Probability and Closure

Jeremy Goodman (Johns Hopkins)

Princeton – November 17, 2023

When theorizing about what people know and rationally believe, contemporary epistemologists often make the idealizing assumptions that knowledge and rational belief are closed under conjunction:

> KCC
> If A knows that p and knows that q, then they know that p and q.
>
> BCC
> If A rationally believes that p and rationally believes that q, then they rationally believe that p and q.

Question: Are these principles any good, even as idealizations?

1 An argument against closure

KCC is inconsistent with the following four attractive principles:

> INDUCTIVE KNOWLEDGE
> You know some propositions whose probability on your evidence is less than 1.
>
> WE KNOW A LOT
> If you know any propositions whose probability on your evidence is less than 1, then the probability on your evidence of the conjunction of all such propositions is low.
>
> ENTAILMENT
> Knowledge entails rational belief.
>
> THRESHOLD
> You rationally believe only propositions whose probability on your evidence is high.

BCC is inconsistent with the following three attractive principles:

> INDUCTIVE BELIEF
> You rationally believe some propositions whose probability on your evidence is less than 1.
>
> WE BELIEVE A LOT
> If you rationally believe any propositions whose probability on your evidence is less than 1, then the probability on your evidence of the conjunction of all such propositions is low.
>
> THRESHOLD
> You rationally believe only propositions whose probability on your evidence is high.

I'm not going to question ENTAILMENT, since the analogue of THRESHOLD for knowledge is independently plausible.

Here is an example illustrating what I kind of knowledge/beliefs I have in mind by INDUCTIVE KNOWLEDGE/BELIEF:

> *Heading for Heads*
> You know 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.

- Claim: you can come to know (and hence rationally believe) that the coin is double-headed by seeing it land heads enough times in a row. (Dorr et al., 2014)

If we are willing to ascribe knowledge such in the above case, we should do the same in the following case:

> *Flipping for Heads*
> You decide to flip a fair coin until it lands heads.

- Claim: for some n, you can know (and hence rationally believe) that the coin will be flipped at most n times. (Dorr et al., 2014)

2 Lockeanism

A simple and popular theory of belief maintains INDUCTIVE BELIEF, WE BELIEVE A LOT, and THRESHOLD by strengthening the latter to the following biconditional:

> LOCKEANISM
> You rationally believe all and only those propositions whose probability on your evidence is high.

In addition to failures of BCC, LOCKEANISM requires rejecting the following three additional plausible principles:

> CONSISTENCY
> The set of propositions that you rationally believe is consistent.

> WEAK BCC
> If A rationally believes that p and rationally believes that q and p and q concern the same subject matter, then A rationally believes that p and q. (Goodman, 2013)

> B=¬E¬K
> A rationally believes that p if and only if it's compatible with their evidence that they know that p.

Examples of failures of CONSISTENCY and B=¬E¬K predicted by LOCKEANISM:

- Believing of every ticket in a lottery that it will lose while believing that one of the tickets will win.
- Believing of every weight (even the most plausible ones) that you don't weigh that much.

Examples of failures of WEAK BCC predicted by LOCKEANISM:

- For some x and y, you rationally believe that you weigh at least x pounds and that you weigh at most y pounds but not that you weigh between x and y pounds.
- For some n, in Flipping for Heads you believe that the coin will be flipped at most n times and that it won't be flipped exactly n times but not that it will be flipped at most $n-1$ times.

On the other hand, an attractive prediction of LOCKEANISM is:

> MODESTY
> You believe that not everything that you believe is true.

3 The normality framework

Here is a simple version of the theory of inductive knowledge and (rational) belief defended in Goodman and Salow (2023).

- Worlds compatible with your evidence differ in their comparative *normality(/plausibility)*:
  - $w \succeq v$ (w is *at least as normal as* v) is reflexive and transitive.
  - $w \ggcurly v$ (w is *sufficiently more normal than* v) is asymmetric, entails being at least as normal, and is extended by being at least as normal (i.e., $w \ggcurly v$ whenever $w \succeq w' \ggcurly v' \succeq v$).
- v is *doxastically accessible* from w iff v is compatible with your evidence in w and v isn't sufficiently less normal than any other u compatible with your evidence in w.
- v is *epistemically accessible* from w iff v is compatible with your evidence in w and is either doxastically accessible from w or is at least as normal as w.
- You know(/believe) that p in world w iff p is true in all worlds v that are epistemically(/doxastically) accessible from w.
  - This guarantees KCC and BCC.
  - Crucially, not all true (justified) beliefs are knowledge – this is because non-actual doxastically inaccessible worlds will be epistemically accessible whenever they are compatible with your evidence and at least as normal as actuality.
- Only truths can be evidence, and (as a simplifying idealization) let's assume that your evidence is the same in all worlds compatible with your evidence.

The normality framework entails none of the worrying features of LOCKEANISM, and (appropriately understood) vindicates INDUCTIVE KNOWLEDGE/BELIEF. So it must reject either WE KNOW/BELIEVE A LOT or reject THRESHOLD.

As we will see, both options are possible. To illustrate them, consider the following hybrid of Flipping for Heads and the Preface Paradox (Makinson, 1965):

> *Racing for Heads*
> Each of n coin flippers has a fair coin. Each will flip their coin until it lands heads. (Goodman and Salow, 2018, 2021)

3.1 Multi-dimensionalism

The basic idea is to think of worlds as differing in normality along many independent dimensions: in this case, one for each coin. Along each dimension, it is less normal for the coin to take longer to land heads, and sufficiently less normal for the coin the coin to take at least m flips longer to land heads. Let m is the least number such that $1 - .5^m$ counts as a 'high' probability.[1]

- $w \succeq v$ iff w is at least as normal as v along every dimension.
- $w \ggcurly v$ iff w it is at least as normal as v and is sufficiently more normal along at least one dimension.

THRESHOLD fails. Every world compatible with your evidence in which any of the coins is flipped more than m times is sufficiently less normal than the worlds compatible with your evidence in which every coin lands heads on the first flip. So no such worlds are doxastically accessible. So you believe that every coin will land heads in at most m flips. But this cannot have high probability on your evidence (assuming there are at least three coins being flipped[2]).

3.2 Probabilism

The basic idea is that, relative to a contextually determined question, we can analyze comparative normality (and hence knowledge and belief) in terms of probabilities.

- $w \succeq v$ (relative to Q) iff w and v agree on your evidence and the probability on this evidence of the answer to Q that is true in w is at least as high as the probability on this evidence of the answer to Q that is true in v.
- $w \ggcurly v$ (relative to Q) iff the probability on the evidence that things are more normal (relative to Q) than v, conditional on things being no more normal (relative to Q) than w, is high.

Fortunately, probabilism allows us to characterize belief directly in terms of probabilities, without mentioning comparative normality:

> The *plausibility* of w (relative to Q) = the probability, given your evidence in w, of the answer to Q that is true in w.

[1] If this threshold $t = .99$, then $m = 7$, since $1 - .5^6 < .99 < 1 - .5^7$.

[2] If we assume that $t > .5$, then $m \geq 2$, so $(1 - .5^m)^3 < 1 - .5^{m-1}$.

> **Fact**: Given probabilism, the set of worlds doxasically accessible from w can be determined by the following procedure. First, include the most plausible (relative to Q) worlds compatible with the evidence in w. If these worlds don't yet have high probability on that evidence, add in all of the most plausible remaining worlds compatible with that evidence. Repeat until the set of included worlds has high probability. Then stop.

Probabilism validates BCC, INDUCTIVE BELIEF, and THRESHOLD. So WE BELIEVE A LOT fails. How it fails depends on the question Q.

Relative to the question *how many times will coin i be flipped*, what you believe about coin i is the same as what you would believe about it according to multi-dimensionalism; but you won't have any non-trivial beliefs about any of the other coins.

The situation is more interesting for other kinds of questions, such at (i) *what will the exact outcome of the whole experiment be*; (ii) *what will the shape of the outcome be* – that is, how many coins will take how long to land (the exact outcome up to isomorphism); (iii) *how many total tails will there be in the experiment as a whole*; (iv) *how long will it be before all the coins have landed heads*; and (v) *how many of the coins will ever land heads at the same time.* For each of these question, we can ask what you believe about a number of issues, such as how many total tails there will be, how many trials the experiment will take (i.e., the number of times that the coin that is flipped the most times will be flipped), and whether all the coins will land heads at once (a claim labelled 'same end' below). The table below records what probabilism says to believe in advance for different choices of Q, for 'high probability' thresholds $t = .75$ and $t = .95$. For more details, including the generalization to continuous probability distributions, see Goodman and Salow (2021).

Q	which worlds are most normal	t	min tails	max tails	min trials	max trials	same end?
(i) exact outcome	all coins land heads first time	.75	0	13	1	14	maybe
		.95	0	18	1	19	maybe
(ii) outcome shape	6×1 flip, 3×2 flips, 1×3 flips	.75	1	15	2	8	no
		.95	0	22	1	12	maybe
(iii) how many total tails	8 or 9 total tails [tied]	.75	5	14	1	15	maybe
		.95	2	18	1	19	maybe
(iv) how long until over	ends on 4th trial	.75	2	50	3	6	maybe
		.95	1	70	2	8	maybe
(v) how many end together	5 flippers get heads at once	.75	3	∞	2	∞	no
		.95	2	∞	2	∞	no

4 Synthesis: probability spheres

- Start with a reflexive, transitive *at least as normal* relation $\succeq$.
  - This is determined as in multi-dimensionalism.
  - For each dimension, normality along that dimension is determined as in probabilism.
  - That is, comparative normality is determined probabilistically from a *set* of questions, one for each dimension.
  - These questions correspond to the pretheoretical notion of subject matter in WEAK BCC. (Lewis, 1988).
- A *sphere* is a set of worlds that contains any world at least as normal as any it contains.
- You *believe* p in w iff for some sphere S, S has high probability on your evidence in w and p is true in every member of S that is compatible with your evidence in w.
- You *know* p in w iff for some sphere S, $w \in S$, S has high probability on your evidence in w, and p is true in every member of S that is compatible with your evidence in w.

Some Facts

1. Probability spheres and LOCKEANISM agree when all worlds are incomparable in normality.
2. Probably spheres and probabilism are equivalent in the special case where there is only one dimension.
3. What you believe about a dimension is independent of what other dimensions there are (in contrast to multi-dimensionalism).
4. CONSISTENCY holds whenever some world compatible with your evidence is at least as normal as every other. So it holds in Racing for Heads modeled with n dimensions, one for each coin.
5. CONSISTENCY fails in Racing for Heads if we add an $n+1$st dimension corresponding to whether all of the coins will land heads in the first m flips, provided there are sufficiently many coins that it has high probability that at least one coin won't.[3]
6. KCC fails too.[4]

[3] Compare Carter and Hawthorne (2021), who similarly use sets of questions to invalidate BCC, but do so in a way that also invalidates closure under entailment.

[4] Carter and Goldstein (2021) model preface cases using a variant of multi-dimensionalism for knowledge (using a more aggressive definition of epistemic accessibility than Goodman and Salow (2023)), and define belief as epistemically possible knowledge. BCC fails in this model, but KCC and CONSISTENCY hold.

Enforcing consistency: grounded probability spheres

- A sphere is *grounded* at w iff every world compatible with your evidence in w that isn't in the sphere is less normal than some world compatible with your evidence in w that is in the sphere.
- We can then ensure CONSISTENCY by replacing that 'sphere' with 'sphere grounded at w' in the definitions of knowledge and belief.

Scorecard

<u>LOCKEANISM</u>

✓ INDUCTION, A LOT, THRESHOLD, MODESTY

✗ CONSISTENCY, WEAK BCC, B=¬E¬K

<u>Multi-dimensionalism</u>

✓ INDUCTION, A LOT, BCC, CONSISTENCY, B=¬E¬K

✗ THRESHOLD, MODESTY

<u>Probabilism</u>

✓ INDUCTION, THRESHOLD, BCC, CONSISTENCY, B=¬E¬K

✗ A LOT, MODESTY

<u>Probability spheres</u>

✓ INDUCTION, A LOT, THRESHOLD, WEAK BBC, B=¬E¬K, MODESTY↔¬CONSISTENCY

✗ BCC, MODESTY, CONSISTENCY

<u>Grounded probability spheres</u>

✓ INDUCTION, A LOT, THRESHOLD, CONSISTENCY, WEAK BBC, B=¬E¬K

✗ BCC, MODESTY; Facts 1&3 for (general) probability spheres

5 How bad are closure failures?

1. 'A knows that p and A knows that q, but A doesn't know that p and q', sounds bad. But this is to be expected. The threshold for 'high' probability is presumably vague. As long as it's vague enough, there won't be any determinately true instances of the above schema, and hence no felicitously assertible or even supposable instance. (Compare: 'Suppose that she was still a child a second ago but isn't a child now – not that anything dramatic happened in that second'.)
2. What about quantified claims like 'Even if these are all students who you know will pass, you still don't know that none of them will fail'? For all I've said this could be determinately true. While it admittedly sounds odd, I don't think it's a disaster.

References

Sam Carter and Simon Goldstein. The normality of error. *Philosophical Studies*, 178(8):2509–2533, 2021.

Sam Carter and John Hawthorne. Normality. July 2021.

Cian Dorr, Jeremy Goodman, and John Hawthorne. Knowing against the odds. *Philosophical Studies*, 170(2):277–87, 2014.

Jeremy Goodman. Inexact knowledge without improbable knowing. *Inquiry*, 56 (1):30–53, 2013.

Jeremy Goodman and Bernhard Salow. Taking a chance on KK. *Philosophical Studies*, 175(1):183–96, 2018.

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*, 132(1):89–145, 2023.

David Lewis. Statements partly about observation. *Philosophical Papers*, 17(1): 1–31, 1988.

D. C. Makinson. The paradox of the preface. *Analysis*, 25(6):205–7, 1965.