HPSC 1001/1901/2101/2901

WHAT IS THIS THING CALLED SCIENCE?

Semester 2, 2020

Lecture 8: Problems with Evidence and Induction

How the problems about evidence look initially: we have all sorts of questions in science, and theories that might answer them, and we look for data that confirms some theories and disconfirms others.

Why did dinosaurs like *Tyrannosaurus rex* go extinct?
What is the charge on a single electron?
How many teenagers smoke cigarettes?
etc....

If you are a logical empiricist, much or even all of science becomes a *search for generalizations* (or 'laws'). Then we have to say how it can be rational to using particular cases to confirm, or disconfirm, generalizations. You have seen a lot of black ravens, and none of any other colors. Does that confirm the hypothesis that all ravens are black? If so, how does this work?

End of last time: in the mid C20, the philosophical problem of evidence in science became the problem of giving an *inductive logic* – a theory of the logical relations between evidence and conclusions. In particular: showing how particular observations can give us reason to believe conclusions that are generalizations, extending to future cases and other unobserved cases.

* I think just about everything in the paragraph is a mistake. But this is how it was approached. And the problems that arise are very interesting even if you have a different view of evidence.

Today: Ravens problem and Goodman's problem.
Two problems that can be (and were) presented with very simple toy cases, but have deep implications.
These were disturbing at the time (mid C20). It's not that they were seen as destroying the LE project, or the project of describing an inductive logic. But they were trouble.
They show how what looked easy turned out to be very hard.

What might an "inductive logic" look like?

Perhaps: every observed case of an *F* that is also *G* provides some support for the generalization that all *F*s are *G*.

That looks like a start.

Ravens Problem (Hempel)

Three ideas that each look fine, considered individually:

(1) All observations of particular black ravens confirm the generalization that all ravens are black.

(2) Any evidence that confirms a hypothesis also confirms any hypothesis that is *logically equivalent* to it.

(3) "All ravens are black" is logically equivalent to "All non-black things are not ravens."

* Logical equivalence: If two sentences are logically equivalent then it's *not possible* for one to be true and the other to be false. ("Possible" is meant here in the broadest sense). "Sam is here and Jo is away" is logically equivalent to "Jo is away and Sam is here." They say exactly the same thing, using different arrangements of words.

Put the three assumptions together.
A white shoe is a non-black thing, and a non-raven.
So it is an instance of this generalization: "All non-black things are not ravens."
(* An "instance" of a generalization is a case that fits it.)
But that hypothesis is logically equivalent to the hypothesis that all ravens are black.
So a white shoe confirms the hypothesis: *All ravens are black.*
You can either accept that conclusion or discard at least one of (1)-(3) above.

I.J. Good: reject (1). Whether an instance of a generalization confirms it depends on background knowledge.
This should have been obvious? It is an application of a holistic view of testing.

Various people, including me: reject (1). It depends on the way the observations were made.
Was there a *genuine test*?

You can also reject (2) or (3).

Or you can accept the conclusion that observations of white shoes confirm that all ravens are black.

Goodman's problem

Introduce it in reverse order from Chapter 3 of *T&R*.

A familiar practical problem in science: the curve-fitting problem.

Suppose you have x, y points.

(Might be: x = years of education, y = lifetime income; or: x = fat level in diet; y = heart attack risk....).

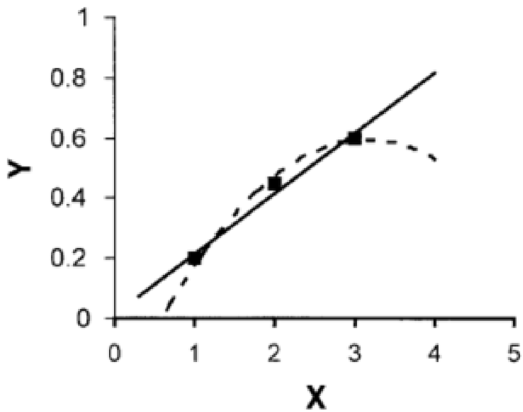

Fig. 3.2
The curve-fitting problem

For any set of points, there is more than one curve that will fit the set (either exactly or approximately). How should you extrapolate? What should you expect for the presently unobserved case when the value of x=5?

Replies. Simplicity? Certainly used in practice.
Why should a simpler curve be preferred?
Role of other background assumptions?
(These might override or support simplicity - eg., you might believe the increase cannot continue linearly forever).

Nelson Goodman, *Fact, Fiction, and Forecast*, 1955:
This problem is *everywhere*, and simplicity does not help (perhaps not much, perhaps not at all).

His own set-up. Immediate goal: to show that there can be no "formal" theory of what makes some inductions good and some bad.

Deductive logic: goodness of an argument is entirely a matter of its *form*, not its particular content.

All *F*s are *G*.	Premises
a is an *F* .	
a is *G*.	Conclusion

The argument is good no matter what you put in for "*F*" and "*G*" (as long as you put in the same thing consistently).

Goodman: induction does not work like that.

All the emeralds observed prior to the year 2021 have been *G*.
All emeralds are *G*

Don't worry about the fact that the conclusion should be qualified; this is not a deductive argument. You can add 'probably' or whatever you like.

But now:

An object is *grue* if and only if it was first observed before the year 2021 and is green, *or* if it was not first observed before 2021 and is blue.

* Note that grue objects do <u>not</u> change color in some mysterious way. Lots of ordinary objects are grue.

The argument above seems fine when *G*=green but bad when *G*=grue. *But the form of the argument is the same.* The same data (the same pile of emeralds) seems to point us towards two contradictory conclusions, depending on how we describe the emeralds (as green or as grue).

That is stage 1 of Goodman's problem. Good and bad inductions can have the same form.

Note that "grue" works perfectly well in *de*ductive arguments. Induction works differently. But how?

Can we exclude inductions using the term "grue" because it is defined in a way that makes reference to time?
Goodman: no.

Another new word:

> An object is *bleen* if and only if it was first observed before the year 2021 and is blue, *or* if it was not first observed before 2021 and is green.

We can use the English words "green" and "blue" to define "grue" and "bleen."

If we do so we must build a reference to time into the definitions of "grue" and "bleen." But suppose we spoke a language that was like English except that "grue" and "bleen" were basic, familiar terms and "green" and "blue" were not. Then we would need to define "green" in a way that involves a reference to time. (A *green* object is one that was first observed before 2021 and is grue, or was not first observed before 2021 and is bleen).

Goodman's own response to all this: Goodness of inductive arguments is language-relative.
Simplicity, also, is language-relative.

---

* My own view: I think the right answer here is related to the right answer to the ravens problem. The relevance of an observation (including instances of generalizations) depends, in many cases, on the procedures being followed and other aspects of the background. (More exactly, there is part of the grue problem that is related to the ravens problem and another part that involves simplicity or something related.)

See "Induction, Samples, and Kinds" on my website as well as Ch 14 of T&R.
See also the book edited by D. Stalker called *Grue!* Frank Jackson's 1975 paper "Grue" is good.