the place. If evidence evaluation is the same, we're done for, and I can stop writing my book, and teaching my courses as I do. Fortunately, I believe, explanatory plausibility is very different from beer preferences. Even though each of us, individual subjects, must rank order alternative accounts for ourselves, it turns out that in a great number of contexts -- courts of law, the natural sciences, and even stories about very similar passages in law journals -- ***subjective*** judgments about plausibility can turn out to be ***inter-subjective***. When all is said and done, when we think about it as free from prejudice and bias as we can be, we discover widespread agreement about what the best explanation is. We are the most intelligent species that has ever existed, and part of being intelligent is being darn good at spotting the best explanation of what's happening around us.

IDEAL AGNOSTICS

I want to share with you an idea that I am very taken with these days. It comes from a contemporary philosopher, as it turns out a very candid Christian philosopher, named Peter van Inwagen. He proposes an audience for arguments (at

least those that occur in philosophical debates) that is psychologically impossible, but is useful to imagine, nonetheless.

> The audience is composed of what we might call ideal agnostics. That is, they are agnostic as regards the subject-matter of the debate. … [E]ach member of the of the audience will have no initial opinion about [the subject of the debate]. … My imaginary agnostics … would very much like to come to some reasoned opinion [on the debate] … indeed to achieve knowledge on that matter if it were possible. … They don't care *which* position … they end up accepting, but they very much want to end up accepting one or the other.[3]

Ideal agnostics are absolutely indifferent -- intellectually, personally, and in every way that might bias them -- about what the best explanation is. But that doesn't mean they don't care. They are also passionately committed to figuring out which explanation is the strongest.

I'm no ideal agnostic, and neither are you. But I think we are both well-served in our discussions and investigations, to pretend that we are. Indeed, I am suggesting that anytime we evaluate another's potential evidence, we try as hard as possible to adopt the position of the ideal agnostic, knowing all along that we will fail in certain respects. When we are presenting our own argument, I would also suggest that we pretend our audience is not composed of partisans, but rather ideal agnostics.

This whole little subsection might strike you as a tedious distraction. I am belaboring all of this because we all carry with us biases that will inevitably affect some of our rank ordering of explanations, especially when two competing accounts are very close to one another. That is the position I find myself in with the current argument. I care very deeply about constitutional law. I have very strong opinions on some

of the judges that Mr. O'Neill helped to get on the Federal Courts, including the Supreme Court. And, as you're going to find out anyway, I am a political liberal.

So now you know, but I'm still hopeful for inter-subjective agreement. I, trying just as hard as I possibly can to be the ideal agnostic, am forced to rank your cousin's hypothesis ahead of Mr. O'Neill's explanation. Four factors point me in that direction. One is, as the *Times* article told us, how rare it is for a law journal to take actions like the *Supreme Court Economic Review* did. Obviously, they thought this was very serious. The second is the pattern of unattributed "quotes" in his other articles. I concede that poor record keeping may be a more general problem with his research methods, but I see a tendency to rely way too much on what other scholars have said. My third consideration is firsthand knowledge of how scholars are trained to be scrupulous about being sure to acknowledge the work of others. But the clincher for me was how the thinking and wording taken from Ms. Daily's review was altered in very slight ways. If he were simply recording what she said for further thought and possible inclusion in his own work, why the subtle changes?

ASSESSMENT OF THE EVIDENCE

I rank order our four explanations in the following order.

t_0. Mr. O'Neill plagiarized from the book review in his *Supreme Court Economic Review* article.

t_1. The similar wording was inadvertent. Notes from his sources, Ms. Daily, Mr. Katyal, and Mr. Caplan, got mixed up with Mr. O'Neill's thoughts and observations.

t_2. Mr. O'Neill and Ms. Daily independently plagiarized from a third source

> t_3. Mr. O'Neill's original manuscript to the *Supreme Court Economic Review*, which was good, honest scholarship, was altered by liberal law student editors seeking to discredit him for his work on behalf of conservatives.

It is a ***total accident*** that this particular ranking was in strict sequential order of the subscripts, so please don't be misled by this coincidence.

The whole purpose of the inference to the best explanation "recipe" is to assess the quality of evidence in an argument. We need to find the ***best*** explanation. The whole test depends on what is in first place. In my best "ideal agnostic" judgment, your cousin's theory was the best explanation, and therefore, her evidence is pretty good. For all the talk about inter-subjectivity and ideal agnostics, I fully realize that some of you will have ranked t_1 ahead of t_0. Those of you who have come to that judgment would say that since there is a better explanation of the facts about the article, your cousin's evidence is weak, maybe so weak that you don't see real evidence at all.

I have been asking my students to use the inference to the best explanation recipe to assess the quality of evidence presented in an argument for more than three decades. The single most common mistake that my students make, including some of the best and most intelligent, is to forget about the purpose of the recipe, and neglect to offer an assessment of the evidence in the argument. They often beautifully schematize it, come up with some challenging rival explanations, offer subtle and insightful comments about how and why they have rank ordered as they have, but then remain silent on the quality of the evidence. I am almost tempted to include a fifth step in the recipe saying something like the following:

5. Conclude your analysis with one of the following two sentences: "*Since the original theory proved to be the **best explanation** of the data in the evidence, the argument's evidence is pretty good (strong, etc.)*" ***or*** "*Since there is a **better explanation** of the data in the evidence, the argument's evidence is weak (poor, non-existent, etc.)*".

Step Four ***requires*** an explicit evaluation of the evidence, as it was presented, and schematized, in the original argument!

WHAT ABOUT TIES?

Suppose, in your best ideal agnostic frame of mind, you came to the conclusion that the plagiarism hypothesis and the poor record keeping rival were equally plausible explanations of all the data you had? What happens in the recipe when the original and one of the rivals are tied for first place?

This is a classic half-full, half-empty kind of dilemma. You might say that since the original is tied as the ***best explanation***, there's some evidence for that conclusion. You might also say, however, that since there's rival explanation that's tied as the best explanation, the evidence is not so hot. I think whichever way we go the message is really the same. The original's being tied for first place allows us to see why someone would offer the argument in its defense in the first place, and why there is some evidence that seems to support it. A rival being tied for first place tells us that the evidence is far from conclusive. Ideally, in such a case, we go out and do a little more investigating, and see if we could discover some new data that would help break the tie. And, indeed, the whole subject of new data is the topic for our next chapter. But, before heading there, let's do a review exercise.

A MAGICAL ENCORE?

Quite by accident I discovered a glitch in the ipod software. On a Saturday night last year my wife and I went to a banquet for the League of Oregon Cities. The entertainment was Pink Martini, a Portland band I like a lot. I had already planned that I was going to ask for two songs when they came back for an encore – "Lilly" and "Que Sera Sera." As it turned out they did "Que Sera Sera" as part of their concert, and there was no chance to ask when they did their encore. On Sunday as we drove back from Portland I plugged in my ipod to listen to them again. I set the settings to "All" and to "Shuffle Songs." This meant that my ipod searched through both of their albums, found all 36 songs and played them in "random" order. That's the glitch! The last two songs were "Lilly" and "Que Sera Sera." The exact encore I had imagined the night before! What are the odds of this? My theory is that these two songs came up last, not randomly, but because of all the Pink Martini songs, I listen to these two the most often. I am thinking of writing to Apple to tell them about the problem.

This crazy philosopher has a theory that there is a glitch in the ipod software. For practice, and to make sure you've got the IBE recipe down pat, take a few minutes, and using all four steps in the inference to the best explanation recipe. assess the quality of evidence he has for this theory.

ENDNOTES

[1] G. Harman, 'The Inference to the Best Explanation,' The Philosophical Review 74:1 (1965), p. 89.
[2] Charles Darwin, *On the Origin of Species: A Facsimile of the First Edition*, (Cambridge, MA: Harvard University Press, 2001).
[3] Peter van Inwagen, *The Problem of Evil* (New York: Oxford, 2006), p. 44.