How to Tell Time with a Broken Clock

Jeremy Goodman and Bernhard Salow 9 January 2025

Inductive Knowledge

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.

Asymmetry: After seeing the double-headed coin land heads 100 times in a row, you can know that it isn't fair. But had you instead selected the fair coin and seen it land heads 100 times in a row, you wouldn't be in a position to know that it wasn't double-headed.

Lots of Heads
A bag contains three coins: one is fair, one is double-headed, and one is 90% biased in favor of heads. You select a coin at random. Rather than inspecting it, you decide to flip it 100 times and record how it lands. The coin lands heads every time.

Gettier cases: You now justifiably believe that the coin is double-headed. If it is, then you know that it is. And if it's the biased coin, you still know that it's not fair. But if it's the fair coin, you don't know that it's not biased, despite justifiably and truly believing this.

The written version of this paper argues that this pattern of knowledge and ignorance isn't captured by the idea that you can know that a possibility doesn't obtain provided that possibility is 'far away', which is a popular way of developing the idea that beliefs are knowledge when they are safe from error.

Normality

Two relations of comparative normality/plausibility:

- w is *at least as* normal/plausible as v
- w is *sufficiently more* normal/plausible than v

Three ideas:

INDUCTIVE KNOWLEDGE: If w is sufficiently more normal than v, then in w you can know that you are not in v.

CONVEXITY: If in w for all you know you're in v, u is at least as normal as v, and it's consistent with your evidence in w that you're in u, then in w for all you know you're in u.

BELIEF: You have justification to believe p when your total evidence is E if and only if p is true in every possibility consistent with E that is not sufficiently less normal than any other possibility consistent with E.

Together, these explain **Heading for Heads** and **Lots of Heads**.

If K = JTB + CONVEXITY we get:

K-NORMALITY: You know that p if and only if p is true in all possibilities compatible with your evidence that are either (i) not sufficiently less normal than any other such possibility, or (ii) at least as normal as the situation you are actually in.

These relations have the formal structure suggested by the terminology:

- *at least as normal as* is transitive and reflexive;
- *sufficiently more normal than* is transitive and asymmetric;
- if one possibility is *sufficiently more normal than* another, it is also *at least as normal as* it;
- the relations can be chained: if w_1 is at least as normal as w_2 which is sufficiently more normal than w_3 which is at least as normal as w_4, then w_1 is sufficiently more normal than w_4.

We'll also assume that *sufficiently more normal than* is well-founded, in the sense that any non-empty set has at least one member that is not sufficiently less normal than any other.

Gettier Cases

K-NORMALITY is well-placed to handle many standard Gettier cases:

Beddor and Pavese [2020, §5], Goodman and Salow [2023, §4]

- It would have been at least as normal for Brown to be in Lisbon as it was for him to be in Barcelona [Gettier, 1963].
- It would have been at least as normal for the stopped clock to be displaying the wrong time as for it to be displaying the correct one [Russell, 1949].
- It would have been at least as normal for Henry to be looking at a fake barn in fake barn county as for him to be looking at a real barn in fake barn county [Goldman, 1976].
- Gloss: a belief that p amounts to knowledge just in case its truth is *appropriately connected to* its justification, in the sense that p is true throughout the smallest 'connected' region of logical space that includes actuality as well as every possibility compatible with what the agent has justification to believe.

A Challenge

> **Broken Clock**
> It feels about noon-ish, and you can see a queue forming outside the cafe, which is slightly more likely to happen on the hour than at other times. You then look at a clock, which reads '12:00'. While it is, in fact, 12:00, the clock you looked at stopped running a day ago.

Goodman and Salow [2023, note 23]

Intuitively, you don't know that it's 12:00. But K-NORMALITY seems to say otherwise: for it would seem that it is 12:00 in every possibility consistent with your evidence that is least as normal as actuality.

Possible escape: coarse-grain normality and insist that it is exactly as normal for the queue to form 11:55 as at 12:00, because they are in the same coarse normality-level [Stalnaker, 2019]. This may work in some cases; but why think that the distinction between coarse normality-levels *never* falls between these two possibilities?

Margins to the rescue?

The most straightforward response to this challenge would be to find a natural normality-theoretic principle that entails that you don't know what time it is in **Broken Clock**. For example:

> MARGINS: If v is less normal but not sufficiently less normal than w, and consistent with your evidence in w, then in w for all you know you're in v.

Carter and Goldstein [2021], Goldstein and Hawthorne [2022], Goodman and Salow [2018, 2021, 2023], Hong [2023]

But this also has costs:

> Bernhard: it predicts knowing without knowing that one knows, knowing that one knows without knowing that one knows that one knows, etc.

Williamson [2000, 2011, 2014]

> Jeremy: we might want knowledge to require high objective chance, and this can be achieved by tweaking the probabilistic account of normality in Goodman and Salow [2021]. But, having done so, MARGINS is overkill, and predicts Moore-paradoxical beliefs of the form *p but I don't know that p*.

Goodman [in preparation]

So let's see if there is another option!

Purely Veridical Gettier Cases

Purely Veridical Broken Clock
It feels about noon-ish, and you can see a queue forming outside the cafe, which is slightly more likely to happen on the hour than at other times. You then look at a clock, which reads '12:00'. You recall that this clock was broken when you last visited a year ago; you think it very likely that maintenance would have fixed it in the meantime, but aren't completely sure. While it is, in fact, 12:00, the clock is still broken.

Claim: one can flesh out the details so that you have justification to believe that it is 12:00, but fall short of having justification to believe the clock is working. If the clock were working, you would have come to know that it is 12:00 without knowing that the clock is working.

Some observations:

- **Purely Veridical Broken Clock** being a Gettier case is inconsistent with the JK (a.k.a. strong belief) principle (that whenever you have justification to believe that p, you have justification to believe that you know that p), and the KK principle (that whenever you know that p, you know that you know that p).

Cf Williamson [2013], Hawthorne [ms]

- Everyone should accept that it is possible to come to know the time by reading a clock that you don't know works. Perhaps this takes some of the sting out of ascribing knowledge in **Broken Clock**.
- Plausibly, **Broken Clock** and **Purely Veridical Broken Clock** stand or fall together: either both are cases of knowledge or both are Gettier cases. So we can defend attributing knowledge in the former by defending attributing knowledge in the latter. This will be our strategy in what follows.

Telling Time with a Broken Clock

Two Clocks
You have two clocks, and you know that one of them works and one of them is broken, but you don't know which is which. You look at the first clock, and it reads 12:00. You then look at the second clock, and see that it also reads 12:00. In fact, the first clock is the one that works and the second clock is the one that is broken.

So it's not impossible to tell time with a broken clock!

Two Clocks is similar to **Purely Veridical Broken Clock**, in that you have two 'pieces of evidence' suggesting that it is 12:00, only one of which is 'tethered' to the truth, but both of which are required for justification. On the other hand:

- In **Two Clocks**, the fact that your untethered evidence is untethered is not a *defeater* for your justification.
- In **Two Clocks**, the tethered evidence is a significant portion of your overall justification, while in **Purely Veridical Broken Clock**, it contributes only marginally.

Can we remove these two disanalogies?

From Clocks to Coins

Two Coins
You know that you are holding two coins: a fair coin and a trick coin, which is either double-headed or double-tailed. But you don't know which coin is fair, nor whether the trick coin is double-headed or double-tailed. You flip the first coin, then the second. Both land heads. In fact, the first coin was double-headed, and the second coin was fair.

Since the question is which trick coin you have, flips of the trick coin correspond to readings of the working clock, while flips of the fair coin correspond to readings of the broken clock.

Two Coins in a Bag
You have a bag that you know contains two coins: one is fair, and the other is a trick coin, either double-headed or double-tailed, but you don't know which. In fact, it is double-headed. You take a coin from the bag at random, flip it, record how it landed, put it back in the bag, and repeat many times. Everything unfolds unremarkably: you observe about 75% heads, about 33% of which come from the fair coin.

You eventually come to know that the trick coin is double-headed. This can happen (i) at the same time as you come to have justification and (ii) as the result of the flip of the fair coin.

If (i) couldn't happen, you could have justification while knowing you don't know. To motivate (ii), consider a variant where many fair and trick coins are flipped all at once.

- Learning that you just flipped a fair coin *would* defeat your justification for believing that the trick coin is double-headed.

Finally, we have:

Many Coins in a Bag
You have a bag that you know contains 50 coins: 49 are fair, and one is a trick coin, either double-headed or double-tailed, but you don't know which. In fact, it is double-headed. You take a coin from the bag at random, flip it, record how it lands, and put it back in the bag, and repeat many times. Everything unfolds unremarkably: you observe about 51% heads, and over 98% of these are the result of flipping fair coins.

Eventually, you still come to know that the trick coin is double-headed. And this can still happen (i) at the same time as you come to have justification and (ii) as the result of the flip of a fair coin.

- Learning that you just flipped a fair coin *would* defeat your justification for believing that the trick coin is double-headed.
- Most of the observed heads that justify your belief that the trick coin is double-headed come from flips of fair coins.

Conclusion

At first blush, **Broken Clock** and **Purely Veridical Broken Clock** feel like Gettier cases. If they really are, then that is a challenge for K-NORMALITY and, more generally, for the KK and JK principles.

But **Purely Veridical Broken Clock** is puzzling for everyone. And it is closely analogous to **Two Clocks/Coins**, **Two Coins in a Bag**, and **Many Coins in a Bag**, all of which are cases in which one knows. So maybe you can know in **Purely Veridical Broken Clock**. And if you can know in **Purely Veridical Broken Clock**, you can know in **Broken Clock** as well.

References

Bob Beddor and Carlotta Pavese. Modal virtue epistemology. *Philosophy and Phenomenological Research*, 101:61–79, 2020.

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

Edmund Gettier. Is justified true belief knowledge? *Analysis*, 23:121–123, 1963.

Alvin Goldman. Discrimination and perceptual knowledge. *The Journal of Philosophy*, 73:771–791, 1976.

Simon Goldstein and John Hawthorne. Knowledge from multiple experiences. *Philosophical Studies*, 179:1341–72, 2022.

Jeremy Goodman. Knowledge and chance. in preparation.

Jeremy Goodman and Bernhard Salow. Taking a chance on KK. *Philosophical Studies*, 175:183–196, 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, 2023.

John Hawthorne. Normality and KK. ms.

Frank Hong. *Uncertain Knowledge*. PhD thesis, University of Southern California, 2023.

Bertrand Russell. *Human Knowledge: Its Scope and its Limits*. Allen & Unwin, 1949.

Robert Stalnaker. Contextualism and the logic of knowledge. In *Knowledge and Conditionals*. Oxford UP, 2019.

Timothy Williamson. *Knowledge and its Limits*. Oxford UP, 2000.

Timothy Williamson. Improbable knowing. In Trent Dougherty, editor, *Evidentialism and its Discontents*. Oxford UP, Oxford, 2011.

Timothy Williamson. Response to Cohen, Comesaña, Goodman, Nagel, and Weatherson on Gettier cases in epistemic logic. *Inquiry*, 56:77–96, 2013.

Timothy Williamson. Very improbable knowing. *Erkenntnis*, 79:971–999, 2014.