Normality Formalized

A Probabilistic Theory of Inductive Knowledge

Jeremy Goodman (USC); joint work with Bernhard Salow (Oxford) — uAnalytiCon, May 14, 2021

based on "Knowledge from Probability" (forthcoming) https://jeremy-goodman.com/KnowledgeFromProbability.pdf

How can we have knowledge that goes beyond what we have observed?

The skeptical answer: we can’t

Nelson Goodman (1955) is representative:

I suppose that the problem of justifying induction has called forth as much fruitless discussion as has any half-way respectable problem of modern philosophy.

Come to think of it, what precisely would constitute the justification we seek? If the problem is to explain how we know that certain predictions will turn out to be correct, the sufficient answer is that we don’t know any such thing.

Nor does it help matters much to say that we are merely trying to show that or why certain predictions are *probable*. Often it is said that while we cannot tell in advance whether a prediction concerning a given throw of a die is true, we can decide whether the prediction is a probable one. But if this means determining how the prediction is related to actual frequency distributions of future throws of the die, surely there is no way of knowing or proving this in advance.

obviously the genuine problem cannot be one of attaining unattainable knowledge or of accounting for knowledge that we do not in fact have.

Three examples of inductive knowledge

Case 1: Scientific measurements

- Bill steps on a scale. It reads 83.4kg. But the scale isn't perfectly reliable. Even if he does weigh exactly 83.4kg, I don't know that he does. But I do learn something about his weight. For example, I at least know that he weights between 80 and 87kg. How is such knowledge possible?

Case 2: predicting the future

- I’m watching a basketball game between the Lakers and the Warriors. I don’t know who will win. But I do know that they will each score between 30 and 200 points. How is such knowledge possible?

Case 3: lawful regularities

- Rocks fall when we drop them. Of course, no matter how many times we drop a rock and see it fall, it won’t logically follow from what we have observed that rocks fall every time they are dropped. And yet by doing enough such experiments we can learn that rocks fall every time they are dropped. How is such knowledge possible?
    - This is the kind of case I will talk about today.
    - We discuss the other two kinds of cases in the paper this talk is based on.

Humean skepticism about induction

"If reason determined us [in making inductive inferences], **it wou'd proceed upon that principle, that instances, of which we have had no experience, must resemble those, of which we have had experience, and that the course of nature continues always uniformly the same**. In order therefore to clear up this matter, let us consider all the arguments, upon which such a proposition may be suppos'd to be founded; and as these must be deriv'd either from knowledge or probability, let us cast our eve on each of these degrees of evidence, and see whether they afford any just conclusion of this nature." -Hume *A Treatise of Human Nature* (1736, 1.3.6)

If this is our task, then it is hopeless

> I suppose that the problem of justifying induction has called forth as much fruitless discussion as has any half-way respectable problem of modern philosophy. The typical writer begins by insisting that some way of justifying predictions must be found; proceeds to argue that for this purpose we need some resounding universal law of the Uniformity of Nature, and then inquires how this universal principle itself can be justified.

- **Problem**: there can be no such principle. We know that induction hasn't always worked. So any principle that implies that it would have worked when it didn't is a principle we know is false, and so are not justified in believing.

- Inductive knowledge doesn't work by deducing conclusions from one's new observations and some background assumption that "nature is uniform", or any other single principle. Few contemporary epistemologists disagree.

Inductive Anti-Dogmatism

- But most contemporary philosophers *do* think that, if you gain inductive knowledge of a hypothesis H on the basis of new observations E, then before making those observations you could know the conditional: If E, then H.

- Call this idea **Inductive Anti-Dogmatism**.
(I will give a more precise formulation of it later.)

- To see how orthodox this principle is, consider Jim Pryor's “dogmatism” about perception. A dogmatist about perceptual knowledge thinks it is possible to learn that a wall is red by seeing that it is red, even if, before looking at it, you couldn't know that, if it would look red, then it is red.

  - Pryor's thinks this shows that perceptual knowledge isn't a kind of inductive knowledge — that, in such cases, you don't come to know that the wall is red on the basis of noticing that it looks red to you.

Here is the plan

- I’ll by outlining a framework for modeling cases of inductive knowledge and (rational) belief that we have defended elsewhere.
    - This framework explains inductive knowledge in terms of the *comparative normality* of possibilities that are compatible with an agent’s evidence.
- I’ll then give an account of the relevant notion of normality in terms of the probability (on the agent’s evidence) of the answers to a question.
- I’ll then explain how the framework can model inductive knowledge of lawful regularities.
- Finally, I’ll explain how the proposal predicts counterexamples to Inductive Anti-Dogmatism, and defend this prediction.

Preview

- Knowledge and belief are closed under entailment.
- Only propositions with sufficiently high probability are known or believed.
- Knowledge and belief are context-sensitive.
- **Relative to a context, knowledge and belief are fully characterized in terms of evidence and probability, notions which even skeptics about induction usually accept.**

- Even if you aren't convinced by our specific models or by the normality framework, I hope to convince you that epistemic logic is a powerful tool for theorizing about induction. Goodman's defeatism was premature.

Knowledge and Normality

Inductive knowledge of laws: a toy model

Heading for Heads
I know a bag contains two coins: one fair, one double-headed. Without looking, I reach in and select a coin. I decide to flip it 100 times and observe how it lands.

- *Claim*: If the coin lands heads every time, then I should believe it is double-headed. And if it really is double headed, then this belief is knowledge.

- But after seeing the coin land heads 100 times, there are two possibilities compatible with my evidence: (d) that it is double headed, and (c) that it is fair and landed heads every time by coincidence. What makes me able to know/rationally believe that (c) is false?

- *Idea*: (c) is ***sufficiently less normal*** than (d).

    - "Normal" in what sense? We'll come to that, but the idea is that we understand this notion of normality by seeing its role in a theory of induction. ("Plausible" sometimes has more helpful connotations than "normal".)

Normality and induction, knowledge and belief

- If my actual situation is w, and v is sufficiently less normal than w, then I can **<u>know</u>** that I am not in v.

- If it is compatible with my evidence that I am in w, and v is sufficiently less normal than w, then I can (rationally) **<u>believe</u>** that I am not in v.

Cf. **Stalnaker (2006,2015,2019), Goodman (2013), Greco (2014), Dutant (2016), Goodman and Salow (2018,forthcoming), Carter (2019), Beddor and Pavese (2020), Littlejohn and Dutant (2020), Carter and Goldstein (2021), Loets (forthcoming), Goldstein and Hawthorne (forthcoming) on knowledge, and Smith (2010,2016,2017,2018) on belief.**

Normality structures

A *normality structure* is a tuple <S,€,W,≥,»> such that

1. S is a non-empty set (of *states*)
2. € is a set of non-empty subsets of S (the *possible bodies of evidence*)
3. W = {<s,E>: s ∈ E ∈ €} (the set of *(centered) worlds*)
4. ≥ is a reflexive and transitive relation on W
   (read ‘w ≥ v’ as ‘w is at least as normal as v’)
5. » is a well-founded relation on W such that, for all worlds w1,w2,w3,w4:
   (a) If w1 » w2, then w1 ≥ w2.
   (b) If w1 ≥ w2 » w3 ≥ w4, then w1 » w4.
   (read ‘w » v’ as ‘w is sufficiently more normal than v’)

Evidential accessibility

- For any world w, $R_e(w)$ is the set of worlds *evidentially accessible* from w. These are all and only the worlds compatible with your evidence in w.
    - A world v is evidentially accessible from w iff w and v agree on your evidence about the state of the world.

$$R_e(\langle s,E \rangle) = \{\langle s^*,E \rangle : s^* \in E\}$$

(In the paper that this talk is based on we discuss ways of not assuming that evidential accessibility is an equivalence relation. But that assumption is harmless for present purposes.)

Doxastic accessibility

- For any world w, $R_b(w)$ is the set of worlds *doxastically accessible* from w. These are all and only the worlds compatible with what you believe in w.
  - A world v is doxastically accessible from w iff:
    (i) v is evidentially accessible from w and,
    (ii) v isn't sufficiently less normal than any other world evidentially accessible from w.

$$R_b(w) = \{v \in R_e(w): \neg\exists u(u \in R_e(w) \,\&\, u \gg v)\}$$

Epistemic accessibility

- For any world w, $R_k(w)$ is the set of worlds *epistemically accessible* from w. These are all and only the worlds compatible with your knowledge in w.

- Two possible definitions (which one we pick won't matter for present purposes):

  - **Stalnakerian**: A world v is epistemically accessible from w iff:
    (i) v is doxastically accessible from w, or
    (ii) v is evidentially accessible from w and at least as normal as w.

    $R_k(w) := R_b(w) \cup \{v \in R_e(w): v \geq w\}$

  - **Williamsonian**: A world v is epistemically accessible from w iff:
    (i) v is doxastically accessible from w, or
    (ii) v is evidentially accessible from w and at least as normal as w, or
    (iii) v is evidentially accessible from w and less normal than w but not sufficiently less normal than w. [i.e., knowledge requires a "margin for error"]

    $R_k(w) := R_b(w) \cup \{v \in R_e(w): v \geq w\} \cup \{v \in R_e(w): w \geq v \text{ and not } w \gg v\}$

What do normality structures represent?

- We model worlds as pairs <s,E>, where s is the state of the world and E is the agent's evidence about the state of the world. All evidence is true, so s $\in$ E.

- We model knowledge and belief using accessibility relations: the agent *knows/ believes* that p in a world w iff p is *true* in all worlds v epistemically/doxastically accessible from v.

    - It follows that agents know/believe everything that is entailed by things they know/believe. (This is an unrealistic idealization, since real people aren't logically perfect, but it is a reasonable one for theorizing about induction.)

- Normality structures also allow us to model the dynamics of knowledge and belief:

  <s,E*> *is the result of discovering* p in <s,E> iff E* = E $\cap$ p

Normality from Probability

The basic idea: more normal = more probable

- **Problem**: probability is sensitive to how we describe the space of possibilities. How many possible outcomes of the Lakers vs Warriors basketball game are there? Two (depending on who wins)? One for every possible score? One for every possible pattern of scoring between the teams? Does it matter which player made which shots?

- These choices will make a big difference to the probabilities of individual possibilities. But it isn't plausible that we need to decide between these different descriptions when we're theorizing about what we can know/rationally believe about the outcome of the game, as long as the possibilities are fine-grained enough to ask the questions we're interested in.

- **Solution**: The probability of a possibility is always relative to a question. Consider a question Q and a world w. Suppose q is the true answer to Q in w. Then relative to this choice of Q, we can identify the level of normality of w with the probability of q in w. Holding fixed the question Q means we can think about possibilities in more or less fine-grained ways without disruptive epistemological consequences.

Being at least as normal

- Relative to a question Q, the *likeliness* of a world w is the probability, given your evidence in w, of the answer to Q that is true in w.

- *Example*: I know I have just shuffled a deck of cards. Let w be some possibility in which (unbeknownst to me) the nine of hearts is on top of the deck. Relative to the question *what is the suit of the top card of the deck*, the likeliness of w is 1/4. Relative to the question *what is the number of the top card of the deck*, the likeliness of w is 1/13.

  - **Normality as Likeliness**
    w ≥ v iff (i) the likeliness of w is at least as high as the likeliness of v, and
    (ii) v is evidentially accessible from w.

Being sufficiently more normal

Basic Idea: An evidentially accessible world is doxastically inaccessible if and only if it is sufficiently improbable that anything so abnormal happens.

- A world w is *most-normal* iff w is at least as normal as any other world evidentially accessible from w.
- A world v is *abnormal* iff the probability in v that things are at least as abnormal as they are in v is sufficiently low.
- **Abnormality Constraint**
  If w is most-normal and v is evidentially accessible from w, then w is sufficiently more normal than v iff v is abnormal.
- <u>Fact</u>: the Abnormality Constrain implies the “Basic Idea” as long as every world evidentially accesses some most-normal world; and that in turn is implied by Normality as Likeliness.

Normality from Probability: technical details

A *probability structure* is a tuple <S,€,W,Q,P,t> such that

1. S,€,W satisfy 1-3 in the definition of a normality structure
2. Q (the *question*) is a partition of S
3. P (the *prior*) is a probability distribution over S such that P(q|E) is defined for all $q \in Q$ and $E \in$ €
4. $t \in [0,1]$

Normality as Likeliness: $w \geq v$ iff $v \in R_e(w)$ and $\lambda(w) \geq \lambda(v)$

where $\lambda(<s,E>) = P([s]_Q \mid E)$ and $[s]_Q$ = the cell of Q containing

Sufficiency: w » v iff $v \in R_e(w)$ and $1 - \tau(v)/\tau(w) \geq t$

where $\tau(<s,E>) = P(\{s^* : <s,E> \geq <s^*,E>\} \mid E)$

Modeling inductive knowledge of laws

Heading for Heads
I know a bag contains two coins: one fair, one double-headed. Without looking, I reach in and select a coin. I decide to flip it 100 times and observe how it lands.

- There are $2^{100}+1$ states: 1 in which the coin is double headed and one for every pattern of heads and tails the coin could have if fair.
- Let c be the state where the coin is fair and lands heads every time by coincidence, and d be the state where the coin is double headed.
- P(d) = .5; P(s) = $.5^{101}$ for all other states. Let t=.99999.
- After seeing the coin land heads 100 times, my evidence is {c,d}. **If the coin is double headed, then I have inductive knowledge that it is double headed:**
    - $\langle c,\{c,d\}\rangle \notin R_k(\langle d,\{c,d\}\rangle)$, even though $\langle c,\{c,d\}\rangle \in R_e(\langle d,\{c,d\}\rangle)$

What do I know before flipping the coin?

Can I know that it won’t land heads every time by coincidence?

- If the question Q is *is the coin fair, and, if so, how many times will it land heads*, then I will know in advance that it won’t land heads every time by coincidence.
    - This is because, given this choice of Q, $1 - \tau(\langle c,S\rangle)/\tau(\langle d,S\rangle) = 1 - .5^{100} \geq .99999$
- But if the question Q* is *is the coin fair, and, if so, what sequence of heads/tails will have*, then I will ***not*** know in advance that the coin won’t land heads every time by coincidence.
    - This is because, given this choice of Q*, $1 - \tau(\langle c,S\rangle)/\tau(\langle d,S\rangle) = 1 - .5 < .99999$
- This shows how knowledge and belief depend on the contextually supplied question.

Inductive Anti-Dogmatism revisited

- As I mentioned at the beginning, many authors accept the principle:
  - **Inductive Anti-Dogmatism**
    If it is compatible with what you know that p and not-q, then you won't come to know q by discovering p.
- In fact, the principle comes from Dorr, Goodman and Hawthorne (2014), in which we defend it!
- But this principle fails in Heading for Heads for Q*. Before flipping, it is compatible with what I know that the coin will land heads every time by coincidence. But after discovering that it lands heads every time, I come to know that this wan't by coincidence.
  - Formally, the point is that $<c,S> \in R_k(<d,S>)$, and $<d,\{c,d\}>$ is the result of discovering $\{c,d\}$ in $<d,S>$, but $<c,\{c,d\}> \notin R_k(<d,\{c,d\}>)$.

What is going on?

- It can be shown that Inductive Anti-Dogmatism holds in normality structures that satisfy the following condition:

    Evidence Neutrality
    <s,E> ≥ <t,E> iff <s,E'> ≥ <t,E'>

- Something like this principle tends to be built into formal models like normality structures; cf. Goodman and Salow (2018), Smith (2016), and models of belief-revision/non-monotonic reasoning using plausibility orders.

- But this principle generally fails in normality structures derived from probability structures: the comparative normality of two states of the world is always relative to a body of evidence.

More failures of Evidence Neutrality

- You (rationally) believe that your car won't break down in the next day. Let n>1 be the least number such that it is consistent with your beliefs that your car will have broken down that many days from now. n-1 days later, your car still works, and you (rationally) believe it won't break down in the next day.

- This case arguably involves a failure of Evidence Neutrality: for all x, <x,{0,1,…}> is not sufficiently more normal than <n,{0,1,…}>; but, for some x, <x,{n,n+1,…}> is sufficiently more normal than <n,{n,n+1,…}>.

- We give more arguments against Evidence Neutrality in "Epistemology Normalized" (ms), which is an extended case for the normality framework that does not appeal to any analysis of normality (in probabilistic terms or otherwise).

Conclusions

- Bayesian epistemologists often hold that, when it comes to inductive hypotheses, the best we can do is assign them high probabilities: knowledge is out of reach.

- I agree that probability is an invaluable tool for theorizing about the epistemology of induction. But this is not because inductive knowledge is impossible.

- On the contrary, I believe that in many cases we can give attractive non-trivial models of inductive knowledge using probabilistic resources that most Bayesians should find congenial.

- These models offer a new window into principles about the dynamics of inductive knowledge and belief. They suggest that some widely held principle about those dynamics ought to be reconsidered.

Thank you!