Induction and Lotteries

Jeremy Goodman (USC)
Johns Hopkins University – October 13, 2022

How is it ever possible to have knowledge (or reasonable belief) that goes beyond what is entailed by your evidence?

- For example, how can we know that an observed regularity continues to hold in unobserved cases, and does so non-accidentally?

A toy model

Heading for Heads
A bag contains two coins: one is fair, one is double-headed. You select a coin at random. Rather than inspecting it, you decide to flip it 100 times and record how it lands. In fact, the coin is double-headed. (Goodman and Salow, 2021, forthcoming)

If inductive knowledge possible at all, it is possible here: *after seeing the coin land heads every time, you know that it is not fair.* But how?

The normality framework

Basic Idea: We have a default entitlement to believe that things aren't too *abnormal*. In the good case, these beliefs are knowledge.

Inductive Belief
If a possibility is sufficiently less normal than any other possibility compatible with your evidence, then it is reasonable to believe that it does not obtain.

Inductive Knowledge
If a possibility is sufficiently less normal than actuality, then you can know that it does not obtain.

Consider *Heading for Heads*. The actual situation, in which you're observing the double-headed coin, is sufficiently more normal than any possibility compatible with your evidence in which you're observing the fair coin and it lands heads every time by coincidence.

So, for every such possibility, you have **Inductive Knowledge** that it does not obtain. Combined with your evidence (that the coin lands heads every time) you can know that it is double-headed.

A puzzle

Heads Up
I have two coins in a bag. One is fair; the other is double-headed. I reach in and randomly select one of the coins. I'm going to flip it 30 times. For every 10 of these flips, I recruit a different volunteer, who will observe those flips and no others. Each volunteer is informed of the setup, and knows in advance which flips they will observe. Unbeknownst to them, I have selected the double-headed coin. You are one of the volunteers.

Learning
After seeing the coin land heads 10 times you know that it is double-headed.

If you learn, everyone learns
If after seeing the coin land heads 10 times you know that it is double-headed, then after seeing the coin land heads 10 times every volunteer knows that it is double-headed.

Frontloading
For any 10 flips: if after seeing the coin land heads on exactly those flips the corresponding volunteer knows that the coin is double-headed, then beforehand they know the material conditional (the coin will land heads on those flips $\supset$ it is double-headed).

If anyone knows, you know
For any 10 flips: if beforehand any volunteer knows the material conditional (the coin will land head on those flips $\supset$ it is double-headed), then you also know that conditional beforehand.

Equivalence
If p and q are obviously logically equivalent given the setup, then you know p if and only if you know q.

Generalization
If every volunteer is someone who you know won't see all heads by coincidence, then you know that every volunteer won't see all heads by coincidence.

These six principles have the following absurd consequence:

Too much knowledge
Beforehand you know the conditional: the coin is fair $\supset$ it will land heads on at most 9 out of 30 flips.

My proposal

- I think we should reject **If anyone knows, you know**.
- At the start, each volunteer knows that they won't see all heads by coincidence. But they don't know the same about volunteers who will observe none of the same flips as they will.
- This requires tweaking **Inductive Knowledge**: we have to make comparative normality *agent-relative*.

Here's the idea. The actual situation, in which the double-headed coin was selected, is sufficiently more normal *for you* than any other situation compatible with the setup in which the fair coin was selected and will land heads on every flip you will observe.

Now consider ten flips you won't observe. The actual situation isn't sufficiently more normal *for you* than every situation compatible with the setup in which the fair coin was selected and will land heads on each of those flips. But is it sufficiently more normal than every such situation *for the person who will observe exactly those flips*.

The lottery paradox

> If it is a fair lottery, can one *know* this particular ticket will not be the winning ticket? This can seem wrong. [...]
>
> Suppose Bill wants to know where Mary will be tomorrow. Bill knows that Mary intends to be in New York. Bill also knows that if Mary's ticket is the winning ticket, she will instead be in Trenton for the award ceremony. But there is only one chance in a million of that. Can't Bill conclude that Mary will be in New York tomorrow and in that way come to know where Mary will be tomorrow? That seems possible. But doesn't it involve knowing her lottery ticket is not going to be a winning ticket? (Harman, 1986)

We ordinarily take ourselves to know that our lives will go in certain ways. A consequence of our lives going these ways is that we won't find ourselves in certain very rare circumstances.

But, given our evidence, very many people are about as likely as we are be in such circumstances. And we think (or at least don't rule out) that at least one of us *will* find ourselves in such circumstances.

Reflecting on this fact, it is tempting to deny that we know as much about our lives as we ordinarily take ourselves to know.

Extending the proposal

Most lottery entrants can know about themselves that they'll lose. But they can't know the same about most of their fellow entrants.

- Who can I know will lose the lottery? *Anyone such that it would be abnormal for me that they win.*

  If the lottery is big enough, merely having thought about a person might be enough for me to know they won't win (if this makes me more closely connected to them than to the vast majority of other entrants).

  This turns out to be in keeping with ordinary people's judgments about these cases; cf. Phillips and Kratzer (2020).

  The big picture is that we have enough knowledge to run our lives, yet the totality of what we know about a lottery's outcome still has high probability given our evidence.

References

Jeremy Goodman and Bernhard Salow. Knowledge from probability. In Joseph Y. Halpern and Andres Perea, editors, *Theoretical Aspects of Rationality and Knowledge 2021 (TARK 2021)*, pages 171–186, 2021.

Jeremy Goodman and Bernhard Salow. Epistemology normalized. *Philosophical Review*, forthcoming.

Gilbert Harman. *Change in View: Principles of Reasoning.* Cambridge, MA: MIT Press, 1986.

Jonathan Phillips and Angelika Kratzer. unpublished experiments on knowledge attributions in lottery cases. 2020.