

Epistemology Normalized*

Jeremy Goodman and Bernhard Salow

June 4, 2022

Curious how much he weighs, Bjorn steps on an inexpensive digital scale. He discovers that it reads 175.4 pounds, but this isn't all that he learns. He also learns something about how much he weighs.

This knowledge has limits. Even if Bjorn weighs exactly 175.4 pounds, he doesn't learn anything so precise. The scale isn't reliable enough for that. Bjorn learns only that his weight is in a certain interval around 175.4 pounds. The size of this interval partly depends on Bjorn's weight, since Bjorn cannot learn what isn't true. While the details of this dependence of knowledge on weight are controversial, some systematic patterns are clear. For example, if the scale overestimates his weight, then anything Bjorn learns is something he would still have learned if the scale had overestimated his weight by less.¹

What explains these contours of Bjorn's knowledge? Here is our account.² Not all situations in which the scale reads 175.4 are equally normal. Some are less normal than others, when they involve a greater disparity between Bjorn's weight and the scale's reading. If such a situation is sufficiently less normal than the one Bjorn is actually in, then he knows that he is not in that situation. And if such a situation either is not sufficiently less normal than any other, or is at least as normal as the situation he is actually in, then for all he knows he is in that situation. As we will show below, this account offers a simple explanation of Bjorn's knowledge and its limits.

Bjorn's case illustrates a much more general phenomenon. In theorizing about what people know, it is often productive to factor their knowledge into two components: *evidential* knowledge of some starting points, such as instrument readings and other background information, and *inductive* knowledge that goes

*Versions of this material were presented at MIT, UC Irvine, Ruhr University Bochum, Cardiff University, Ural Federal University, and the Universities of Barcelona, Helsinki, and Oxford. We'd like to thank Sam Carter, Kevin Dorst, Julien Dutant, Rachel Fraser, Dmitri Gallow, Simon Goldstein, Ben Holguín, Wes Holliday, Josh Knobe, Jennifer Nagel, Josh Pearson, Jim Pryor, and Tim Williamson for helpful discussion of these ideas, and Cian Dorr, Daniel Kotsi, Maria Lasonen-Aarnio, and the editors and reviewers for this journal for comments on earlier drafts. We are especially grateful to Harvey Lederman and one extraordinary reviewer, who each gave extensive comments that improved the paper immensely, and to John Hawthorne, whose influence on this project has been pervasive.

¹Williamson (2011, 2013a, 2014, 2021), Cohen and Comesaña (2013), Goodman (2013), Weatherson (2013), Stalnaker (2015), Carter (2019), Dutant and Rosenkranz (2020), Carter and Goldstein (2021) and Goldstein (2022) develop formal models of cases like this. All of them agree on (the relevant analogues of) the above claims about Bjorn's knowledge.

²Related ideas about knowledge and normality are discussed in Stalnaker (2006, 2015, 2019), Goodman (2013), Greco (2014), Dutant (2016), Goodman and Salow (2018), Carter (2019), Littlejohn and Dutant (2020), Beddor and Pavese (2020), Carter and Goldstein (2021), Loets (2022), and Goldstein and Hawthorne (2022, forthcoming). Smith (2010, 2016, 2017, 2018a,b) and Carter and Hawthorne (ms) develop related ideas about epistemic justification, Hawthorne (2021) about epistemic uses of 'ought', and Carter (2022) about assertability. Connections between normality and belief have also been influential in the literature on non-monotonic reasoning; see Kraus et al. (1990) and Makinson (1993, §3).

beyond one’s evidence. This conception of inductive knowledge is very broad: for example, we do not assume that such knowledge always involves projecting observed regularities to unobserved cases. It is also non-committal on the psychological mechanisms underpinning such knowledge: for example, we do not assume that all inductive knowledge is inferred from one’s evidence.

This paper offers a general framework for theorizing about inductive knowledge in terms of the comparative normality of situations compatible with one’s evidence. While there is an immense literature on induction, much of it deploying sophisticated formal methods, very little of this work is concerned with *knowledge*, as opposed to belief or probability. No doubt this is in part because the inherent fallibility of induction has led some to question whether inductive knowledge is even possible. For example, when articulating the “old” (traditional) problem of induction, Goodman (1955, p. 61) disparagingly writes that “obviously the genuine problem cannot be one of attaining unattainable knowledge or of accounting for knowledge that we do not in fact have.” By contrast, we think that inductive knowledge is not merely attainable but commonplace. As elsewhere in epistemology, a better methodology – and the one we will be adopting here – is to bracket sweeping skeptical impulses and see whether a stable picture of human knowledge can emerge.

Section 1 presents the normality framework and shows how it accounts for the basic contours of Bjorn’s knowledge. Section 2 clarifies the framework and draws some comparisons to other approaches. Sections 3 and 4 are the core of the paper: the first shows how different models of cases like Bjorn’s found in the literature correspond to different combinations of assumptions about the structure of comparative normality and about how comparative normality determines what one knows; the second explores these options more systematically and considers what they imply about the connection between knowledge and belief. Section 5 outlines two competing treatments of situations that are abnormal in some respects but not in others, one of which analyzes normality in terms of probability. Section 6 shows how the framework can be extended to model the dynamics of knowledge and belief in response to new evidence. Section 7 applies these models to inductive knowledge of the future and of laws of nature. Section 8 argues for some surprising knowledge dynamics, which help to further illustrate the significance of the distinction between evidential and inductive knowledge. Section 9 concludes.

1 The Normality Framework

This section lays out the basic framework linking knowledge, evidence, and comparative normality. This framework will form the common core of three more specific proposals we will consider in sections 3 and 4.

The framework makes the familiar idealizing assumption that a person’s knowledge can be characterized by the set of situations that, for all they know, they are in – the set of situations *epistemically accessible* for them – and that a person’s evidence is characterized by the set of situations that, for all their

evidence implies, they are in – the set of situations that are *evidentially accessible* for them. That is, we assume that a person knows a proposition if and only if that proposition is true in all situations that are epistemically accessible for them.³ We write $R_K(w)$ for the set of situations epistemically accessible from w and $R_E(w)$ for the set of situations evidentially accessible from w .

Building on Goodman and Salow (2018, 2021), we appeal to two relations of comparative normality: that of one situation being *at least as normal as* another (notated \succeq), and that of one situation being *sufficiently more normal than* another (notated \gg). We make the following structural assumptions about these relations: (i) being at least as normal is reflexive and transitive, (ii) being sufficiently more normal is irreflexive and implies being at least as normal, and (iii) these relations can be chained together, so that, if w_1 is at least as normal as w_2 , w_2 is sufficiently more normal than w_3 , and w_3 is at least as normal as w_4 , then w_1 is sufficiently more normal than w_4 .⁴ Against this backdrop, let the *normality framework* be the conjunction of the following five claims:

FACTIVITY

w is epistemically accessible from w .

$$w \in R_K(w)$$

EVIDENTIAL KNOWLEDGE

If v is not evidentially accessible from w , then v is not epistemically accessible from w .

$$w \notin R_E(w) \Rightarrow w \notin R_K(w)$$

INDUCTIVE KNOWLEDGE

If w is sufficiently more normal than v , then v is not epistemically accessible from w .

$$w \gg v \Rightarrow v \notin R_K(w)$$

INDUCTIVE LIMITS

If v is evidentially accessible from w and no u evidentially accessible from w is sufficiently more normal than v , then v is epistemically accessible from w .

$$(v \in R_E(w) \wedge \forall u \in R_E(w)(u \not\gg v)) \Rightarrow v \in R_K(w)$$

CONVEXITY

If v is evidentially accessible from w and is at least as normal as some u epistemically accessible from w , then v is epistemically accessible from w .

$$(v \in R_E(w) \wedge \exists u \in R_K(w)(v \succeq u)) \Rightarrow v \in R_K(w)$$

³Cf. Hintikka (1962). Situations, in our sense, settle all eternal truths about the course of history, what time it is, and what evidence the relevant agent has; so they are more like the “centered worlds” of Lewis (1979) and “cases” of Williamson (2000) than like “possible worlds” as traditionally conceived (which only take a stand on eternal propositions) or “situations” in the literature on situation semantics (which take a stand on even less; cf. Kratzer (2019)).

⁴(i)–(iii) imply that being sufficiently more normal is also asymmetric and transitive; unlike Goodman and Salow (2018, 2021) we do not assume that it is well-founded (see note 20).

FACTIVITY says that everything you know is true. EVIDENTIAL KNOWLEDGE says that your knowledge includes your evidence. INDUCTIVE KNOWLEDGE says that any possibility sufficiently less normal than the one you are in is one that you know you are not in; it thus says that certain facts about comparative normality are sufficient for your knowledge to go beyond your evidence. By contrast, INDUCTIVE LIMITS says that certain patterns of comparative normality are necessary for your knowledge to go beyond your evidence: in order to know that an evidential possibility does not obtain, some other evidential possibility must be sufficiently more normal than it. Finally, CONVEXITY says that, to the extent that knowledge can go beyond your evidence, its contours must follow those of comparative normality: every evidential possibility at least as normal as some epistemic possibility must itself be an epistemic possibility.

To see the framework in action, let us apply it to Bjorn. We assume that his evidence implies that the scale reads 175.4 but doesn't imply anything about how much he weighs; we also assume that his evidence is the same in all situations compatible with his evidence. We represent these situations by real numbers, each corresponding to a situation in which Bjorn weighs that many pounds and the scale reads 175.4; for simplicity, we ignore all non-weight differences between situations. Supposing that Bjorn actually weighs α pounds, we make three assumptions about these situations' comparative normality:

1. A situation in which the scale overestimates Bjorn's weight by a given amount is at least as normal as one in which the scale overestimates his weight by a greater amount (and likewise for underestimation).

$$(x < y \leq 175.4 \text{ or } 175.4 \leq y < x) \Rightarrow y \succeq x$$

2. The actual situation is sufficiently more normal than some situations in which the scale overestimates Bjorn's weight (and likewise for underestimation).

$$\exists x(x < 175.4 \text{ and } \alpha \succ x) \text{ and } \exists x(175.4 < x \text{ and } \alpha \succ x)$$

3. The situation in which the scale's reading is perfectly accurate is not sufficiently more normal than every situation in which it overestimates Bjorn's weight (and likewise for underestimation).

$$\exists x(x < 175.4 \text{ and } 175.4 \not\succeq x) \text{ and } \exists x(175.4 < x \text{ and } 175.4 \not\succeq x)$$

Combined with these assumptions, the normality framework implies most of the advertised features of Bjorn's knowledge of his weight. By (2), α is sufficiently more normal than some situations in which the scale overestimates his weight; by (1), and the chaining principle (iii), α is also sufficiently more normal than any situation in which the scale overestimates his weight by more. So, by INDUCTIVE KNOWLEDGE, Bjorn's knowledge puts a lower bound on his weight. Similar reasoning shows that his knowledge also puts an upper bound on his weight. CONVEXITY – together with (1) – ensures that his knowledge cannot go beyond the placing of such upper and lower bounds; so the weights compatible with his knowledge form an interval. FACTIVITY ensures that this interval always

contains α . (1) ensures that the situation in which Bjorn weighs 175.4 pounds is at least as normal as any other; (3) thus ensures that there are some situations involving over- and under-estimation that are not sufficiently less normal than any other. So INDUCTIVE LIMITS ensures that the interval representing Bjorn’s knowledge also contains 175.4, and some values on either side of 175.4.

These predictions are notable because the normality framework and (1)-(3) leave open many other questions about the structure of Bjorn’s knowledge, questions about which there is disagreement in the literature. This is both because (1)-(3) don’t settle all questions about the comparative normality of the situations in question, and because the normality framework doesn’t settle all questions about what a person knows given their evidence and facts about situations’ comparative normality. In section 3, we will look at how both the case-specific assumptions (1)-(3) and the normality framework more generally can be elaborated to deliver more detailed epistemological predictions, including our earlier claim that decreasing the extremity of the scale’s error should not decrease what Bjorn knows. But before delving into those details, let us consider some broader features of the normality framework, and how they compare to other accounts of inductive knowledge that may be more familiar.

2 Situating the Framework

One useful reference point for the present project is Lewis’s (1996) well-known theory of knowledge. Like the normality framework, Lewis employs a notion of evidential accessibility (he refers to this as one situation being “uneliminated” by the evidence at another) and endorses EVIDENTIAL KNOWLEDGE. His various rules for what one can “properly ignore” play roles analogous to the normality framework’s other principles. The “rule of actuality” (that one can never properly ignore the subject’s actual situation) is equivalent to FACTIVITY. The rules of “attention” and “belief” play a role similar to INDUCTIVE LIMITS: they set limits on inductive knowledge, by guaranteeing that certain situations (those that we are attending to, and those that the subject lends substantial credence to) are epistemically accessible whenever they are evidentially accessible. The rules of “method”, “reliability”, and “conservatism” play a role analogous to INDUCTIVE KNOWLEDGE: they ensure the possibility of inductive knowledge, by maintaining that certain situations (those in which perception, testimony, memory, statistical reasoning, inference to the best explanation, or other ordinary belief-forming processes go wrong) are not epistemically accessible (unless one of the other rules requires them to be). And the “rule of resemblance” plays a similar role to CONVEXITY, telling us that certain situations’ being epistemically accessible guarantees that others are too.

Another helpful reference point is a theory that, inspired by theories of knowledge as belief that is safe from error, maintains that we know whatever our evidence safely indicates. We can make this idea precise by appealing to the notion of two situations’ being *sufficiently close* to one another, and saying that v is epistemically accessible from w if and only if it is both evidentially

accessible from w and the two situations are sufficiently close.⁵

Like the normality framework, both of these accounts approach their topic at a high level of abstraction. For example, they don't mention that in order to know something you need to believe it. Consequently, they also don't capture the fact that you fail to know when you believe only for bad reasons. And since they characterize one's knowledge in terms of what is true in all epistemically accessible situations, they have the implausible consequence that knowing some propositions implies knowing every proposition that is entailed by them. While a full theory of inductive knowledge would address these important issues, there is still much that can be said while abstracting away from them. Moreover, the main extant strategies that safety-based theorists have used to address these problems can be co-opted within the normality framework fairly straightforwardly.⁶ So we will set these issues aside in what follows.

Even at this high level of abstraction and idealization, there are important differences between the normality framework and its competitors. Perhaps the most salient difference from Lewis's account is that it contains no analogue of his "rule of attention". But another, arguably more important, difference is that where Lewis's rules feature a large grab bag of different notions, the principles of the normality framework feature only the relations of comparative normality and evidential accessibility. This makes the normality framework significantly more parsimonious, and hence more amenable to systematic investigation.⁷

By contrast, the simple safety theory sketched above is at least as parsimonious and tractable as the normality framework. Both theories are also naturally developed using a similarly non-reductive methodology. Williamson (2009, pp. 9-10), for example, explicitly rejects the possibility of giving an account of

⁵Sosa (1999) and Williamson (2000) argue that safe belief is necessary for knowledge; Williamson (2009) treats it as sufficient as well. The theory discussed here departs from theirs by focusing not on whether a belief is safe, but rather on what some evidence safely indicates. Bacon (2014, 2020) uses a theory roughly like the one sketched here to theorize about the kinds of cases we discuss in section 7; assuming that "sufficiently close" and "saliently resembles" are interchangeable, it is also equivalent to a simplification of Lewis's (1996) theory that drops all rules except those of "actuality" and "resemblance". Smith (2016, chapter 7) discusses the notion of what some evidence safely indicates in the context of a theory of justification.

⁶Central to these strategies is the reification of belief-states. Adapting the strategy of Hawthorne and Lasonen-Aarnio (2009), we could treat normality-based theories of epistemic accessibility as instead accounts of those situations in which a belief-state must have a true content in order to amount to knowledge. (Their idea is that belief-states can have their contents contingently, so that b might have a necessarily true content without amounting to knowledge because it could have had a different content that was false – for example, a true belief expressed by "He's Bill" won't amount to knowledge if it has a false content in any evidentially accessible situation at least as normal as actuality.) Alternatively, we could adapt the strategy of Williamson (2009) and treat the relation of normality relations and 'evidential' accessibility as pairs of a situation and a belief-state, and reinterpret normality-based theories of epistemic accessibility as theories of which belief-states must have true contents in which situations in order for a given belief-state to amount to knowledge in a given situation. (For example, INDUCTIVE KNOWLEDGE could be replaced with the claim that b amounts to knowledge in w whenever, for all $\langle b', v \rangle$, if (i) b' is formed at v on the same evidential basis as b is in w and (ii) $\langle b', v \rangle$ is not sufficiently less normal than $\langle b, w \rangle$, then b' has a true content in v .)

⁷This is not to suggest that the Lewisian framework is completely resistant to formal investigation; see Holliday (2015) and Salow (2016).

“closeness” which does not itself presuppose epistemic notions such as knowledge. Rather, following Lewis’s (1973) attitude to the notion of comparative similarity central to his theory of counterfactuals, Williamson maintains that the relevant notion is to be understood in terms of its role in his theory. Since one could adopt a similar approach to comparative normality, one might wonder whether the views are indeed genuine competitors.

Fortunately, substantive accounts of closeness and comparative normality are not needed to distinguish the simple safety theory from the normality framework. We have seen how the normality framework makes natural predictions about Bjorn given fairly minimal assumptions about comparative normality. The simple safety theory is less promising here. For whatever it takes for two situations to be sufficiently close, being sufficiently close is a symmetric relation. So provided the relevant situations in which the scale reads 175.4 pounds are all evidentially accessible from each other (as we have been assuming), the simple safety theory predicts that epistemic accessibility is a symmetric relation on those situation. And this prediction is undesirable: if in v Bjorn does in fact weigh approximately 175.4 pounds while in w he weighs much less, we want v to be epistemically accessible from w (since we may suppose that, even if the scale overestimated his weight significantly, Bjorn wouldn’t be able to tell this from the reading) but not *vice versa* (since Bjorn should learn a fair bit about his weight when the scale is working well).⁸

Of course, proponents of a safety-based epistemology might reasonably complain that the simple safety theory is oversimple, and go on to articulate some non-symmetric notion of one possibility being close to another. Still, it is striking that the normality framework has a clearly non-symmetric relation at its core, while the language of “closeness” that one finds in discussions of safety strongly suggests a symmetric relation. For this reason, we find the normality framework to be, *prima facie*, a more promising account of inductive knowledge.

We mentioned above that one might theorize about comparative normality in a non-reductive spirit, maintaining that the notion is to be understood primarily in terms of its role in a theory of inductive knowledge. We will mostly adopt this approach (although we will consider a possible analysis of comparative normality in terms of comparative probability in section 5). Importantly, our use of the word “normal” to describe the operative relations between situations is largely in deference to the literature. While we think the connotations of “normal” are often helpful, they can sometimes be misleading or distracting. Some may find it more intuitive to think of \succeq and \succ as relations of comparative *plausibility*; readers with this sensibility are welcome to mentally substitute “plausible” for “normal” throughout.⁹

⁸Compare Magidor’s (2018) argument that, if knowledge is belief that is safe from error, then safely envatted brains-in-vats can know that they are envatted (although, unlike us, she does not treat this as a *reductio*), and Bacon’s (2014) view that, if a fair coin is going to land heads 1000 times in a row, then you’re in a position to know now that it won’t land heads approximately half of the time.

⁹Compare the literatures on non-monotonic reasoning and on belief revision: the former tends to be framed in terms of comparative “normality” and the latter in terms of comparative

A notable structural feature of the normality framework is that both \succeq and \succ are binary relations between situations. Drawing on Carter (2019) and Williamson (2021), one might object that this feature of the framework ignores the fact that what is normal is clearly a contingent matter: if what is normal is different in different situations, shouldn't comparative normality be relativized to such situations?¹⁰ A flat-footed response to this objection is that it appeals to a pretheoretical notion of what is “normal”, which is of dubious relevance to the current project.¹¹ However, we think a more satisfying response is available. Carter and Williamson are theorizing about what is normally the case – a property of propositions. This may well be a contingent matter, depending, for example, on contingent facts about what is and isn't typical. But our discussion is conducted in terms of relations of comparative normality between situations. Since situations take a stand on what is and isn't typical, we can respect the relevance of typicality to normality without postulating contingency in situations' comparative normality. For example, while it may be contingent whether it is normal for Anna to be up after midnight, due in part to contingency in her routines, it does not follow that it is contingent whether a situation in which Anna goes to bed late as usual is more normal than one in which Anna goes to bed early despite typically going to bed late.

Given the relevance of such background conditions to situations' comparative normality, we should revisit an assumption we made earlier about Bjorn: that we may harmlessly pretend that there is only one evidential possibility in which he has any given weight. A more realistic model would have different possibilities that agree on Bjorn's weight and on the scale's reading but differ about the scale's reliability, since his evidence needn't settle exactly how reliable his scale is. And since large errors are less abnormal on less reliable scales, this raises the worry that some situations with large errors will be fairly normal, and hence epistemically accessible, in which case Bjorn will learn very little about his weight; cf. Loets (2022, p. 177). In reply: people who lack evidential knowledge of their scales' reliability typically still have inductive knowledge that their scales are not too unreliable. This knowledge could be based on testimonial evidence,

“plausibility”, but this is seen by many in these fields as a merely terminological difference.

¹⁰Formally, we could introduce relativization to a third situation that settles the standards of normality by (a) substituting “normal_w” for “normal” in INDUCTIVE KNOWLEDGE, INDUCTIVE LIMITS, and CONVEXITY, and (b) requiring, for all w , that the structural constraints (i)-(iii) from section 1 hold for being at least as/sufficiently more normal_w. In Goodman and Salow (2021, appendix A), we show that the probabilistic account of comparative normality discussed in section 5 requires such relativization if TRANSPARENCY (see below) fails. However, allowing for widespread situation-relativity even in the presence of TRANSPARENCY, as in Holliday's (2015) formalization of relevant-alternatives theories of knowledge, blurs the line between normality-based views and safety-based ones.

¹¹Loets (2022) argues that English constructions using the word “normal” are a poor guide to epistemologically relevant notions of normality. Relatedly, a number of recent theories of knowledge and justified belief formulated in terms of what is “normal” are substantively quite different from ours; cf. Goldman (1986), Lepin (2007), Graham (2012, 2017), Ball (2013), Peet and Pitcovski (2018), and Horvath and Nado (2021), who, unlike us and the authors cited in note 2, do not use normality to characterize those situations in which a belief must be true in order to amount to knowledge or be justified.

or obtained oneself (say, by making repeated measurements of standardized weights to learn how reliable the scale is). Either way, the evidentially accessible situations in which the scale is highly unreliable will be sufficiently less normal than actuality; so, by INDUCTIVE KNOWLEDGE, they won't threaten Bjorn's knowledge of his weight. More realistic models will thus make broadly similar predictions to those of the simple models we will be exploring here.

More generally, the anti-skeptical predictions of the normality framework are fairly robust with respect to different ways of drawing the distinction between evidential and inductive knowledge. For example, although we don't share such a conception of perceptual knowledge, those who think of perceptual knowledge as inductive, and of one's perceptual evidence as confined to facts about one's experiences, can use the normality framework to explain how Bjorn can have inductive knowledge that his scale reads 175.4. Whether this knowledge is evidential or inductive makes little difference to whether Bjorn can gain further inductive knowledge about how much he weighs.

In what follows we will make the following assumption about evidence:

TRANSPARENCY: Your evidence entails what your evidence is.

$$v \in R_E(w) \Rightarrow R_E(v) = R_E(w)$$

This is a natural idealization in a case like Bjorn's, where we can think of one situation as evidentially accessible from another if and only if the scale gives the same reading in both situations. Lewis (1996) and Stalnaker (2015, 2019) endorse TRANSPARENCY quite generally, deriving it from a substantive account of evidential accessibility as sameness of certain aspects of the relevant person's internal state. We are not endorsing any such account. In particular, we want to leave open that, in ordinary circumstances, our evidence can include facts about our immediate surroundings, such as a scale reading, which conflicts with TRANSPARENCY given FACTIVITY, EVIDENTIAL KNOWLEDGE, and the possibility of misperception; see Williamson (2000, chapter 8). But even if TRANSPARENCY is false, it remains a reasonable idealization when thinking about the structure of our purely inductive knowledge. People are not in fact infallible or omniscient about the readings of their scales, but pretending that they are does little harm when our interest is in how observing those readings allows them to know how much they weigh. How the normality framework interacts with failures of TRANSPARENCY is an important topic for future work.

3 Characterizing Epistemic Accessibility

While the normality framework places strong constraints on how knowledge is related to evidence and normality, it falls short of an explicit characterization of epistemic accessibility. This section and the next present three such characterizations. Each gives necessary and sufficient conditions on epistemic accessibility in terms of evidential accessibility and comparative normality. And each, when combined with the normality framework's structural constraints on comparative normality, implies the framework's five principles about epistemic accessibility.

The present section uses the Bjorn example to illustrate these three options; it also shows how different proposals from the literature discussing similar examples can be situated within the normality framework.

This is not the only possible approach to giving necessary and sufficient conditions for epistemic accessibility in terms of evidential accessibility and comparative normality. An alternative strategy would be to impose additional structural constraints on comparative normality that enable the five principles of the normality framework to fully pin down epistemic accessibility. In particular, consider the following two principles (where one situation is *more normal than* another if it is at least as normal as the other but not *vice versa*, and two situations are *comparable* if one is at least as normal as the other):

COLLAPSE

If v is more normal than w , then v is sufficiently more normal than w .

$$(v \succeq w \wedge w \not\succeq v) \Rightarrow v \gg w$$

COMPARABILITY

If v is evidentially accessible from w , then w and v are comparable.

$$v \in R_E(w) \Rightarrow (v \succeq w \vee w \succeq v)$$

Given the normality framework, COLLAPSE and COMPARABILITY imply that $R_K(w) = \{v \in R_E(w) : v \succeq w\}$ – the epistemically accessible situations are the evidentially accessible situations that are at least as normal as actuality.¹²

But while adding COLLAPSE and COMPARABILITY to the normality framework yields a simple characterization of epistemic accessibility and an improvement in overall parsimony (by allowing us to operate with a single notion of comparative normality), we think there are strong reasons to reject it. This is because we think there are strong abductive reasons to reject COLLAPSE, as well as powerful, though less decisive, reasons to question COMPARABILITY. While the strongest such reasons will emerge from examples considered in later sections, the basic ideas can be illustrated already in the case of Bjorn.

Consider four competing hypotheses, depicted in figure 1, about the comparative normality of Bjorn’s evidential possibilities. All four hypotheses respect COMPARABILITY, and the first three respect COLLAPSE. To motivate rejecting COLLAPSE, we will argue that (a) makes problematic epistemological predictions, and that (d) is a more natural way of thinking about relations of comparative normality than either (b) or (c). (We will return to COMPARABILITY later.)

The normality relations depicted in (a) are an instance of a more general approach, according to which all normality relations are reducible to a normal/abnormal dichotomy. On this picture, all normal situations are equally normal, all abnormal situations are equally normal, and any normal situation is sufficiently more normal than any abnormal one. The normality framework

¹²Stalnaker (2019) accepts both COLLAPSE and this characterization of R_K ; he rejects COMPARABILITY, so this characterization isn’t forced by the normality framework, but he also holds that there is always some evidential possibility that is at least as normal as every other, which combined with COLLAPSE, the reflexivity and transitivity of being at least as normal, and the characterization of R_K entails the rest of the normality framework.

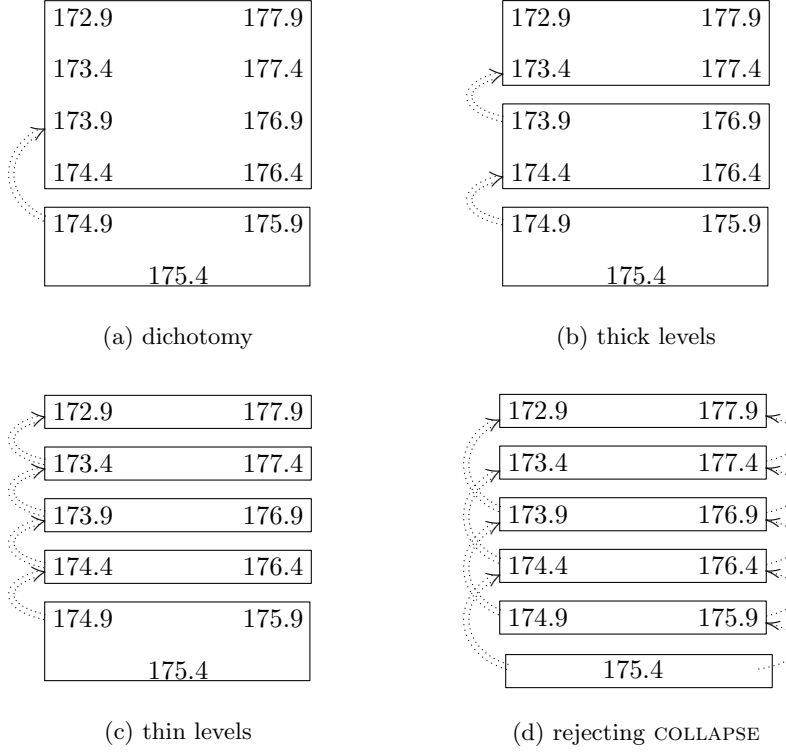


Figure 1: Bjorn and COLLAPSE

Pictorial conventions: Every depicted situation is evidentially accessible from every other. A thin(/thick) arrow from w to v means that w is at least as normal as(/sufficiently more normal than) v ; situations in the same box are equally normal; an arrow from one box to another means that the relevant relation holds between every situation in the first box and every situation in the second box. Normality relations implied by the structural conditions (i)-(iii) from section 1 are not depicted.

then implies that, in abnormal situations, we know only what is entailed by our evidence, while in normal situations we know what is entailed by our evidence together with the fact that conditions are normal. Greco (2014) proposes a version of this view. It can also be seen as a descendant of the traditional reading of Hume, according to which the problem of justifying induction reduces to the problem of justifying belief in a single proposition, namely the “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*” (1739, 1.3.6.4, emphasis original).

Inductive knowledge, however, admits finer gradations than this proposal allows, as pointed out by Carter (2019) and Bird and Pettigrew (2021). There can be three situations in each of which you have the same evidence but in no two of which do you have the same knowledge. For example, the reading on Bjorn’s scale could have a middling error, so that he knows less than he would had that reading been perfectly accurate, but more than he would had the scale erred even more extremely.

The normality relations depicted in (b) and (c) represent the two most natural ways of responding to this problem while preserving COLLAPSE.¹³ Both agree with (a) about what Bjorn knows in the most normal situations, where the scale’s error is small, but unlike (a) they allow him to have non-trivial knowledge when the scale’s errors are bigger. However, neither option is very natural. The basic issue is that, given COLLAPSE and COMPARABILITY, the most normal situations must correspond to a non-trivial interval of weights, since even in the most favorable situations Bjorn’s knowledge is characterized by a non-trivial interval. All of these situations are equally normal despite the magnitude of the scale’s error differing between them. Since not all evidentially accessible situations are equally normal, this means that some but not all increases in the scale’s error make for decreases in normality. This is sufficiently unnatural to motivate exploring alternatives.

One such alternative, in which COLLAPSE fails, is depicted in (d): x is at least as normal as y just in case x is at least as close to 175.4 as y is, and that x is sufficiently more normal than y just in case x is at least c pounds closer to 175.4 than y is.¹⁴ (For concreteness, (d) depicts $c = 1$.) The uniquely most normal situation on this model has Bjorn’s weight matching the displayed weight of 175.4 pounds; the interval of most normal situations that featured in the COLLAPSE-respecting models (a)-(c) are now, in (d), the situations that are not sufficiently less normal than the most normal one – that is, those in which his weight lies in the (open) interval $175.4 \pm c$.

Given such failures of COLLAPSE, the characterization of epistemic accessibility as $R_K(w) = \{v \in R_E : v \succeq w\}$ no longer follows from the principles of the normality framework; in fact, it now conflicts with INDUCTIVE LIMITS in all situations where the magnitude of the scale’s error is less than c . The minimal change required to respect INDUCTIVE LIMITS is to include among the epistemically accessible situations not only those evidential possibilities which are at least as normal as the actual situation, but also those which are not sufficiently less normal than any other evidential possibilities. This leads to the following characterization of epistemic accessibility, from Goodman and Salow (2018):

¹³Stalnaker (2015) defends what is in effect a version of option (b); Goldstein (2022) defends a more complicated view that embraces similar discrete jumps in normality. Cohen and Comesaña (2013) and Goodman (2013) defend something akin to option (c), postulating a non-trivial interval of maximal normality that decreases continuously outside of that interval (although the former reject COMPARABILITY and the latter rejects COLLAPSE); Williamson (2013b) criticizes them on this basis.

¹⁴Our preferred model of the case, described in section 6, differs slightly from this one, but in ways that are inessential to the present discussion.

K-NORMALITY: v is epistemically accessible from w if and only if v is evidentially accessible from w and either (i) no u evidentially accessible from w is sufficiently more normal than v or (ii) v is at least as normal as w .

$$R_K(w) = \{v \in R_E(w) : \forall u \in R_E(w)(u \not\geq v)\} \cup \{v \in R_E : v \succeq w\}$$

K-NORMALITY makes the same predictions about knowledge given the COLLAPSE-violating hypothesis (d) as the normality framework alone makes given the COLLAPSE-respecting hypothesis (c). However, while the normality framework is *consistent* with this pattern of epistemic accessibility given (d), the failure of COLLAPSE means that the framework does not entail it. For the framework is also consistent with the following simpler characterization of epistemic accessibility, from Goldstein and Hawthorne (2022, p.1346):

I-NORMALITY: v is epistemically accessible from w if and only if v is evidentially accessible from w and w is not sufficiently more normal than v .

$$R_K(w) = \{v \in R_E : w \not\geq v\}$$

I-NORMALITY and K-NORMALITY make different predictions given hypothesis (d). Following Williamson (2013b), say that Bjorn has *cliff-edge knowledge* if he weighs x pounds and either knows that he weighs at least x pounds or knows that he weighs at most x pounds. I-NORMALITY implies that Bjorn never has cliff-edge knowledge, while K-NORMALITY implies that he has cliff-edge knowledge in every situation where the scale's error is at least c . More generally, I-NORMALITY makes predictions about Bjorn's knowledge similar to Williamson's (2011, 2013a, 2014) and Goodman's (2013) treatments of similar example. This contrast between K-NORMALITY and I-NORMALITY is depicted in the right side of figure 2, which shows the epistemological predictions of both proposals given hypothesis (d) about the relevant situations' comparative normality.

It is controversial whether there is anything wrong with cliff-edge knowledge. While the normality framework does not settle this debate, it does offer a new perspective on what is at stake. For given COLLAPSE-violating hypotheses such as (d), the question of whether to allow cliff-edge knowledge is equivalent to a purely normality-theoretic question: must evidentially accessible situations that are less normal but not sufficiently less normal than actuality be epistemically accessible? An affirmative answer to this question corresponds to the following principle (where one situation is *insufficiently less normal than* another if it is less normal than it but not sufficiently less normal than it):

MARGINS: If v is evidentially accessible from w , and insufficiently less normal than w , then v is epistemically accessible from w .

$$(v \in R_E(w) \wedge w \succeq v \wedge v \not\geq w \wedge w \not\geq v) \Rightarrow v \in R_K(w)$$

This means that a ban on cliff-edge knowledge can be motivated independently of the idea that knowledge requires belief that is safe from error (which is how Williamson motivates it). For it can be derived from MARGINS given the relevant failures of COLLAPSE. And MARGINS can be independently motivated as a consequence of I-NORMALITY, an especially simple characterization of epistemic




Figure 2: Bjorn and MARGINS

Pictorial conventions: The right diagram depicts epistemic accessibility given the pattern of comparative normality depicted on the left. It only depicts what is epistemically accessible from the situations in bold. Triple arrows indicate what is accessible under all three of K-, KM-, and I-NORMALITY; double arrows indicate what is accessible under both KM- and I-NORMALITY; and single arrows indicate what is accessible under I-NORMALITY only.

accessibility. A ban on cliff-edge knowledge thus cannot be undermined merely by objecting to a safety condition on knowledge.¹⁵

At the same time, those committed to the possibility of cliff-edge knowledge can take heart in the fact that there is a principled way to combine that possibility with one's knowledge being inevitably inexact. Williamson (2013b) criticizes this combination on the grounds that it requires epistemic accessibility to have the suspicious formal features that comparative normality has according to hypotheses (b) and (c) above: while some errors in the scale are epistemically significant, all small enough errors are epistemically inert. But we have just seen how K-NORMALITY provides a principled basis for that prediction, by showing how it might arise from a natural hypothesis (d) about the structure of comparative normality.

The reason cliff-edge knowledge has been so widely discussed is that, as Williamson (2000) shows, the possibility of such knowledge is a consequence of the 'KK' principle, that knowing a proposition implies knowing that you know it. On this basis Williamson and others have attacked KK, by arguing against the possibility of cliff-edge knowledge in cases where KK predicts it. On the other side of the debate, a number of authors have appealed to versions of K-NORMALITY to defend the possibility of cliff-edge knowledge.¹⁶ This is natural,

¹⁵Neta and Rohrbaugh (2004) respond to Williamsonian qualms about cliff-edge knowledge by raising potential counterexamples to safe belief being necessary for knowledge. Beddor and Pavese (2020) argue that these counterexamples – and others from Kelp (2009) – involve situations in which things could easily have been substantially less normal than they actually are, which would prevent them telling against MARGINS.

¹⁶See Greco (2014), Stalnaker (2015, 2019), and Goodman and Salow (2018); for critical




Figure 3: Bjorn without COMPARABILITY

since, given TRANSPARENCY, K-NORMALITY entails KK. Whether one sees this as an attraction or a cost of K-NORMALITY depends on what one thinks about the independent plausibility of KK and about the status of TRANSPARENCY, issues which are beyond the scope of this paper.¹⁷

Having explained how COLLAPSE fails in the case of Bjorn, and highlighted an important choice that arises as a result, we turn now to COMPARABILITY. One might hold that, in situations where the scale significantly underestimates his weight, Bjorn can still know that it does not significantly overestimate it, and likewise that, in situations where the scale significantly overestimates his weight, Bjorn can still know that it does not significantly underestimate it. Cohen and Comesaña (2013) defend a version of this claim. But it is inconsistent with COMPARABILITY, since these situations are mutually evidentially accessible, and COMPARABILITY implies (given FACTIVITY and CONVEXITY) that mutually evidentially accessible situations cannot be mutually epistemically inaccessible. However, if we maintain that situations involving overestimation are incomparable with situations involving underestimation, as depicted in figure 3, then K-NORMALITY vindicates Cohen and Comesaña’s idea that situations involving significant overestimation and situations involving significant underestimation are mutually epistemically inaccessible.

This prediction does not generalize to I-NORMALITY. In fact, I-NORMALITY has unacceptably skeptical consequences when combined with the incomparability hypothesis, since it predicts that, if the scale even slightly underestimates Bjorn’s weight, then *any* amount of overestimation will be compatible with his knowledge, so his knowledge places no lower bound on his weight. I-NORMALITY

discussion, see Carter (2019), Bird and Pettigrew (2021), Williamson (2021), and Loets (2022).

¹⁷Although nominally a defense of KK, Goodman and Salow (2018) also shows how the kind of cases discussed in section 7 put significant pressure on that principle (building on Dorr et al. (2014)). We think some arguments against KK, such as those of Radford (1966) and Williamson (2021), are best seen as turning on failures of TRANSPARENCY; arguably, this is true even of the argument in Williamson (2000, chapter 4). We discuss the status of KK in the normality framework at greater length in Goodman and Salow (ms a), where we describe cases showing that, even assuming K-NORMALITY, the transitivity of evidential accessibility is neither necessary nor sufficient for the transitivity of epistemic accessibility.

is thus implausible if COMPARABILITY can fail in this way.

Does this mean that MARGINS, and a ban on cliff-edge knowledge, are hostage to COMPARABILITY? Not necessarily. For we can insist that evidentially accessible situations insufficiently less normal than actuality must always be epistemically accessible (so that MARGINS holds), without requiring the same of situations incomparable with actuality. More precisely, consider:¹⁸

KM-NORMALITY: v is epistemically accessible from w if and only if v is evidentially accessible from w and either (i) no u evidentially accessible from w is sufficiently more normal than v , or (ii) v is at least as normal as some u evidentially accessible from w that w is at least as normal as but not sufficiently more normal than.

$$R_K(w) = \{v \in R_E(w) : \forall u \in R_E(w)(u \not\geq v)\} \cup \{v \in R_E(w) : \exists u \in R_E(w)(v \geq u \wedge w \geq u \wedge w \not\geq u)\}$$

As illustrated in figure 3, KM-NORMALITY agrees with K-NORMALITY that situations involving significant overestimation and situations involving significant underestimation are mutually inaccessible, and agrees with I-NORMALITY in ruling out cliff-edge knowledge.

This completes our discussion of how COLLAPSE and COMPARABILITY might fail in the case of Bjorn, and what characterizations of epistemic accessibility look natural if they do. Our investigation of the normality framework will be guided by the hypothesis that one of the three characterizations we have introduced – I-, K- or KM-NORMALITY – is correct. This disjunctive hypothesis already non-trivially strengthens the normality framework. For given our minimal assumptions from section 1 about the comparative normality of Bjorn’s evidential possibilities, each of I-, K- and KM-NORMALITY implies our earlier claim that, if the scale overestimates his weight, then Bjorn knows at least as much as he does in any situation in which it overestimates his weight by more. This is because all three characterizations imply that, if you have the same evidence in w and in v and w is at least as normal as v , then you know at least as much in w as you do in v : if $R_E(w) = R_E(v)$ and $w \geq v$, then $R_K(w) \subseteq R_K(v)$. This principle is not a consequence of the normality framework alone. For it is consistent with the framework that I-NORMALITY characterizes what Bjorn knows if he weighs 174.9 pounds while K-NORMALITY characterizes what he knows if he weighs 174.4 pounds; given the normality relations depicted in figures 2 and 3, Bjorn would then know less in the more normal situation.

We should note that we ultimately do not think that COMPARABILITY fails in Bjorn’s case; this is largely because, as we will see in section 6, a model

¹⁸Goodman and Salow (2018, footnote 10) considers characterizing epistemic accessibility as $R_K(w) = \{v \in R_E(w) : \forall u \in R_E(w)(u \not\geq v)\} \cup \{v \in R_E(w) : v \geq w \vee (w \geq v \wedge w \not\geq u)\}$. This entails MARGINS but fails to entail CONVEXITY, and so makes implausible predictions about the case depicted in figure 4 below: (2,1) wouldn’t be accessible from (1,2) even though (2,2) would be. Carter and Goldstein (2021) offer the even simpler characterization $R_K(w) = \{v \in R_E(w) : v \geq w \vee (w \geq v \wedge w \not\geq u)\}$, which entails MARGINS but neither CONVEXITY nor INDUCTIVE LIMITS. Given the normality relations in figure 3, this makes the clearly wrong prediction that, if the scale overestimates Bjorn’s weight, then he knows that it didn’t underestimate it.

that preserves COMPARABILITY allows for an attractive generalization to slightly more complex variants of the case. We will consider more compelling potential counterexamples to COMPARABILITY in section 5. But before doing so, we will discuss the three characterizations we have just introduced more systematically.

4 The normality landscape

The last section introduced three characterizations of epistemic accessibility:

$$\begin{aligned} \text{I} : R_K(w) &= \{v \in R_E(w) : w \not\asymp v\} \\ \text{K} : R_K(w) &= \{v \in R_E(w) : \forall u \in R_E(w)(u \not\asymp v)\} \cup \\ &\quad \{v \in R_E(w) : v \succeq w\} \\ \text{KM} : R_K(w) &= \{v \in R_E(w) : \forall u \in R_E(w)(u \not\asymp v)\} \cup \\ &\quad \{v \in R_E(w) : \exists u \in R_E(w)(v \succeq u \wedge w \succeq u \wedge w \not\asymp u)\} \end{aligned}$$

Each of these has the feature, advertised at the beginning of the last section, that together with the normality framework’s structural constraints on comparative normality, it implies the normality framework’s five principles about epistemic accessibility. In this section we will explore these options further, showing why they are natural and illustrating some of the trade-offs they exhibit. We also show how the normality framework can be used to model a notion of justified belief, and use this to explore how the framework handles Gettier cases.

We will begin by reformulating these three explicit definitions of epistemic accessibility in a more intuitive way. Given the structural constraints on comparative normality, they are equivalent to the following implicit definitions:

I-NORMALITY: Epistemic accessibility is the most *inclusive* relation compatible with EVIDENTIAL KNOWLEDGE and INDUCTIVE KNOWLEDGE.

K-NORMALITY: Epistemic accessibility is the most *restrictive* relation compatible with FACTIVITY, INDUCTIVE LIMITS, and CONVEXITY.

KM-NORMALITY: Epistemic accessibility is the most *restrictive* relation compatible with MARGINS, INDUCTIVE LIMITS, and CONVEXITY.

These facts underlie our nomenclature. The “I” indicates that, except when one of the framework’s principles mentioned implies otherwise, there is an *ignorance default*: epistemic accessibility relates as *many* situations as possible, and hence people know as little as possible. By contrast, the “K” indicates that, except when one of the principles mentioned implies otherwise, there is a *knowledge default*: epistemic accessibility relates as *few* situations as possible, and hence people get to know as much as possible. The “M” stands for MARGINS.

I-NORMALITY and K-NORMALITY represent opposite extremes: I-NORMALITY has people know as little as possible given the normality framework, while K-NORMALITY has them know as much as possible. And KM-NORMALITY is a natural intermediate position, given the following two facts:

- KM-NORMALITY and K-NORMALITY are equivalent given COLLAPSE.
- KM-NORMALITY and I-NORMALITY are equivalent given COMPARABILITY.

Considered at a purely formal level, I-NORMALITY is in many respects the most attractive of the three options. Its explicit characterization of epistemic accessibility is encouragingly simple, non-disjunctive, and parsimonious, featuring only one relation of comparative normality.¹⁹ By contrast, both K-NORMALITY and KM-NORMALITY have disjunctive characterizations of epistemic accessibility that feature both normality relations. And neither characterization is as simple as I-NORMALITY, with KM-NORMALITY being especially complicated.

However, for reasons that began to emerge in section 3 and that will be amplified in section 5, I-NORMALITY has problematic skeptical consequences if there are counterexamples to COMPARABILITY. (They are ‘skeptical’ in that they imply that we do not have anywhere near as much knowledge as we ordinarily take ourselves to have.) We thus think that I-NORMALITY is tenable only when combined with COMPARABILITY. This makes for interesting dialectical connections between COMPARABILITY and MARGINS. In particular, it lends some abductive support to the biconditional: COMPARABILITY \leftrightarrow MARGINS. For if COMPARABILITY is true, then I-NORMALITY is tenable; provided it is tenable, it arguably displays greater general theoretical virtues than its competitors; and I-NORMALITY entails MARGINS. In the other direction, if COMPARABILITY fails, then the question of whether MARGINS is true becomes the question of which of K-NORMALITY and KM-NORMALITY is true; K-NORMALITY is arguably preferable on grounds of its comparative simplicity; and K-NORMALITY leads to counterexamples to MARGINS given failures of COLLAPSE. But these are only some of the many considerations for and against COMPARABILITY and MARGINS. So while we reject COLLAPSE and reject I-NORMALITY without COMPARABILITY, we think all other consistent combinations are worthy of serious consideration.

Let us now turn to K- and KM-NORMALITY. According to both views, your total knowledge is a disjunction of two propositions: one that depends only on your evidence, and another which is sensitive to where you are situated among your evidential possibilities. The first disjunct is the same for both proposals, and is required by INDUCTIVE LIMITS. The second disjuncts are different, and result from the interplay of CONVEXITY with either FACTIVITY or MARGINS. The remainder of this section explores this structure more systematically. We first show how the normality framework can model belief as well as knowledge, and that INDUCTIVE LIMITS corresponds to the claim that knowledge requires belief. We then explain how CONVEXITY addresses the Gettier problem of explaining why not all justified true beliefs are knowledge. While the normality framework’s sparse resources cannot distinguish what one believes, what one justifiably believes, and what one has justification to believe, running these notions together fits the level of idealization we are operating with already, which abstracts away from the psychological underpinnings of knowledge.

¹⁹I-NORMALITY also implies that knowledge factors into independent evidential and inductive components: $R_K(w) = R_E(w) \cap \{v : w \not\asymp v\}$. In this way it resembles the Lewisian and simple safety theories discussed in section 2; see also Xu and Chen (2018).

Say that a situation is *doxastically accessible* for someone just in case everything they believe is true in that situation. We write $R_B(w)$ for the set of situations doxastically accessible from w . As with knowledge, we make the idealizing assumption that people believe every proposition that is true in every doxastically accessible situation. So a characterization of belief reduces to a characterization of doxastic accessibility.

The natural characterization of doxastic accessibility is as follows:²⁰

BELIEF: v is doxastically accessible from w if and only if v is evidentially accessible from w and no u evidentially accessible from w is sufficiently more normal than v .

$$R_B(w) = \{v \in R_E(w) : \forall u \in R_E(w)(u \not\asymp v)\}$$

This should look familiar: it is the first disjunct in the K- and KM-NORMALITY characterizations of epistemic accessibility. As such, BELIEF in effect identifies (justified) belief with the evidential component of knowledge: ‘evidential’ in the sense that what one believes only depends on one’s evidence, and a ‘component’ in the sense that knowing something requires believing it. BELIEF also implies that belief ‘aims at’ or ‘aspires to’ knowledge, in the sense that belief and knowledge coincide in ‘optimal’ situations that are at least as normal as every situation evidentially accessible from them.²¹

We can now consider how the normality framework responds to the Gettier problem. To do so, consider what would happen if we defined knowledge as justified true belief, so that $R_K(w) = R_B(w) \cup \{w\}$. Call this view JTB-NORMALITY. As with K- and KM-NORMALITY, it can be equivalently characterized in terms of a knowledge-default: given the normality framework’s structural constraints on comparative normality, JTB-NORMALITY is equivalent to the claim that epistemic accessibility is the most restrictive relation compatible with FACTIVITY

²⁰Goodman and Salow (2018) characterizes belief using this definition of doxastic accessibility. Smith (2016) characterizes justification similarly. (His official view deploys only an *at least as normal* relation, so in comparing it to ours we treat him as presupposing COLLAPSE.) There are also natural views that agree with BELIEF about doxastic accessibility without identifying what you believe with what is true in all doxastic possibilities. One is the view, defended by Lenzen (1978), Stalnaker (2006), Rosenkranz (2018), Carter and Goldstein (2021), and Dutant (forthcoming), that a person believes p just in case for all they know they know that p ; this entails BELIEF given TRANSPARENCY and the normality framework. Another such view is that a person believes p in w just in case there is a set of evidential possibilities in which p is true that (i) contains every evidential possibility that is at least as normal as any of its members, and (ii) is such that every evidential possibility not in it is sufficiently less normal than one of its members. It is modeled on Lewis’s (1973, 1981) truth conditions for counterfactuals, which are designed to be well-behaved in the presence of infinite chains of worlds ever closer to actuality, and it differs from the account assumed here only if there can be infinite chains of evidential possibilities each of which is sufficiently more normal than the previous one.

²¹While we here follow the tradition in epistemology of thinking of belief as ‘aiming at’ knowledge, we agree with Hawthorne et al. (2016), Dorst (2019), and Holguín (2022) that ‘believe’ in English expresses something considerably weaker. Goodman and Holguín (forthcoming) argue that ‘be sure’ expresses the weakest doxastic attitude that aims at knowledge. (While they deny that being sure is necessary for knowing, the cases they consider arguably involve failures of TRANSPARENCY, and so are beyond the scope of this paper.)

and INDUCTIVE LIMITS. However, JTB-NORMALITY is not a version of the normality framework, since it does not entail CONVEXITY.

JTB-NORMALITY is vulnerable to Gettier-style counterexamples. According to Williamson (2015), the blame rests on its disjunctive characterization of epistemic accessibility. If this diagnosis were correct, it would be trouble for K- and KM-NORMALITY, since they also characterize epistemic accessibility disjunctively. Moreover, Williamson’s suggestion that disjunctive accounts of epistemic accessibility are vulnerable to Gettier-style counterexamples is not *ad hoc*. To see why, consider an account of knowledge as justified *safe* belief, according to which one knows a proposition just in case it is true in all situations that are either doxastically accessible or sufficiently close to the actual situation. This account predicts that any disjunction of a justified false belief and an unknown but safely true proposition will be known. That prediction is false: such disjunctions are typically not known, despite being true and justifiably believed, for the same reason as in classic Gettier cases.²² This is to be expected if disjunctive characterizations of epistemic accessibility beget Gettier problems.

But we think there is a better diagnosis of why both JTB-NORMALITY and the justified safe belief account are vulnerable to Gettier-style counterexamples – namely, that neither implies CONVEXITY. In support of this diagnosis, consider Bjorn. Let ‘around 175.4’ denote the smallest closed interval such that Bjorn believes that his weight in pounds is in that interval. Now consider a situation in which this belief is false: the scale is significantly in error, so Bjorn’s true weight x is not around 175.4. Then Bjorn has a justified true belief that he weighs either around 175.4 pounds or exactly x pounds. Intuitively, though, this is not something that he knows. This is not predicted by JTB-NORMALITY given any of the hypotheses about the comparative normality of Bjorn’s situations considered in section 3. But it is easily predicted by the normality framework, as explained in section 1. CONVEXITY was crucial to that explanation, as it ensures that the epistemically accessible situations correspond to an interval of weights containing 175.4 pounds and Bjorn’s actual weight, and hence contains some non-actual weights that are not around 175.4 pounds.

This diagnosis generalizes. If at 12pm you come across a clock that (unknownst to you) is broken but happens to read 12pm, then intuitively your justified true belief that it is within a couple minutes of 12pm will not be knowledge; cf. Russell (1949). This is again naturally explained by the normality framework, assuming that it is at least as normal to come across such a stopped clock at 12:15pm as it is to come across it at 12pm, as you actually did.²³ For CONVEXITY

²²Cf. Gettier (1963). While the simple safety theory described in section 2 avoids this problem, it does so at the cost of offering no clear account of justified belief; see Bacon (2014), who for this reason treats ‘epistemic accessibility’ not as an account of what people know but rather as an account of what they are in a position to know by believing accordingly.

²³Given MARGINS, we need only the weaker assumption that some evidentially accessible situation in which the time is 12:15pm is either at least as or *insufficiently less* normal than the actual situation. By contrast, given K-NORMALITY, if the most normal evidential possibilities in which the clock is broken are all ones in which the time is within a few minutes of 12pm (say, because you’ve noticed a queue forming at the coffee shop, which is slightly more normal at the top of the hour), you *will* come to know that it’s within a few minutes of 12pm by

(together with FACTIVITY) then implies that such situations are epistemically accessible, and hence for all you know it is 12:15pm. JTB-NORMALITY fails to make this prediction, as such situations are neither doxastically accessible (since you believe the clock isn't broken) nor actualized.

Consider a third example, from Pritchard (2012). Temp thinks his thermometer works; in fact, it is broken, its readings are fluctuating randomly, and someone is covertly adjusting the thermostat to match its readings. Intuitively, Temp doesn't learn anything about the temperature by reading the broken thermometer. Again, CONVEXITY explains why. As with Russell's clock, reading a broken thermometer that happens to be right doesn't yield knowledge when there is nothing weird going on behind the scenes, since there are evidential possibilities at least as normal as actuality in which the thermometer is broken but the temperature is different. And adding weird goings on behind the scenes won't change this, since these possibilities will remain evidentially accessible and at least as normal as actuality. Importantly, this is true even if the weird goings on render Temp's belief 'safe from error', in the sense that he couldn't easily have been mistaken about the temperature; cf. Beddor and Pavese (2020).

There is an interesting respect in which K-NORMALITY treats Gettier phenomena differently from KM- and I-NORMALITY. Following Williamson (2013a), say that a *purely veridical Gettier case* is a situation where someone has a justified true belief that falls short of knowledge, even though they do not have justification to believe any false proposition. Suppose Bjorn's actual weight is within the interval of weights compatible with his beliefs but near the edge of that interval. MARGINS then implies that this is a purely veridical Gettier case: Bjorn's justified true belief about how much he weighs is not knowledge. By contrast, K-NORMALITY never allows for purely veridical Gettier cases.²⁴

Before moving on, we should address a worry about the normality framework. We noted at the end of section 3 that each of K-, KM-, and I-NORMALITY entails that, if you have the same evidence in w and in v , and w is at least as normal as v , then you know at least as much in w as you do in v . In other words, circumstances cannot be abnormally conducive to inductive knowledge: holding your evidence fixed, you never know more in less normal situations.²⁵ Carter (2019) suggests that this is problematic, writing that "abnormal conditions can be beneficial to the acquisition of knowledge" (p.1797). Explaining our disagreement with Carter helps to illustrate how the normality framework differs from reliabilist theories of knowledge, which are widely thought to make knowledge too easy to come by when abnormally reliable processes are at work; cf. BonJour (1980). The normality framework avoids this prediction by making background evidence about one's inductive methods epistemically significant.

reading the broken clock. This is an important consideration in favor of MARGINS, although we argue in Goodman and Salow (ms) that it is not as decisive as it may at first appear.

²⁴Williamson (2013a,b) suggests, and Cohen and Comesaña (2013) seem to grant, that examples such as Goldman's (1976) 'fake barn' cases are purely veridical Gettier cases. We disagree, since presumably Henry justifiably believes that he isn't surrounded by fake barns. So we don't think such thought experiments provide independent support for MARGINS.

²⁵Circumstances may be abnormally conducive to *evidential* knowledge; by allowing this, the normality framework accommodates cases like "Enhancement" in Valaris (forthcoming).

Consider what happens when, unbeknownst to Bjorn, his scale is abnormally good, in that the errors on its readings are likely to be smaller than on ordinary scales. Carter claims about such cases that “Intuitively, it should be *easier* to acquire knowledge of the value of the relevant physical parameter” (p.1798, emphasis ours). We disagree: not only do we deny that Bjorn having an abnormally good scale allows him to know *more* about his weight than he would had a normal scale given the same reading; we also think that the abnormality of his scale can cause him to know *less* about his weight than he otherwise would have. For an error that would be extreme for a laboratory-grade scale might be perfectly normal for a consumer-grade scale. So suppose that, although Bjorn justifiably believes that his scale is consumer-grade, it is actually laboratory-grade, and its error on this occasion is extreme for such a scale but would not be extreme for a consumer-grade scale. Intuitively, if for all he knows the error is extreme and for all he knows the scale is consumer-grade, then for all he knows the error is extreme and the scale is consumer-grade. This claim follows from FACTIVITY, CONVEXITY, and the natural assumption that a situation in which the scale has normal quality (consumer-grade) but is abnormally in error is at least as normal as the actual situation, in which the scale is both of abnormally high quality (laboratory-grade) and abnormally in error. And from this claim, it follows that Bjorn knows less than he would have known were the scale consumer-grade, since in such a situation the error would have been perfectly normal, and his belief that the error is not extreme would have amounted to knowledge.

5 Comparability

This section considers arguments for and against COMPARABILITY. We will begin with an important argument that COMPARABILITY has overly skeptical implications, which can be avoided by adopting a *multi-dimensionalist* conception of comparative normality. We then consider a possible response to this argument, according to which relations of comparative normality are determined by probabilities in a context-sensitive way, and explain some independent motivations for this *probabilist* approach. We will not come to any firm conclusions about COMPARABILITY here, since its plausibility ultimately turns on whether a reasonably constrained account of the needed context-sensitivity can be given. Fortunately, all of the cases discussed in later sections are naturally modeled in ways that obey COMPARABILITY, so it will not matter for that discussion whether the principle holds in general.

The argument that COMPARABILITY has overly skeptical implications is best brought out by cases in which, intuitively, there are orthogonal ways for things to go epistemically awry. Suppose that, in addition to weighing himself on his scale, Bjorn also takes his temperature with a thermometer, which displays 98.6°F. Plausibly, malfunctions in the two devices are epistemically independent: as long as the thermometer gives an accurate reading, Bjorn knows that his temperature is approximately 98.6°F – even if the scale grossly misreports his weight; and as long as his scale gives an accurate reading, Bjorn knows that he weighs

approximately 175.4 pounds – even if the thermometer grossly misreports his temperature. Denying such independence is worryingly skeptical: since errors do occur, allowing errors of one kind to generate ignorance about independent subject matters threatens much of the knowledge we take ourselves to have.

Now consider a pair of ‘single malfunction’ situations that agree on the readings of both the scale and the thermometer, but that differ in which reading is the accurate one and which reading is grossly in error. If the foregoing is correct, then neither situation is epistemically accessible from the other. But since they agree on the instruments’ readings, they are evidentially accessible from each other. Such a pattern of accessibility is incompatible with COMPARABILITY, which implies that, whenever two situations are evidentially accessible from each other, one of them must be epistemically accessible from the other. In this way COMPARABILITY prevents intuitively orthogonal dimension of epistemic variation from being actually orthogonal, arguably predicting too little knowledge in worlds like ours where things are abnormal in some but not all respects.²⁶

Suppose we accept the conclusion of this argument, and reject COMPARABILITY. Doing so allows us to give a natural model of how there can be orthogonal epistemic errors. The basic idea is that, in such cases, normality is multi-dimensional. If there are many dimensions along which situations vary in normality (such as the accuracy of the scale and the accuracy of the thermometer), then for w to be at least as normal as v is for it to be at least as normal on all dimensions, and for w to be sufficiently more normal than v is for it to be at least as normal and to be sufficiently more normal on at least one dimension.

To get a feel for this proposal, consider a more abstract but simpler example. Suppose that there are two dimensions along which situations vary in normality, and that along each dimension situations can be either ordinary (0), odd (1), or bizarre (2). Ordinary is sufficiently more normal than bizarre; odd is in between, but neither sufficiently less normal than ordinary nor sufficiently more normal than bizarre. We then have nine situations evidentially accessible from each other, depicted at the top of figure 4. COMPARABILITY fails in this model; for example, neither of (0,2) and (2,0) is at least as normal as the other. As a result, it is consistent with the normality framework that neither is epistemically accessible from the other. Indeed, this is a consequence of both K- and KM-NORMALITY: even when things are bizarre along one dimension, you can still know that they aren’t bizarre along the other dimension. In this way the normality framework can accommodate the kind of epistemic independence that anti-skeptical considerations arguably motivate.

This multi-dimensional account is no help given I-NORMALITY. Indeed, given I-NORMALITY, rejecting COMPARABILITY only exacerbates the skeptical problem; for example, by making (2,0) epistemically accessible from (0,1). Intuitively, this means that any abnormality along one dimension makes inductive knowledge impossible along any orthogonal dimension. This is the same problem that

²⁶Rott (2004) and Stalnaker (2006) give a structurally similar argument against the ‘defeasibility analysis’ of knowledge. Baltag and Smets (2008) model the defeasibility analysis using plausibility orders, showing, in effect, that it is a version of the normality framework committed to COMPARABILITY, COLLAPSE, TRANSPARENCY and the principle STATISM from section 6.




Figure 4: multi-dimensionality without COMPARABILITY

arose when combining I-NORMALITY with the COMPARABILITY-violating model of Bjorn depicted in figure 3. These skeptical consequences justify our earlier claim that I-NORMALITY is untenable if COMPARABILITY is rejected.

While the above considerations provide a powerful argument against COMPARABILITY, they are not the end of the story. This is because there are also powerful considerations in the other direction. In particular, COMPARABILITY is a consequence of natural accounts of comparative normality in terms of comparative probability. Recall our first assumption in section 1 about the comparative normality of Bjorn’s evidential possibilities: that a situation in which the scale overestimates his weight by a given amount is at least as normal as any situation in which the scale overestimates his weight by a greater amount. Why should this be? A natural explanation is that, given Bjorn’s evidence, more extreme errors in the scale are *less probable* than less extreme ones. Indeed, if this weren’t so (say, because Bjorn had prior evidence that the scale was miscalibrated and hence disposed to appreciably overestimate his weight), then the assumption would lose its plausibility. In the next two sections we will describe a number of further thought experiments where comparative normality and comparative probability also seem to be aligned. Taken together, these cases are suggestive of some general connection between normality and probability.

In Goodman and Salow (2021) we show how this basic thought can be fleshed out into an attractive probabilist theory of comparative normality. For present

purposes we can focus on the following three features of the theory. First, its account of one situation being *at least as normal as* another both vindicates COMPARABILITY and involves a certain kind of context-sensitivity. Second, this independently motivated context-sensitivity can be used to respond to the argument that COMPARABILITY has overly skeptical implications. Third, its account of one situation being *sufficiently more normal than* another vindicates the independently attractive principle that only highly probable propositions can be known, which the COMPARABILITY-rejecting strategy fails to vindicate. We will explain these three features in turn.

The basic idea is that one evidentially accessible situation is at least as normal as another just in case it is at least as probable as the other given your evidence. There are two complications. The first is that these probabilities are not of particular maximally specific situations, but rather of classes of relevantly equivalent situations: what matters is not how likely it is that every aspect of a particular situation obtains, but how likely it is, say, that Bjorn weighs as much as he does in a given situation. For this reason, the account has an equivalence relation on situations as an additional parameter. We can think of this as a *question*, whose complete answers correspond to equivalence classes under the relation. The probabilities that determine two situations' comparative normality are then the probabilities of those situations' respective answers to the question.²⁷

The second complication concerns cases where the question has a continuum of answers, each of which has probability zero. *How much does Bjorn weigh* is such a question, since there are a continuum of weights that Bjorn might have, given his evidence. In cases like this, what matters for two situations' comparative normality isn't their associated probabilities but rather their associated probability *densities*. Suppose, for example, that the evidential probabilities of Bjorn having various weights are given by a 'bell curve' centered at 175.4. Then the comparative normality of two situations is given by the comparative heights of this curve for the weights he has in those situations. This vindicates our aforementioned assumption about which of Bjorn's situations are at least as normal as which others. The account also vindicates COMPARABILITY, since for any two probabilities (or probability densities) one is at least as great as the other.

Although *how much does Bjorn weigh* is a natural question when considering Bjorn's knowledge of his weight, it is not a natural question when considering his knowledge of his temperature. So in order for the probabilistic account of normality to be plausible, the relevant question should be context dependent. The relation expressed by "know" will then be different in different contexts: different standards for ignoring irrelevant distinctions among situations will be operative, leading to different assignments of probabilities to situations, and hence to different normality relations, which then determine different epistemic accessibility relations.

This context-sensitivity suggests a natural response to the worry that COM-

²⁷For other probabilistic, contextualist, and question-sensitive accounts applied directly to belief, see Leitgeb (2014) and Holguín (2022).

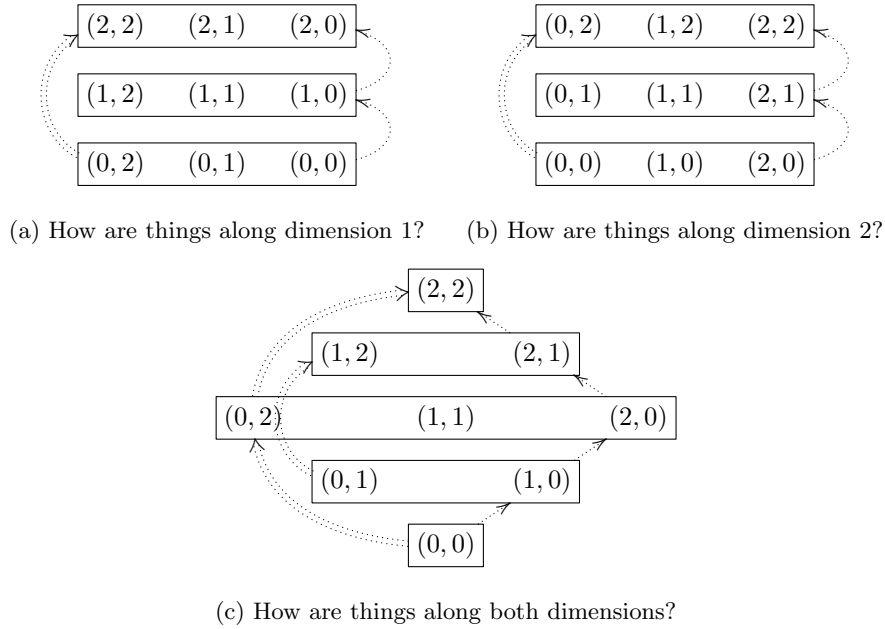


Figure 5: contextualism

PARABILITY leads to skepticism. The idea is that usually, when we focus on what Bjorn knows about his weight, context supplies the question *how much does he weigh*, ignoring non-weight differences among evidential possibilities (such as differences in his temperature). Likewise, when we focus on what Bjorn knows about his temperature, the contextually relevant question is *what is his temperature*, ignoring non-temperature differences among evidential possibilities (such as differences in his weight). So while there is no context in which accurate-scale/malfunctioning-thermometer situations and accurate-thermometer/malfunctioning-scale situations are mutually epistemically inaccessible (since COMPARABILITY holds in all contexts), a typical use of “Bjorn knows he weighs around 175.4 pounds” is in a context relative to which it expresses a proposition that is true in the former situation, and a typical use of “Bjorn knows his temperature is around 98.6°F” is in a context relative to which it expresses a proposition that is true in the latter situation. Figure 5 illustrates this contextualist response to the skeptical threat posed by COMPARABILITY, again using a simplified normal/odd/bizarre two-dimensional example.

To derive these predictions, we need a probabilistic account of one situation being sufficiently more normal than another. Here is one natural proposal: w is sufficiently more normal than v just in case the evidential probability of things being more normal than v , given that they are no more normal than w , is sufficiently high. This delivers the normality relations depicted in figure 5 if (i) 0 is ten times as probable as 1 and 1 is ten times as probable as 2 (for each

dimension), (ii) the dimensions are probabilistically independent, and (iii) the threshold for ‘sufficiently high’ probability is .95.

One attraction of this account is that, in cases involving continuous probability distributions like Bjorn’s, it predicts that belief is characterized by the *highest posterior density region*, which is the standard way in Bayesian statistics of giving a qualitative summary of such distributions; cf. Kruschke (2014).²⁸ Another attraction is that, as explained in the next section, it predicts the rough pattern of failures of COLLAPSE that we argued for in section 3. Finally, and perhaps most importantly, the probabilist account validates the plausible principle that, if someone knows or justifiably believes that p , then the probability of p given their evidence must be sufficiently high.

The multi-dimensionalist account, by contrast, predicts the opposite. Consider what happens when we move from two dimensions to tens or hundreds of dimensions. As with two dimensions, the multi-dimensionalist account predicts that, if things are normal along every member of some set of dimensions, then one knows that things are not bizarre along any of those dimensions. If the set is big enough, then this will be knowledge of something improbable (even if for any given dimension it is highly probable that things are not bizarre along it).

This high-dimensional example is reminiscent of Makinson’s (1965) preface paradox. It seems that a historian might have inductive knowledge of most of the claims in their new book, despite the conjunction of those known claims having low probability given the evidence. The multi-dimensional approach can accept this description of the case, but the probabilist account cannot. Probabilism offers a different, contextualist diagnosis: although for most claims in the book there is a natural context relative to which the author knows them, there is no single context relative to which the author knows all of them. And unlike the multi-dimensional account, probabilism allows that there are also natural contexts relative to which the author can know that not all of the claims in their book are true. We give a formal model of this kind of case in Goodman and Salow (2021, section 5).²⁹

We have no settled view about which of the multi-dimensionalist and proba-

²⁸This distinguishes the present account from that of Goldstein and Hawthorne (2022), whose ‘Probabilistic Margin for Error’ (a version of which we also express sympathy for in Goodman and Salow (2021, note 10)) entails that people sometimes believe less than this.

²⁹Carter and Goldstein (2021) propose a different, non-contextualist way of accommodating knowledge that things are not in all respects normal. In effect, they propose that, in addition to the ‘first-order’ dimensions of normality, there is a further ‘second-order’ dimension representing how much abnormality is present across the first-order dimensions, and situations which are maximally normal along every first-order dimension are (when there are enough dimensions) sufficiently abnormal along this second-order dimension. These dimensions cannot be combined into an overall normality ordering, since first-order normality and second-order normality sometimes conflict with each other. But they can still be used to characterize the epistemically accessible situations as those evidentially accessible situation that are not sufficiently less normal along any dimension, including the second-order one. This allows Carter and Goldstein to construct models in which one can know that things are not ordinary along every first-order dimension. But, like our multi-dimensionalist picture and unlike our probabilistic one, they still predict that, if things are normal along each of a certain set of dimensions, then one knows that things are not bizarre along any of those dimensions.

bilist approaches is correct. Adjudicating between them is an important question for further work. The remainder of this paper remains neutral on this issue by focusing on cases where both approaches plausibly agree in their predictions.

6 Dynamics

In this section we will show how the normality framework can be used to model how people’s knowledge changes when they get new evidence.³⁰ Doing so requires introducing models with additional structure. As an illustration, we will show how these models can be used to make predictions about the dynamics of Bjorn’s knowledge if he weighs himself more than once. This is a simple version of a central problem in inductive epistemology, namely how to combine independent observations to estimate the value of a parameter. In the next section we will apply these models to other core kinds of inductive knowledge.

A tempting thought is that getting more evidence is a matter of fewer situations becoming evidentially accessible. But ordinary processes of evidence-gathering – like weighing yourself – aren’t like this. Before you’ve weighed yourself, your evidence (ordinarily) will entail that you haven’t weighed yourself yet; after you’ve weighed yourself, your evidence will entail that you have weighed yourself already. So no situation evidentially accessible after you weigh yourself was evidentially accessible before you weighed yourself. Part of what is going on is that evidential accessibility encodes not only facts about what your evidence entails about your environment, but also facts about what your evidence entails about what your evidence entails. Ordinarily, when we think about ending up with *more* evidence than we started with, what we mean is that our later evidence entails more *about a given subject matter* than our earlier evidence did. This will be our guiding idea here.

To implement this idea, we model situations as having additional structure, representing how things are with respect to a given subject matter and what your evidence implies concerning that subject matter. Formally, we start with a set of *states*, each of which corresponds to a complete specification of how things are with respect to the relevant subject matter. Situations are then ordered pairs $\langle x, Y \rangle$, where x is the actual state of the world, and Y is the set of states of the world compatible with your evidence. Not all such pairs are possible situations; at a minimum, x must be a member of Y , since the truth must be compatible with your evidence. A situation $\langle x', Y' \rangle$ is evidentially accessible from a situation $\langle x, Y \rangle$ just in case $Y = Y'$ – that is, just in case they agree on what evidence you have about the state of the world.³¹ In our running example of Bjorn weighing himself, we can treat states of the world as encoding facts about how much he weighs and what the scale reads when he steps on it.

³⁰In Goodman and Salow (ms b) we explore these models’ predictions for the dynamics of belief, and compare them to more familiar models of belief revision. Baltag and Smets (2008) discuss how models similar to ours are related to work in dynamic epistemic logic.

³¹This implies TRANSPARENCY, and it is not clear whether the current strategy for modelling the dynamics of knowledge can be extended to cases where TRANSPARENCY fails.

Say that a situation $\langle x', Y' \rangle$ is *the result of discovering p* in a situation $\langle x, Y \rangle$ just in case $x' = x$ and $Y' = Y \cap p$. This notion of discovery allows us to make precise our earlier dynamic locutions about evidence accumulation, by allowing us to model acquiring a new piece of evidence p as a change in one's situation.³² Where there is no danger of confusion we ignore the distinction between a set of states and a proposition that is true in all and only the situations in which the state of the world is a member of that set. When we speak of what someone knows or believes 'about the state of the world', or state dynamic principles such as NO DEFEAT (discussed below), we are confining our attention to such propositions, and setting aside propositions that have additional implications about what evidence one has.

To see this machinery in action, consider an elaboration of the Bjorn case in which he weighs himself on two scales, one after another. The scales are made by different manufacturers, neither of which Bjorn has heard of, so he has no reason to think either scale is more accurate than the other or that their errors will be correlated. States of the world correspond to triples $\langle x, y_1, y_2 \rangle$, where x is Bjorn's actual weight, y_1 is the reading of the first scale, and y_2 is the reading of the second scale. For each state s , there are three situations. In the first, Bjorn's evidence is characterized by the set of all states; in the second, his evidence is characterized by the set of all states that agree with s about the reading of the first scale; in the third, his evidence is characterized by the set of all states that agree with s about the readings of both scales. We want to know how Bjorn's beliefs and knowledge evolve as he moves between these situations.

A natural strategy for thinking about the normality relations between these situations is to treat them as parasitic on normality relations between their component states. Two states' comparative normality is determined by comparing the sum of the squares of the scales' errors; sufficient differences in normality correspond to differences in these sums that exceed a given amount. Formally, the state $\langle x, y_1, y_2 \rangle$ is at least as normal as the state $\langle x', y'_1, y'_2 \rangle$ just in case $(x - y_1)^2 + (x - y_2)^2 \leq (x' - y'_1)^2 + (x' - y'_2)^2$, and sufficiently more normal just in case $(x - y_1)^2 + (x - y_2)^2 + c^2 < (x' - y'_1)^2 + (x' - y'_2)^2$, where c is a positive constant. A situation $\langle x, Y \rangle$ is at least as normal as (/sufficiently more normal than) a situation $\langle x', Y' \rangle$ just in case x is at least as normal as (/sufficiently more normal than) x' . This is a natural model both on account of its qualitative predictions and on account of the fact that it agrees with the predictions of the probabilistic account of comparative normality mentioned in the last section, given natural probabilistic assumptions.³³ We will explain these points in turn.

³²This notion of discovery is logical, not temporal (although we sometimes use temporal idioms to describe it for ease of exposition). Being in w sometime before you are in v neither implies nor is implied by the existence of any p such that discovering p in w results in v . Indeed, when we forget things, a temporally prior situation w may be the result of discovering p in a later situation v . Discovering p is also not a one-to-one operation: there can be distinct situations w and w' such that v is the result of discovering p in both of them. So there is no well-defined operation of "forgetting p ".

³³Strictly speaking, the probabilistic account predicts these normality relations among mutually evidentially accessible situations while making other pairs of situations incomparable, since probabilism requires incomparabilities between mutually evidentially inaccessible situa-

The model implies that, before weighing himself, Bjorn believes that the scales' readings will not differ by more than $\sqrt{2}c$. But he does not yet believe anything about his weight. After then seeing the first scale read y_1 , he believes that his weight lies in the (closed) interval $y_1 \pm c$. Finally, after also seeing the second scale read y_2 , he believes that his weight lies in the interval $\frac{y_1+y_2}{2} \pm \frac{c}{\sqrt{2}}$. If both readings are perfectly accurate (so that $x = y_1 = y_2$), then his beliefs amount to knowledge before weighing himself, after weighing himself once, and after weighing himself a second time. Given K-NORMALITY, the same is true as long as the scales' errors are not too large (so that $(x - y_1)^2 + (x - y_2)^2 \leq c^2$). By contrast, given MARGINS, any error in either scale's reading implies that (i) after weighing himself once, Bjorn has some beliefs about how much he weighs that are not knowledge, (ii) even before weighing himself, Bjorn has some beliefs about how far apart the two readings will be that do not amount to knowledge, and (iii) after weighing himself twice, all of his beliefs amount to knowledge just in case his true weight is the average of the two readings (so that $x = \frac{y_1+y_2}{2}$).

An important feature of this model is that it generalizes smoothly to more than two measurements, as well as to a single measurement.³⁴ In the case of a single measurement, it yields a version of the model from section 3 that accepts COMPARABILITY but rejects COLLAPSE. However, unlike the normality relations depicted in figure 2, which correspond to the accessibility relations proposed by Williamson (2013a), the sensitivity to the square of the scale's error rather than the absolute value of its error means that, in absolute terms, the differences in error-magnitude needed to make for sufficiently greater abnormality decrease as the error-magnitudes increase; cf. Goodman (2013, §2).

Comparing the one-scale and two-scale versions of the model reveals an important feature: Bjorn's knowledge about his weight after two measurements is not the conjunction of what he would have known about it had he performed only the first measurement and what he would have known about it had he performed only the second measurement. He may know more than is entailed by that conjunction (if both measurements are highly accurate) or he may know less (if one measurement is highly accurate but the other is highly inaccurate). This is a good prediction: additional readings really are epistemically beneficial when they are highly accurate, and epistemically harmful when they are highly inaccurate; cf. Goldstein and Hawthorne (2022).

A central question about the dynamics of knowledge is whether getting new evidence can destroy prior knowledge. The following principle rules this out:

NO DEFEAT: If you know p , then you still know p after discovering q .

Whether our model validates this principle depends on which version of the normality framework we adopt. Given K-NORMALITY, Bjorn can lose knowledge about how much he weighs. Consider any state $\langle x, x, y \rangle$ where $|x - y| \leq c$. After weighing himself once, Bjorn's knowledge about his weight is characterized by

tions for reasons explained in Goodman and Salow (2021, note 8). We ignore this subtlety in what follows, since it makes no epistemic difference.

³⁴For n measurements, $\langle x, y_1, \dots, y_n \rangle \succeq \langle x', y'_1, \dots, y'_n \rangle$ iff $\sum_i (x - y_i)^2 \leq \sum_i (x' - y'_i)^2$, and $\langle x, y_1, \dots, y_n \rangle \gg \langle x', y'_1, \dots, y'_n \rangle$ iff $\sum_i (x - y_i)^2 + c^2 < \sum_i (x' - y'_i)^2$.

the interval $x \pm c$. After weighing himself again, it is characterized by the interval $\frac{x+y}{2} \pm \frac{c}{\sqrt{2}}$. Bjorn therefore loses knowledge whenever the former interval fails to contain the latter, which will be the case whenever $|x - y| > (2 - \sqrt{2})c$.

By contrast, given MARGINS, the model predicts that Bjorn never loses knowledge. This fact is an instance of a more general result about models in which evidential possibilities' comparative normality is determined by their underlying states. Such models satisfy the following condition:

STATISM: Comparative normality is evidence-independent:

$$\begin{aligned} \langle x, Y \rangle \succeq \langle x', Y \rangle & \text{ if and only if } \langle x, Y' \rangle \succeq \langle x', Y' \rangle, \text{ and} \\ \langle x, Y \rangle \succ \langle x', Y \rangle & \text{ if and only if } \langle x, Y' \rangle \succ \langle x', Y' \rangle. \end{aligned}$$

The general result is that MARGINS, COMPARABILITY and STATISM together entail NO DEFEAT.³⁵ Since the above model satisfies COMPARABILITY and STATISM, it also validates NO DEFEAT given MARGINS.

A surprising prediction of the model is that what Bjorn knows after weighing himself once depends on what the second scale will read when he weighs himself again. For example, after a perfectly accurate reading from the first scale, he knows more if it is going to be followed by a perfectly accurate reading on the second scale than he does if it is going to be followed by a reading on the second scale whose error is greater than c (given K-NORMALITY) or merely greater than 0 (given MARGINS). While this prediction takes some getting used to, we think it is ultimately the right result, given the naturalness of the model overall.³⁶

We mentioned above that our model can be motivated by combining natural probabilistic assumptions with the probabilistic account of comparative normality sketched in section 5 and developed in Goodman and Salow (2021). We think such a probabilistic treatment is plausible in the present case even if it is rejected as a general account of comparative normality.³⁷ We will now explain the four required assumptions about the probabilities given Bjorn's evidence.

The first assumption is that, prior to weighing himself, the probabilities for different errors on the first scale are characterized by a Gaussian probability

³⁵Section 8 describes counterexamples to NO DEFEAT given MARGINS and COMPARABILITY that arise from failures of STATISM. Here is a formal counterexample arising from failures of COMPARABILITY given MARGINS and STATISM. There are three states a, b, c , which we treat as the primary bearers of normality relations. a is sufficiently more normal than b , both of which are incomparable with c . Given KM-NORMALITY, in $\langle c, \{a, b, c\} \rangle$ you know that b doesn't obtain; but in $\langle c, \{b, c\} \rangle$ (after discovering that a doesn't obtain) you no longer know this.

³⁶One could resist the prediction by adopting the multi-dimensionalist approach from section 5, with the error on observed scales and the error on unobserved scales being independent dimensions of normality. The most attractive version of this strategy rejects STATISM, so that as more scales are read those scales' errors become comparable along a single dimension. Such a model can agree with ours about what Bjorn believes about his weight at every stage, as well as about what he knows after reading both scales. See also note 40.

³⁷In particular, even those who prefer the multi-dimensionalist treatment of the case where Bjorn both weighs himself and takes his temperature should reject a parallel treatment of Bjorn weighing himself on two scales. This is because a parallel treatment wrongly predicts that his knowledge about his weight after two measurements is the conjunction of what he would have known about his weight had he performed only the first measurement and what we would have known about his weight had he performed only the second measurement.

density function (a normal distribution, or ‘bell curve’) centered on 0 (that is, with accurate readings being most likely, and with overestimations being as likely as equally large underestimations), and the probabilities for different errors on the second scale are characterized by the same distribution (since Bjorn has no relevant evidence distinguishing the scales). The second assumption is that these distributions are independent, so that the probability of one scale having a particular error is the same given any hypothesis about the error of the other scale. The third assumption is that these probabilities are unchanged after Bjorn weighs himself the first time (as any pair of errors remains consistent with his evidence). The final assumption is that, after weighing himself a second time, his probabilities for different pairs of errors are the result of conditioning his prior probabilities on his new evidence, which now entails that the difference between the scales’ errors equals the observed difference between their readings. These four assumptions, together with the probabilistic account of comparative normality, entail the normality relations advertised above.

The Gaussian probability distributions featuring in these assumptions are ubiquitous in scientific modeling, for many reason. Especially relevant here is the fact that the probability distributions described above are the only ones such that, after any sequence of measurements, Bjorn’s most likely weight is the average of those measurements.³⁸ This is of course an unrealistic idealization – in a realistic case, prior to weighing himself Bjorn will have some non-trivial probabilities, and indeed non-trivial knowledge, about how much he weighs, which will not be completely swamped by his subsequent observations.³⁹ We can make the model more realistic by assuming that Bjorn’s probabilities about his weight prior to stepping on the first scale are also given by a Gaussian distribution, effectively treating them as if they resulted from an earlier measurement. The resulting post-measurement probabilities will then still be Gaussian, so that the general shape of the normality relations, in terms of the sums of squares of errors, will be the same, and the core qualitative predictions will be preserved.⁴⁰

³⁸ Another relevant feature of Gaussians, to do with the Central Limit Theorem, is mentioned in section 7. Jaynes (2003, ch. 7) catalogues the above mentioned and many other remarkable properties that are unique to Gaussian distributions, and explains why as a result Gaussians are the uniquely natural probability distributions for a wide range of applications.

³⁹ Indeed, our assumptions allow Bjorn to be initially completely ignorant about how much he weighs only by entailing that, prior to weighing himself the first time, he lacks any well-defined probabilities about how much he weighs.

⁴⁰ Our account here is closely related to that of Goldstein and Hawthorne (2022), who also model knowledge about a given quantity based on multiple observations, and do so in terms of comparative normality relations that are generated from probabilities. The main difference is that they predict failures of NO DEFEAT even given MARGINS, and do not predict current knowledge to be sensitive to future observations. In effect, their proposal is the result of applying our probabilistic account of comparative normality relative to the question *how much does Bjorn weigh*, rather than (as in our model) relative to the question *what will the scales’ errors be*. So those sympathetic to a contextualist probabilistic account of normality can hold that there is no substantive disagreement between us – each model correctly describes what Bjorn “knows” in different natural contexts. However, if one is not such a contextualist (and Goldstein and Hawthorne’s framework does not incorporate any such context-sensitivity), then we think our model is preferable, since, unlike theirs, it allows Bjorn to have non-trivial inductive knowledge not only about his weight but also about the scales’ future readings.

7 Prediction and generalization

This section considers two applications of the dynamic framework from the last section. Although the cases are highly artificial, involving repeated coin flips, they illustrate how the normality framework can be used to model non-trivial knowledge about the chancy future and inductive knowledge of lawful regularities. However, we will see that doing so requires rejecting Martin Smith’s influential ‘explanationist’ conception of comparative normality.

Here is the first case, from Dorr et al. (2014):

Flipping for Heads

A coin flipper will flip a fair coin until it lands heads. Then he will flip no more.

Suppose you have evidential knowledge that this is the setup. The coin will in fact land heads on the first flip. Dorr et al. argue that, on pain of skepticism, in such a case you can know that the coin will not be flipped 1000 times. We agree, and following Goodman and Salow (2018) suggest the following model. The states are s_1, s_2, \dots , where s_i is the state in which the coin lands heads on the i th flip. (We ignore the possibility of the coin landing tails forever.) Like the model in the last section, normality relations between situations are parasitic on the normality relations between their underlying states, thereby satisfying STATISM: s_i is at least as normal as s_j just in case $i \leq j$, and s_i is sufficiently more normal than s_j just in case $i + c < j$, where $0 < c \ll 1000$ is “the largest real number such that a $.5^c$ chance still qualifies as substantial” (p. 187). This model predicts that, if your evidence is that the coin will be flipped at least n times, then what you believe is that it will be flipped no more than $n + c$ times.

Like the model of Bjorn’s repeated weighings discussed in the previous section, this model again illustrates the incompatibility between K-NORMALITY and NO DEFEAT. Suppose the coin will land heads on the second toss. Since this outcome is not sufficiently less normal than any other, K-NORMALITY implies that all of your beliefs amount to knowledge. So before the experiment, you know that the coin won’t be tossed more than $c + 1$ times. This knowledge is lost when you see the coin land tails on the first toss; you then know only that it won’t be tossed more than $c + 2$ times (since there are no longer any evidentially possible states sufficiently more normal than s_{c+2}). Such defeat is again a distinctive prediction of K-NORMALITY: MARGINS implies that all you know to begin with is that the coin won’t be tossed more than $c + 2$ times (since the actual outcome s_2 isn’t sufficiently more normal than s_{c+2}).⁴¹

This picture of what it is reasonable to believe about the outcomes of chancy processes is very different from the one defended by Smith (2016, 2017, 2018a). He writes that “Given that a coin is fair, it would be equally normal for it

⁴¹**Flipping for Heads** closely resembles the surprise exam paradox, and the choice of whether or not to accept K-NORMALITY or MARGINS is analogous to the choice between diagnosing the surprise exam paradox as involving a failure of the KK thesis, as Williamson (2000, chapter 6) does, or a failure of NO DEFEAT, as Kripke (2011) does; see also Holliday (2016), who suggests that which diagnosis is appropriate may depend on the details of the case.

to be flipped and land heads and for it to be flipped and land tails” (2018a, p. 731). Whether this marks a disagreement between us is a delicate question, since Smith operates with only one relation of comparative normality where we have two. His notion of one situation being more normal than another plays a role in his theory of justified belief analogous to our notion of one situation being sufficiently more normal than another. And we can agree with Smith that the outcome of a fair coin flip cannot by itself make for a *sufficient* difference in the normality of two situations, and the model just described respects this judgment. Where we disagree with Smith is over the stronger principle that the outcome of a fair coin flip cannot by itself make for *any* difference in the normality of two situations. Smith’s commitment to this stronger principle is evident from the fact that he thinks that you shouldn’t believe anything about the outcome of the experiment beyond what is entailed by your evidence. We reject this principle, since we think s_1, \dots, s_{c+2} are a case of small differences in normality adding up to a sufficient difference in normality.

Smith arrives at his view by appealing to a particular conception of normality: “To describe an event or a situation as ‘abnormal’ ... is to mark [it] out ... as something that would require special explanation if it were to occur or come about” (Smith, 2018b, p.3). We agree with Smith that striking outcomes of repeated coin flips sometimes have no special explanation (that is, none that wouldn’t be shared by more mundane outcomes), and hence don’t require special explanation. We thus reject the supposed connection between being sufficiently less normal and requiring a special explanation.⁴²

Our second case further illustrates what is at stake in our disagreement with Smith. Dorr et al. (2014) write, and we agree, that “Surely, if you could ever learn anything non-trivial about objective chances, you could learn that a certain double-headed coin is not fair by flipping it repeatedly, seeing it land heads each time, and eventually inferring that it is not fair.” So consider:

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.

Surely you can learn that the coin is double-headed after seeing it land heads every time. But this is impossible given STATISM (which Smith’s framework presupposes, and which arguably follows from his preferred conception of normality) and Smith’s desideratum that two situations are equally normal whenever they differ only in the outcome of a fair coin flip.

⁴²Cases like **Flipping for Heads**, in which we have justified beliefs about the future that aren’t entailed by eternal generalizations, are a challenge for a broad class of explanationist accounts of induction, such as the familiar idea that induction is a matter of drawing out the consequences of the best explanation of one’s evidence; see Byerly (2013). (McCain (2014) replies by appealing to the idea that we can justifiably believe that things will unfold normally; Byerly and Martin (2015) respond that things unfolding normally isn’t generally a consequence of the best explanation of our evidence.) Elliott (2021) also argues that this explanationist epistemology cannot explain our knowledge of objective chances.

Smith appears willing to bite the bullet here, since he explicitly defends the view that discovering that someone has won a lottery fifty years in a row would not justify you in believing the lottery was rigged (2016, chapter 3.4). We think this reaction is unacceptably skeptical. For, as Bacon (2014) points out, **Heading for Heads** is a simplified model for how we learn about laws of nature: we observe a particular outcome over and over and eventually conclude that the regularity we see arises from an underlying law as opposed to random chance. So denying that one can learn that a coin is double-headed by flipping it repeatedly seems to entail skepticism about scientific knowledge.

Smith’s view that the outcomes of coin flips make no difference to situations’ comparative normality also threatens the possibility of gaining inductive knowledge about measurable magnitudes using scientific instruments. This is because the errors in instrument readings arise, in part, from adding up the impacts of many independent chance outcomes that occur in the measuring process. The resulting ‘noise’ is then analogous to the heads/tails disparity among many independent coin flips. In this way skepticism about heads/tails disparities extends to skepticism about noise in the measurement processes. Since such noise is the result of adding up many independent chance outcomes, it will have an approximately Gaussian distribution (by the Central Limit Theorem), meaning that extreme errors will have a small but non-zero chance of occurring, and hence (by Smith’s lights) are epistemically possible. In this way a skeptical position about our knowledge of coin flips threatens the possibility of appreciable instrument-based knowledge, such as Bjorn’s inductive knowledge about his weight; cf. Goodman (2013, §3).

Once we allow the outcomes of coin flips to make a difference to situations’ comparative normality, we can give a simple model of **Heading for Heads** according to which, after discovering that the coin lands heads every time, you will know that it is double-headed. This model obeys STATISM (like our earlier models of Bjorn’s repeated weighings and of **Flipping for Heads**), so it can be specified in terms of normality relations between states. Among states in which the coin is fair: one in which it lands heads n times is at least as normal as one in which it lands heads m times just in case the chance of a fair coin landing heads exactly n times out of 100 is at least as high as the chance of it landing heads exactly m times out of 100; and it is sufficiently more normal just in case the probability of having a heads:tails ratio as extreme as $m:(100 - m)$, conditional on having a ratio at least as extreme as $n:(100 - n)$, is sufficiently low. The state in which the coin is double-headed is as normal as any state in which it is fair and lands heads exactly 50 times. This model predicts both that you cannot know that the coin is double-headed at the outset, and that you will know it is double-headed after seeing it land heads enough times.⁴³ Straightforward generalizations of this model likewise predict that you can learn how many sides of a die are painted red and how many are painted green by observing the

⁴³While this model illustrates the tenability of these predictions, it is oversimple, since it predicts that, after seeing the coin land heads the first fifty times, you won’t yet know that it is double-headed. Fortunately, the probabilist strategy outlined in section 5 yields more natural models that allow for such knowledge, as described in Goodman and Salow (2021, §4).

outcome of enough rolls of the die. So such models illuminate our knowledge not only of law-like generalizations, but also of facts about the objective chances.

8 Inductive dogmatism

This section argues that upgrading merely inductive knowledge to evidential knowledge can change what you know, and that these knowledge dynamics can be explained by the evidence-dependence of comparative normality. At the same time, these dynamics help to motivate the distinction between evidential and inductive knowledge in a way that is independent of the normality framework. The cases we will discuss also present counterexamples to NO DEFEAT that don't turn on which version of the normality framework one adopts, and likewise undermine the following widely endorsed principle:

INDUCTIVE ANTI-DOGMATISM

If p together with what you know doesn't entail q , then you won't know q after discovering p .

While this principle has been less discussed than NO DEFEAT, it is often simply taken for granted.⁴⁴ It says that, if you know q after discovering p , then you already knew the material conditional $p \supset q$ before discovering p .

INDUCTIVE ANTI-DOGMATISM is so named because violations of it would be an analogue for inductive knowledge of what Pryor (2000) calls “dogmatism” about perceptual justification. But even Pryor, who defends dogmatism about perception, in effect presupposes this principle. This is because he takes his position to be incompatible with an inductive epistemology of perception – that is, with a view that likens perceptual knowledge to Bjorn's knowledge of his weight (with facts about one's experiences playing the role of facts about the scale's readings). We agree with both Pryor and his anti-dogmatist opponents that ordinary perceptual knowledge does not threaten INDUCTIVE ANTI-DOGMATISM, since we think perception ordinarily provides evidential rather than merely inductive knowledge. Our reservations lie elsewhere.

The dynamics we are interested in arise in cases like the following:

Flipping for All Heads

A coin flipper will simultaneously flip 100 fair coins until they all simultaneously land heads. Then he will flip no more.

Bus Driver

You've been riding the same bus every day for months, and noticed that it's been the same driver every day. But people move on.

⁴⁴The principle is adapted from Dorr et al. (2014), who call it “inferential anti-dogmatism” and assume it without argument. It is explicitly defended by Bacon (2014); see also Hawthorne (2002) on “The Explainer”. Chalmers (2012) defends a related “frontloading principle” that “if one knows M with justification from E [...], then one can have conditional knowledge of M given E with justification independent of E ” (p. 162).

Decay

A radioactive atom is created. It will eventually decay.

Light Bulb

You've installed a brand new light bulb. But nothing lasts forever.

We will analyze these in stages, starting from judgments about the dynamics of belief and then drawing out their implications for the dynamics of knowledge.

In each of the above cases there is an event that you both justifiably believe will have happened after a long while and justifiably believe won't happen for a little while. And if you then discover, after a little while, that the event hasn't happened yet, you should adjust your beliefs about when it will happen. These two judgments already have the surprising consequence that you should sometimes change what you believe as a result of discovering something you already believed (which is ruled out by most theories of belief revision). And while these dynamics can be captured using the formalism introduced in section 6, they cannot arise in models satisfying STATISM, and hence require rejecting the evidence-independence of comparative normality.

Most importantly, these belief dynamics have striking implications for knowledge. Consider **Flipping for All Heads**. Let n be the least number such that you initially believe that the experiment won't take more than that many trials. You also initially believe that the experiment will take more than one trial. Now consider what happens when you then discover that the experiment doesn't end on the first trial. This discovery should shift your beliefs: given the symmetries in the objective chances, $n + 1$ should now be the least number such that you believe that the experiment won't take more than that many trials. So discovering something you already believed requires you to give up a belief. And this could happen in a case where, initially, all of your beliefs are knowledge.⁴⁵ So we have a counterexample to the following principle:

NO DEFEAT FROM WHAT YOU ALREADY KNOW

If you know q and know p , then you'll still know p after discovering q .

This example shows that upgrading merely inductive knowledge to evidential knowledge can change what you know, and thereby motivates the distinction between evidential and inductive knowledge in a way that is independent of the normality framework.⁴⁶ It is especially striking that upgrading some inductive knowledge to evidential knowledge can be epistemically deleterious, causing you to lose knowledge you previously had. Moreover, unlike the counterexamples to

⁴⁵The existence of an evidentially accessible situation in which all of your beliefs are knowledge can be motivated in several ways: it follows from K-NORMALITY, from the plausible principle that one never has justification to believe anything of the form ' p and I don't know that p ', and from the assumption that some epistemically accessible situation is at least as normal as any other (which is plausible in this case).

⁴⁶Zardini (2017) uses a similar example to make roughly this point, framed as an objection to Williamson's thesis that all knowledge is evidence. For other arguments that inductive knowledge isn't evidence, see Littlejohn (2011), Dunn (2014), and Bacon (2014).

NO DEFEAT discussed in the previous two sections, this conclusion cannot be avoided by rejecting K-NORMALITY in favor of MARGINS. Defeat is inevitable.

If upgrading a piece of inductive knowledge to evidential knowledge can destroy another piece of knowledge, can it also create new inductive knowledge? If so, this would yield a counterexample to INDUCTIVE ANTI-DOGMATISM, since it would yield a counterexample to the following even weaker principle:

NO LEARNING FROM WHAT YOU ALREADY KNOW

If you know q and don't know p , then you still won't know p after discovering q .

We will now argue that this principle fails in the above examples too.

The argument is largely parallel to the previous one. Let m be the greatest number such that you initially believe that the experiment will take at least that many trials. After discovering that the experiment didn't end on the first trial, $m + 1$ is now the greatest such number. So discovering something you already believed (that the experiment wouldn't end on the first trial) generates a new belief (that the experiment will take at least $m + 1$ trials). This is a counterexample to the belief-analogue of NO LEARNING FROM WHAT YOU ALREADY KNOW, and thus to the belief-analogue of INDUCTIVE ANTI-DOGMATISM.

Again, this result about the dynamics of belief extends to the dynamics of knowledge. Consider a version of the case where, after discovering that the experiment didn't end on the first trial, all of your beliefs amount to knowledge. So you have gained new knowledge by discovering something you already believed. Moreover, it is hard to deny that this prior belief was also knowledge. For given the large numbers involved, m is plausibly much greater than 1. Moreover, discovering that the experiment doesn't end on the first trial doesn't seem to vastly improve your epistemic position. So since after the discovery you know that the experiment won't end in the next $m - 1$ trials, presumably before that discovery you knew at least that the experiment wouldn't end on the very next trial. If this is right, then discovering something you already knew resulted in new knowledge, in violation of INDUCTIVE ANTI-DOGMATISM.

Similar arguments can be made about **Decay**, **Bus Driver**, and **Light Bulb**. The latter two cases are more delicate, since as time passes the chances of the bus driver soon retiring, or of the light bulb soon breaking, increase. So after a year you might believe that your light bulb is closer to breaking than you did when you first installed it. But it is still plausible that, at each point, you believe that the bulb will last a little longer, which suffices for our argument. Indeed, denying that in **Bus Driver** you should believe each day that you will have the same driver tomorrow seems to imply a problematic form of skepticism about 'enumerative induction' (that is, about the possibility of knowing that the next observation will fit a pattern in one's previous observations).

We do not reject INDUCTIVE ANTI-DOGMATISM lightly: it is a natural principle considered on its own, and it also follows from STATISM, which is an attractively strong constraint on dynamic models in the normality framework. So if the above judgments couldn't be accommodated within natural, constrained models, that would be a reason to revisit those judgments. In Goodman and

Salow (2021, appendix C) we give a formal model of **Decay**, showing how the above behavior can arise in a principled way when normality is analyzed in probabilistic terms, as explained in section 5. That model is too involved to describe here. But many of the striking features we’ve highlighted are already predicted by the probabilist account of comparative normality in simple variants of **Flipping for Heads**, as we will now explain.⁴⁷

Suppose that the coin will in fact land heads on the first flip. Instead of watching the experiment, you discover afterwards (by testimony) something you already knew inductively: that the coin was flipped no more than $c + 1$ times. According to the probabilistic account, discovering this can give you new knowledge that the coin was flipped no more than c times. This is because, given your stronger evidence, this proposition now has above threshold probability; as a result, the coin landing heads after $c + 1$ flips is now sufficiently less normal than it landing heads on the first flip. So there is a general theoretical basis for rejecting NO LEARNING FROM WHAT YOU ALREADY KNOW that does not directly appeal to judgments about thought experiments.

Alternatively, suppose that you discover after the experiment (again by testimony) that the coin was flipped either once or more than $c + 1$ times. This destroys your knowledge that it was flipped no more than $c + 1$ times, since, given your stronger evidence, this claim now has below threshold probability. So there is also a general theoretical basis for failures of NO DEFEAT that turn on failures of STATISM, rather than on failures of COMPARABILITY or MARGINS.⁴⁸

That failures of INDUCTIVE ANTI-DOGMATISM are predicted by a probabilistic account of normality is particularly notable because probabilistic considerations are often invoked to support it and related principles. The basic observation, influentially highlighted by White (2006), is that conditioning a probability distribution on the antecedent of a material conditional cannot increase the probability of that conditional. Suppose that the probabilities given a person’s evidence after discovering p are the result of conditioning their prior probabilities on p . Then the probabilistic *bona fides* of the material conditional $p \supset q$ cannot be better after discovering p than they were beforehand. In light of this fact, it can seem bizarre that the conditional is knowable only after discovering p . But the probabilistic account of comparative normality shows that one can reconcile conditioning as an account of the dynamics of probabilities with the surprising dynamics of knowledge and belief we have motivated here.⁴⁹

⁴⁷More generally, probabilism predicts rampant failures of STATISM; our probabilist models of Bjorn and of **Flipping for Heads** only obey STATISM because of very special probabilistic symmetries with respect to the possible bodies of evidence we considered.

⁴⁸This is not, however, a counterexample to NO DEFEAT FROM WHAT YOU ALREADY KNOW, and such counterexamples are predicted by the probabilistic account only in relatively complex situations, as explained in Goodman and Salow (ms b).

⁴⁹Weatherson (2014) discusses other putative cases of coming to know a proposition by making a discovery that lowers its evidential probability.

9 Conclusion

The ambition of this paper has been to do for inductive knowledge what Lewis (1973) did for counterfactual dependence. The normality framework provides a unified setting in which a range of theories of inductive knowledge can be productively compared and in which general principles about inductive knowledge can be derived from structural conditions on comparative normality, just as Lewis showed how different principles of counterfactual logic correspond to different structural conditions on worlds' comparative similarity. More importantly, just as Lewis's framework gave philosophers new tools to think systematically about patterns of counterfactual dependence as they relate to causation, dispositions, laws of nature, and so on, the normality framework suggests natural models of what people know in a range of independently interesting cases. In particular, we have seen how it can be used to build attractive models of knowledge gained through multiple readings of imperfect instruments, knowledge about the future, knowledge of laws of nature, knowledge spanning multiple subject matters, and knowledge about complex patterns that can be broken down into independent chance events. Finally, just as Lewis theorizes directly about counterfactuals, rather than analyzing them in terms of other modal notions like laws of nature, the normality framework characterizes knowledge directly, rather than analyzing it in terms of belief. At least at the present level of idealization, inductive knowledge may do more to illuminate inductive belief than the other way around.

References

- Andrew Bacon. Giving your knowledge half a chance. *Philosophical Studies*, 171:373–397, 2014.
- Andrew Bacon. Inductive knowledge. *Noûs*, 54:354–388, 2020.
- Brian Ball. Knowledge is normal belief. *Analysis*, 73:69–76, 2013.
- Alexandru Baltag and Sonja Smets. A qualitative theory of dynamic interactive belief revision. In Giacomo Bonanno, Wiebe van der Hoek, and Michael Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory*. Amsterdam University Press, 2008.
- Bob Beddor and Carlotta Pavese. Modal virtue epistemology. *Philosophy and Phenomenological Research*, 101:61–79, 2020.
- Alexander Bird and Richard Pettigrew. Internalism, externalism, and the KK principle. *Erkenntnis*, 86:1713–1732, 2021.
- Laurence Bonjour. Externalist theories of empirical knowledge. *Midwest Studies in Philosophy*, 5:53–73, 1980.
- Ryan Byerly. Explanationism and justified beliefs about the future. *Erkenntnis*, 78:229–243, 2013.

- Ryan Byerly and Kraig Martin. Problems for explanationism on both sides. *Erkenntnis*, 80:773–791, 2015.
- Sam Carter. Higher order ignorance inside the margins. *Philosophical Studies*, 176:1789–1806, 2019.
- Sam Carter. Degrees of assertability. *Philosophy and Phenomenological Research*, 104(1):19–49, 2022.
- Sam Carter and Simon Goldstein. The normality of error. *Philosophical Studies*, 178:2509–2533, 2021.
- Sam Carter and John Hawthorne. Normality. ms.
- David Chalmers. *Constructing the World*. Oxford University Press, 2012.
- Stewart Cohen and Juan Comesaña. Williamson on Gettier cases and epistemic logic. *Inquiry*, 56:15–29, 2013.
- Cian Dorr, Jeremy Goodman, and John Hawthorne. Knowing against the odds. *Philosophical Studies*, 170:277–287, 2014.
- Kevin Dorst. Lockeans maximize expected accuracy. *Mind*, 128:175–211, 2019.
- Jeffrey Dunn. Inferential evidence. *American Philosophical Quarterly*, 51:203–213, 2014.
- Julien Dutant. How to be an infallibilist. *Philosophical Issues*, 26:148–171, 2016.
- Julien Dutant. Knowledge-first evidentialism about rationality. In Fabian Dorsch and Julien Dutant, editors, *The New Evil Demon Problem*. Oxford UP, forthcoming.
- Julien Dutant and Sven Rosenkranz. Inexact knowledge 2.0. *Inquiry*, 63:812–830, 2020.
- Katrina Elliott. Inference to the best explanation and the new size elitism. *Philosophical Perspectives*, 35:170–188, 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.
- Alvin Goldman. *Epistemology and Cognition*. Harvard UP, 1986.
- Simon Goldstein. Fragile knowledge. *Mind*, 131:487–515, 2022.
- Simon Goldstein and John Hawthorne. Knowledge from multiple experiences. *Philosophical Studies*, 179:1341–72, 2022.
- Simon Goldstein and John Hawthorne. Counterfactual contamination. *Australasian Journal of Philosophy*, forthcoming.

- Jeremy Goodman. Inexact knowledge without improbable knowing. *Inquiry*, 56:30–53, 2013.
- Jeremy Goodman and Ben Holguín. Thinking and being sure. *Philosophy and Phenomenological Research*, forthcoming.
- 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. Normality, evidence, and KK. ms a.
- Jeremy Goodman and Bernhard Salow. Belief revision normalized. ms b.
- Jeremy Goodman and Bernhard Salow. Telling time with a broken clock. ms.
- Nelson Goodman. *Fact, Fiction, and Forecast*. Harvard UP, 1955.
- Peter Graham. Epistemic entitlement. *Noûs*, 46:449–482, 2012.
- Peter Graham. Normal circumstances reliabilism: Goldman on reliability and justified belief. *Philosophical topics*, 45:33–61, 2017.
- Daniel Greco. Could KK be OK? *Journal of Philosophy*, 111:169–197, 2014.
- John Hawthorne. Deeply contingent a priori knowledge. *Philosophy and Phenomenological Research*, 65(2):247–269, 2002.
- John Hawthorne. The epistemic use of ought. In John Hawthorne and Lee Walters, editors, *Conditionals, paradoxes, and probability: Themes from the philosophy of Dorothy Edgington*. Oxford University Press, 2021.
- John Hawthorne and Maria Lasonen-Aarnio. Knowledge and objective chance. In Patrick Greenough and Duncan Pritchard, editors, *Williamson on knowledge*. Oxford UP, 2009.
- John Hawthorne, Daniel Rothschild, and Levi Spectre. Belief is weak. *Philosophical Studies*, 173:1393–1404, 2016.
- Jaakko Hintikka. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Ithaca: Cornell University Press, 1962.
- Ben Holguín. Thinking, guessing, and believing. *Philosophers’ Imprint*, 22, 2022.
- Wesley Holliday. Epistemic close and epistemic logic I: Relevant alternatives and subjunctivism. *Journal of Philosophical Logic*, 44:1–62, 2015.

- Wesley Holliday. Simplifying the surprise exam. *UC Berkeley Working Paper in Philosophy*, 2016.
- Joachim Horvath and Jennifer Nado. Knowledge and normality. *Synthese*, 198: 11673–94, 2021.
- David Hume. *A Treatise of Human Nature*. Oxford UP, 1739.
- E. T. Jaynes. *Probability Theory: the Logic of Science*. Cambridge University Press, 2003.
- Christoph Kelp. Knowledge and safety. *Journal of Philosophical Research*, 34: 21–31, 2009.
- Angelika Kratzer. Situations in natural language semantics. In *The Stanford Encyclopedia of Philosophy (Summer 2019 Edition)*. 2019.
- Sarit Kraus, Daniel Lehmann, and Menachem Magidor. Nonmonotonic reasoning, preferential models, and cumulative logics. *Artificial Intelligence*, 44: 167–207, 1990.
- Saul Kripke. On two paradoxes of knowledge. In *Philosophical Troubles*. Oxford UP, 2011.
- John Kruschke. *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan*. Academic Press/Elsevier, 2014.
- Hannes Leitgeb. The stability theory of belief. *Philosophical Review*, 123:131–171, 2014.
- Wolfgang Lenzen. *Recent Work in Epistemic Logic*. Acta Philosophica Fennica, 1978.
- Jarrett Leplin. In defense of reliabilism. *Philosophical Studies*, 134:31–42, 2007.
- David Lewis. *Counterfactuals*. Blackwell, 1973.
- David Lewis. Attitudes *de dicto* and *de se*. *Philosophical Review*, 88:513–543, 1979.
- David Lewis. Ordering semantics and premise semantics for counterfactuals. *Journal of Philosophical Logic*, 10:217–234, 1981.
- David Lewis. Elusive knowledge. *Australasian Journal of Philosophy*, 74:549–567, 1996.
- Clayton Littlejohn. Evidence and knowledge. *Erkenntnis*, 74:241–262, 2011.
- Clayton Littlejohn and Julien Dutant. Justification, knowledge, and normality. *Philosophical Studies*, 177:1593–1609, 2020.

- Annina Loets. Choice points for a theory of normality. *Mind*, 131(521):159–191, 2022.
- Ofra Magidor. How both you and the brain in a vat can know whether or not you are envatted. *Aristotelian Society Supplementary Volume*, 92:151–181, 2018.
- David Makinson. The paradox of the preface. *Analysis*, 25:205–207, 1965.
- David Makinson. Five faces of minimality. *Studia Logica*, 52:339–379, 1993.
- Kevin McCain. Evidentialism, explanationism, and beliefs about the future. *Erkenntnis*, 79:99–109, 2014.
- Ram Neta and Guy Rohrbaugh. Luminosity and the safety of knowledge. *Pacific Philosophical Quarterly*, 85:396–406, 2004.
- Andrew Peet and Eli Pitcovski. Normal knowledge: Towards an explanation-based theory of knowledge. *Journal of Philosophy*, 115:141–157, 2018.
- Duncan Pritchard. Anti-luck virtue epistemology. *Journal of Philosophy*, 109:247–279, 2012.
- James Pryor. The sceptic and the dogmatist. *Noûs*, 34:517–549, 2000.
- Colin Radford. Knowledge: by examples. *Analysis*, 27:1–11, 1966.
- Sven Rosenkranz. The structure of justification. *Mind*, 127:309–338, 2018.
- Hans Rott. Stability, strength and sensitivity: Converting belief into knowledge. *Erkenntnis*, 61:469–493, 2004.
- Bertrand Russell. *Human Knowledge: Its Scope and its Limits*. Allen & Unwin, 1949.
- Bernhard Salow. Lewis on iterated knowledge. *Philosophical Studies*, 173:1571–1590, 2016.
- Martin Smith. What else justification could be. *Noûs*, 44:10–31, 2010.
- Martin Smith. *Between Probability and Certainty: What Justifies Belief*. Oxford UP, 2016.
- Martin Smith. Why throwing 92 heads in a row is not surprising. *Philosophers’ Imprint*, 17:1–8, 2017.
- Martin Smith. Coin trials. *Canadian Journal of Philosophy*, 48:726–741, 2018a.
- Martin Smith. The logic of epistemic justification. *Synthese*, 195:3857–3875, 2018b.
- Ernest Sosa. How to defeat opposition to Moore. *Philosophical Perspectives*, 13:141–154, 1999.

- Robert Stalnaker. On logics of knowledge and belief. *Philosophical Studies*, 128: 169–99, 2006.
- Robert Stalnaker. Luminosity and the KK thesis. In S. Goldberg, editor, *Externalism, Self-Knowledge, and Skepticism*. Cambridge UP, 2015.
- Robert Stalnaker. Contextualism and the logic of knowledge. In *Knowledge and Conditionals*. Oxford UP, 2019.
- Markos Valaris. Normality, safety and knowledge. *Philosophy and Phenomenological Research*, forthcoming.
- Brian Weatherson. Margins and errors. *Inquiry*, 56:63–76, 2013.
- Brian Weatherson. Probability and scepticism. In Dylan Dodd and Elia Zardini, editors, *Scepticism and Perceptual Justification*. Oxford University Press, 2014.
- Roger White. Problems for dogmatism. *Philosophical Studies*, 131:525–557, 2006.
- Timothy Williamson. *Knowledge and its Limits*. Oxford UP, 2000.
- Timothy Williamson. Probability and danger. *The Amherst Lecture in Philosophy*, 4:1–35, 2009.
- Timothy Williamson. Improbable knowing. In Trent Dougherty, editor, *Evidentialism and its Discontents*. Oxford UP, Oxford, 2011.
- Timothy Williamson. Gettier cases in epistemic logic. *Inquiry*, 56:1–14, 2013a.
- Timothy Williamson. Response to Cohen, Comesaña, Goodman, Nagel, and Weatherson on Gettier cases in epistemic logic. *Inquiry*, 56:77–96, 2013b.
- Timothy Williamson. Very improbable knowing. *Erkenntnis*, 79:971–999, 2014.
- Timothy Williamson. A note on Gettier cases in epistemic logic. *Philosophical Studies*, 172:129–140, 2015.
- Timothy Williamson. The KK principle and rotational symmetry. *Analytic Philosophy*, 62:107–24, 2021.
- Zhaoqing Xu and Bo Chen. Epistemic logic with evidence and relevant alternatives. In Hans van Ditmarsch and Gabriel Sandu, editors, *Jaakko Hintikka on Knowledge and Game-Theoretical Semantics*. Springer, 2018.
- Elia Zardini. $K \not\subseteq E$. *Philosophy and Phenomenological Research*, 94:540–557, 2017.