Philosophy of Science

Lecture 11: Bayesianism and Evidence

Instructor: David McElhoes, PhD

Today's Agenda

- Goal: To understand the problems with the notion of scientific evidence, and the Bayesian solution (along with its drawbacks)

- Breakdown
  1. The Problem of Scientific Evidence
  2. The Bayesian Theory of Evidence
  3. Subjectivism about Probability
  4. Objectivism about Probability

1. The Problem(s) of Scientific Evidence

From Induction to Evidence

- Earlier, we characterized science as being reliant on induction, in order to provide evidence for theories.
  - From particular observations, we conclude that a generalization PROBABLY holds (or that the next observation we make will turn out a certain way).
- This thought presumes that giving evidence requires a kind of argument.
  - This is a natural idea: to justify that something is true is to give a reason to think that it's true. A reason is an argument.
- But perhaps we can challenge that. Perhaps we can give a theory of evidence that allows us to say that an observation O count as evidence/confirmation for a hypothesis H without any appeal to any argument whatsoever.
  - The idea is to give a non-argument theory of scientific justification.
- The most widely used theory of evidence/confirmation is the ***hypothetico-deductive*** method.
  - Although it is still grounded in logic, it characterizes evidence without appeal to argument. Instead of argument, we simply think about the consequences of conditional expresssions: "if…then" expressions.

Hypothetico-Deductivism

- The idea goes back to logical positivists

- **Hypothetico-Deductivism**: the idea that a hypothesis is confirmed (or corroborated?) when its logical consequences turn out to be true.
    - "All swans are white," implies that the next swan one sees is going to be white.
    - If we suppose that the next swan you see IS indeed white, then that is taken to be evidence that ALL swans are white.

- For the hypothetico-deductivist, to give evidence for a hypothesis doesn't require an argument; it simply requires that the consequences of a hypothesis indeed turn out to be true.
    - This, in fact, is a routine method of reasoning in science; we even teach it to freshmen students enrolled in "critical thinking" classes!!

- But there are SERIOUS (and simple!) problems with this way of thinking.

Glymour's Problem(s)

- Clark Glymour points out, if H is a hypothesis that we want to test, H is going to imply H or P, where P is anything at all (e.g., my shoes are untied).
  1. H implies (H or P)
  2. So, if hypothetico-deductivism is correct, then an observation that P (e.g., my shoes are untied) counts as confirmation that H is true.
     - But that's nuts! This means that any hypothesis is confirmed by anything and everything!
- Also: since anything implied by H is also implied by H&P (e.g., where P = Santa Claus exists), it follows that, on the hypothetico-deductive method, anything that confirms H (by virtue of being an implication that is observed to be true), also confirms that Santa Clause exists.
  1. If H implies E, then H&P also implies E
  2. If hypothetico-deductivism is true, then an observation that E counts as confirmation that H&P is true.
  3. If H&P is true, then P is true.
  4. Thus, any observation that confirms H also confirms that P (e.g., that Santa Claus exists).
     1. But that's surely nuts!

The Ravens Problem

- Logical Positivists/Empiricists supposed that all observed instances of an R being B (e.g. a Raven being black) confirms the generalization “All R’s are B’s.”

- But consider this:
  1. If an observation O confirms a generalization “All R’s are B’s”, then it also confirms anything that is logically equivalent to that generalization, e.g., “All non-B’s are non-R’s”
  2. Thus, (from 1) if an observation O*, confirms “All non-B’s are non-R’s” then O* also confirms “All R’s are B’s.”
  3. If the hypothetico-deductive theory of confirmation is correct, then an instance of a non-B that is a non-R confirms that “All R’s are B’s”.
  4. But: my blue left shoe, by virtue of being blue, is a non-black thing, and by virtue of being a shoe is a non-raven; i.e., my blue left show is a non-B that is a non-R.
  5. Thus, if the hypothetico-deductive theory of confirmation is correct, then my blue left shoe confirms that all ravens are black.
     1. This is crazy! Nelson Goodman, in a moment of cheekiness, calls this “indoor ornithology” We can test hypotheses about birds without every having seen a single bird!

- It appears that something is very wrong with the hypothetico-deductive theory of confirmation!!
    - Yikes! That’s actually a pretty standard way to do science!

2. The Bayesian Theory of Evidence

Evidence and Conditional Probability

- The Bayesian account of evidence has nothing to do with logical consequence. Rather, it involves the notion of a conditional probability (also called posterior probability).
  - **<u>Prior Probability</u>:** The probability that a claim, h, is true.
    - E.g., P(h) is the probability of the truth of hypothesis h.
  - **<u>Conditional Probability (or: Posterior Probability)</u>**: The probability of a claim, h, given some other claim, e.
    - E.g., P(h|e) is the probability of the hypothesis h, given the truth of observation-claim e.
- Using these notions, Thomas Bayes proved the following theorem:
  - $P(h|e) = \frac{P(e|h)P(h)}{P(e|h)P(h)+P(e|not-h)P(not-h)}$
    - The probability of h given e is the product of the probability of e given h and the probability of h alone, divided by the product of the probability of e given h and the probability of h alone, plus the product of the probability of e given not-h and the probability of not-h alone.

The Bayesian Theory of Evidence

- Bayesians think that we can give a formal theory of scientific evidence using these notions.
- Since P(h|e) and P(h) both have numeric values (between 0 and 1), a comparison between those probabilities tells us the difference that e makes to the probability of our hypotheses.
  - This gives us a theory of evidence!

- **Bayesian Theory of Evidence**: observation e is evidence (i.e., confirms) a hypotheses h if and only if P(h|e) > P(h).
  - What a simple analysis!

An Example:

- Suppose you think that the prior probability of the Winnipeg Blue Bombers winning the CFL's prestigious Grey Cup (our h) is .3, and their not winning (our not-h) is .7
  - P(h) = .3
  - P(not-h) = .7 [since P(h)+P(not-h) = 1]
- Suppose you hear that their opponent, the Saskatchewan Roughriders, lost their best player to a travelling circus (our e).
  - Suppose that this is happy news for Bomber's fans, since they've only won 2/10 games against the Roughriders when they've fielded their star player.
- You then suppose that the probability that the Roughrider's best player was lost, given that the Bombers win the Cup is .8.
  - P(e|h) = .8
- You also suppose that the probability that the Roughrider's best player was lost, given that the Bombers will LOSE the cup is only .1
  - Since the bombers used to win 9/10 games before the RR drafted that rascal!
  - P(e|not-h) = .1
- These suppositions are enough to tell us the conditional probability that the Bombers will win, given that their opponent's best player would be lost.
  - P(h|e) = (.8)(.3) / [(.8)(.3)+(.1)(.7)]   = .77
- So, since .77 > .3, it follows that the loss of the star player counts as evidence that the Bombers will win.

So What Are Probabilities?

- So far so good…
- One of the main challenges facing Bayesians is to give some account of these probabilities.
- Two pressing questions:
  - What are probabilities?
    - Out there in the world as features of reality? (objectivism)
    - Degrees of belief? (subjectivism)
  - How do we get our prior probability values to put into the Bayesian equation?

3. Subjective Probabilities

Credence: Degrees of Belief

- In order to use Bayesian methods, we must have some grasp of where these prior probabilities come from… why THESE probabilities?
  - In answering these questions, the Bayesians are generally in agreement that these probabilities represent degrees of belief (credence) – they come from us and our prior beliefs.
  - That is, Bayesians are generally subjectivists.

- **Subjectivism**: the idea that the probability of h is nothing more than a measure of an ideally rational person's credence in h (their degree of belief that h is true).
  - The ideally rational person is not a single real or unreal person; but rather a class of fictional individuals who are ideally rational.
  - A prior probability of h is the degree of belief that an ideally rational person has in h

Three Questions

- This leaves us to wonder:
  - (1) What is an ideally rational person?
  - (2) How do we figure out what their credence in h is?
  - (3) What happens if different ideally rational people have different credences regarding the same hypothesis?

The Ideally Rational Person

- The **ideally rational person** is any person who has a perfectly coherent set of beliefs: their network of beliefs/credences conform to the following four Axioms:
  - A1: All probabilities are numbers between 0 and 1 inclusive.
  - A2: If a proposition is a tautology (trivially or logically true), then it has a probability of 1
  - A3: If h and h* are exclusive alternatives (they cannot both be true), then P(h or h*) = P(h) + P(h*)
  - A4: P(h|j) = P(h&j)/P(j), provided that P(j) > 0
    - Bayes' Theorem follows from A4
- Why these axioms? Because if you don't follow these axioms, you leave yourself open to taking bad bets… bets that GUARANTEE a loss.
  - "Dutch bookies" (p. 207 of PGS)

Determining the Credence of the Ideally Rational

- Bayesians (subjectivists) understand credences in an ideally rational person in terms of how they make bets in gambling situations.
  - Specifically, their credence in h is determined by what they take to be (subjectively) fair odds.
- **Subjectively fair odds**: the odds that would be offered by the bookie that would make the (ideal) gambler willing to take either side of the bet.
    - (the odds at which the gambler thinks she cannot take advantage of the bookie)
  - To be willing to bet on h at (fair) odds of X:N (e.g., 3:1) is to be willing to risk losing $X if h is false, in return for a gain of $N if h is true.
  - When X is big and N is small, that indicates a lot of confidence that h is true.
  - If your subjectively fair odds for a bet on h are X:1, then your degree in belief in h is X/(X+1)
- To generalize: P(h) = X/(X+N) (where X:N are the subjective fair odds)

The Arbitrariness Problem for Bayesianism

- The main problem facing Bayesians is that the subjective fair odds can differ between ideally rational people.
  1. The subjective fair odds that one begins with are completely arbitrary.
  2. If the odds are arbitrary, then the credences based off them are arbitrary
  3. If the credences are arbitrary, and credences are probabilities (as the subjectivist Bayesian claims), then the prior probabilities in Bayesian calculations are arbitrary.
  4. If Bayesianism is correct, then e’s providing evidence for h is determined by the prior probabilities
  5. Thus, if Bayesianism is correct, then e’s providing evidence for h is arbitrary!
     - Science is not supposed to be arbitrary: whether something provides evidence for something else should not be arbitrary.

4. Objective Probabilities

Probabilities as Aspects of the World

- Although Bayesians are usually subjectivists, one might think to reap the benefits of Bayes' theorem without regard for degrees of belief.

- The *Objectivist* says that the probabilities are somehow related to the world itself, independent from what human's believe.
  - E.g., That there is a mind-independent fact of the matter whether or not a fair-coin toss has a 50/50 chance of landing on heads.

- But what are objective chances?

Frequentism and Objective Chance

- The popular idea is to understand chance in terms of the number of actual occurrences divided by the number of possible occurences, throughout the entire history of the universe (from beginning to end).

- **Frequentism**: the probability that h = the number of actual occurrences divided by the number of possible occurrences.
  - P(h) = #h/#possible h's.
  - E.g. a fair coin will land heads, P(heads) = the number of times a tossed coin actually lands heads divided by the number coin tosses that have, and will ever, be made.

Problems for Frequentism

- Problem 1 (back to induction): Frequentism tells us that probabilities are determined, in part, by what happens in the future.
  - But cannot see into the future, and so to justify our probability claims, we must look only to the past.
  - We would then be assuming that the future will resemble the past, and we've fallen back into the problem of induction (we've been trying to give a theory of evidence that doesn't rely on induction).
- Problem 2 (uncertain evidence): builds on Problem 1.
  - (1) we cannot know, with any certainty, what the actual probabilities are.
    - b/c we can't see into the future.
  - (2) our knowledge that something counts as evidence requires that we know whether or not the posterior probability is greater than the prior probability
  - (3) Thus, we cannot know with any certainty what counts as evidence for what.
    - We have a theory of evidence that fails to provide science with what they need to produce genuine knowledge about the way the world works!

Discussion Questions

- In this lecture, we learned about Bayesianism…
  - An alternative approach to understanding evidence that does not rely on inductive arguments, but rather, mathematical theorems governing conditional probabilities.
  - We also studied its problems.

- Discussion Questions:
  1. How should supporters of the classic view of evidence – hypothetico-deductivism – react to the Glymore Problems and Ravens problem? Explain.
  2. How should the Subjectivist Bayesian react to the abtirariness objection? Which premise should they reject? Or should they abandon subjectivism?
  3. How should the Objectivist react to the problems for frequentism, bearing in mind that the problems are connected to one another?
  4. Can the objectivist position be adequately defended by simply suggesting that prior probabilities are primitive features of reality, not to be analyzed in terms of frequencies? Why or why not?
  5. What view do you prefer: hypothetico-deductivism or Baysianism? Why?