Getting Accurate about Knowledge

Sam Carter & Simon Goldstein

Forthcoming, *Mind*

Abstract

There is a large literature exploring how accuracy constrains rational degrees of belief. This paper turns to the unexplored question of how accuracy constrains knowledge. We begin by introducing a simple hypothesis: increases in the accuracy of an agent's evidence never lead to decreases in what the agent knows. We explore various precise formulations of this principle, consider arguments in its favor, and explain how it interacts with different conceptions of evidence and accuracy. As we show, the principle has some noteworthy consequences for the wider theory of knowledge. First, it implies that an agent cannot be justified in believing a set of mutually inconsistent claims. Second, it implies the existence of a kind of epistemic blindspot: it is not possible to know that one's evidence is misleading.

1   Accuracy

Not all evidence is equally accurate. Weather reports, weighing scales, and world maps can all vary in how faithfully they represent reality. Where they do, the accuracy of evidence obtained from them will vary correspondingly.

How is accuracy related to knowledge? The last decades have witnessed an explosion of research using accuracy to articulate a theory of rational degrees of belief.[1] Likewise, a common approach to theorizing about perception assigns accuracy a central role in characterizing the contents of perceptual experience.[2] However, there is surprisingly little work explaining how accuracy interacts with what an agent can know. This paper fills that gap.

It is natural to think that what you can know is partly a matter of how accurate your evidence is. Other things being equal, increasing accuracy increases what you can know.

   ACCURACY   The more accurate your evidence, the more you can know.

Imagine you are handed a map of a region you've never visited. The closer the region's geography conforms to the map, the more accurate the map is. An appealing idea is that increases in the map's accuracy cannot introduce barriers to knowledge. However accurate the map is, you don't know anything from the map that you could not know if the region's geography better fit the map.

---

[1]See Joyce 1998 for an initial exposition, and Leitgeb and Pettigrew 2010*a*,*b* and Pettigrew (2016) for significant recent contributions.

[2]See, e.g., Siewert (1998); Chalmers (2006); Siegel (2010).

Or take another case. Imagine you are given a class register which records each student as present or absent. The register contains an error wherever an absent student is marked present, or a present student absent. However many errors there are in the register, you don't get to know anything about class attendance that you could not know if some errors had been corrected.

ACCURACY articulates this idea by positing a connection between how closely the world conforms to your evidence and how much you get to know. In this way, it constitutes a substantive proposal about the relationship between evidence and knowledge. We investigate this proposal in what follows. Our primary goal is exploratory, not polemical. Rather than defending a specific position on our various formulations of ACCURACY, we instead clarify their consequences and show how they interact with various theories of evidence, accuracy, and knowledge. In doing so, we hope to highlight some surprising implications of what may, initially, appear a relatively innocuous principle. This project can be understood as identifying and critically assessing ideas tacit in a significant strand of recent work in epistemology, including Williamson (2013); Goodman (2013); Goodman and Salow (2018, manuscript), and Littlejohn and Dutant (2020). This work has generally limited its attention to knowledge in specific domains (e.g., perceptual knowledge, inductive knowledge, etc.). Our aim, in contrast, is to assess the tenability of endorsing ACCURACY generally.

The paper proceeds as follows: §2 explores what exactly ACCURACY says, developing different formulations of the principle which vary in strength. §3 presents three arguments in favor of ACCURACY, involving (i) anti-skepticism; (ii) normality; and (iii) Gettierization. §4 considers a variety of substantive conceptions of when evidence is accurate, drawing on existing literatures on truthlikeness and rational credence. Finally, §5 addresses the consequences of ACCURACY for justification and knowledge. First, ACCURACY suggests that justified beliefs must be mutually consistent. Second, ACCURACY requires the existence of epistemic blindspots regarding one's own inaccuracy, ruling out the possibility of knowing that one's evidence is misleading.

2   Accuracy principles

ACCURACY relates accuracy, evidence, and knowledge. To see what it says, we need a bit more precision. Given a domain of worlds, W, we represent the epistemic position of a designated agent in terms of a family of operations on W. We identify propositions with sets of worlds and assume that for any set of propositions, P, its intersection, $\bigcap P$, is also a proposition. To represent what is epistemically accessible to the agent, we let K be a function which maps each world to the set of propositions the agent is in a position to know there. That is, $p \in K(w)$ iff the agent is in a position to know p at w.[3] To represent the agent's evidential state, we let E be a partition of W according to the agent's

[3]Given concerns of the kind raised in Heylen (2016) and Hawthorne (Forthcoming) we do not assume that what an agent is in a position to know is closed (that is, it is not required

evidence. That is, $E(w)$ is the set of worlds at which the agent's evidence is the same as at w. To model accuracy (the gradable property, not the principle) we introduce a relation, $\leq$, of comparative accuracy over worlds. $w \leq v$ iff the agent's evidence is at least as accurate at v as it is at w. We assume that $\leq$ is reflexive and transitive (that is, it is a pre-order).

There are a variety of ways an accuracy ordering might be determined. In §4, we consider two broad families of approaches. On the first, an agent's evidence is represented by a set of propositions. How accurate the evidence is at a world is a matter of how well the set of propositions approximate the way things are at that world. On the second, in contrast, an agent's evidence is represented by a probability measure. How accurate the evidence is at a world is a matter of how close that measure comes to correctly predicting how things are at the world.

Considering concrete approaches to accuracy can help us to theorize about its formal features. Different approaches will yield accuracy orders with different structural properties. However, we do not intend to treat accuracy as a purely theoretical posit. We have a pre-theoretic grasp on the conditions under which, e.g., a register or map will be more or less accurate. The success of different approaches to accuracy will ultimately depend on how well they capture our pre-theoretic understanding of accuracy. And this pre-theoretic understanding is, we want to suggest, sufficient to give us at least a provisional grasp of what different precifications of ACCURACY entail.

We will interpret ACCURACY as silent about differences in knowledge between worlds with different evidence. There is no obvious generalization to be made regarding the difference in what you can know from a more detailed but less accurate map and what you can know from a less detailed but more accurate map of the same region. Accordingly, our interpretation of ACCURACY compares only what is known across worlds with the same evidence.

We now introduce our first precisification of ACCURACY. STRONG ACCURACY says that if w and v have the same evidence and the evidence is at least as accurate at v as at w, then anything the agent can know at w she can know at v.

   STRONG ACCURACY   For all $v \in E(w)$, if $w \leq v$, then $K(w) \subseteq K(v)$.

It is important to note that STRONG ACCURACY concerns not what an agent in fact knows, but what she is in a position to know. Knowledge depends not only on evidence but also on belief. An agent may know less at v than w simply because she believes less at v. Accordingly, if it is to be at all plausible, it crucial that STRONG ACCURACY is framed in terms of an epistemic state which does not entail belief.

Even taking this into account, a natural objection to ACCURACY is that it still makes what an agent can know overly dependent upon evidence. In particular,

---

that, for $X \subseteq K(w), \bigcap X \in K(w)$). However, we do assume that an agent can be a position to know p only if p is true (that is, it is $w \in \bigcap K(w)$).

it implies that in any pair of cases which agree on how accurate one's evidence is, one is in a position to know exactly the same propositions on the basis of that evidence. Yet can't other factors, such as luck, influence what you can know?

One response to the latter problem is to treat STRONG ACCURACY as a thesis about knowledge in an idealized setting, one in which such factors are presumed to be absent. Even if it holds only under special conditions, investigating consequences of the strong version of the principle might nevertheless teach us something about the relationship between evidence and knowledge more generally.

Alternatively, we can weaken STRONG ACCURACY to allow for the possibility of factors besides accuracy affecting how much one knows. WEAK ACCURACY says that for any world, w, it is possible to find a world v with the same evidence which is at least as accurate as w, and where at least as much can be known as at any world z no more accurate than w.

> WEAK ACCURACY   There is some $v \in E(w)$ such that $w \leq v$ and for all $z \in E(w)$, if $z \leq w$, then $K(z) \subseteq K(v)$.

WEAK ACCURACY is weaker than STRONG ACCURACY. It says that increases in accuracy are compatible with knowing everything that could be known if the evidence was less accurate. In this way, WEAK ACCURACY says that accuracy is not a barrier to knowledge.

It is also worth considering even weaker versions of ACCURACY. VERY WEAK ACCURACY reverses the scope of quantifiers in WEAK ACCURACY:

> VERY WEAK   For all $z \in E(w)$, if $z \leq w$, then there is some $v \in E(w)$
> ACCURACY   such that $w \leq v$ and $K(z) \subseteq K(v)$.

VERY WEAK ACCURACY is weaker than WEAK ACCURACY. It requires only that for each world z less accurate than w, it is possible to find a world v at least as accurate as w where at least as much can be known. VERY WEAK ACCURACY is very weak. It is compatible with there being two worlds, v_1 and v_2, which are less accurate than w, even though there is no world as accurate as w where you know at least as much as you do in both v_1 and v_2.

Throughout, we use 'ACCURACY' to refer to the informal principle stated in §1, understood as the disjunction of the precise principles above. We refer to ACCURACY when the differences between each formulation are irrelevant.

ACCURACY's implications depend on the structure of comparative accuracy. We focus on two structural properties in particular. CONNECTEDNESS says that any two worlds with the same evidence can be compared for accuracy (so that $\leq$ is a total pre-order over each cell of E). DIRECTEDNESS is weaker than CONNECTEDNESS. It says that for any two worlds with the same evidence, there is a world with the same evidence at which that evidence is at least as accurate as at both.

CONNECTEDNESS   If $v \in E(w)$ then either $v \leq w$ or $w \leq v$.

DIRECTEDNESS   If $v \in E(w)$ then there is some $z \in E(w)$ such that $v \leq z$ and $w \leq z$.

Later, we'll see that different conceptions of accuracy disagree about what structural properties it has. And, in turn, these differences have important downstream implications for what accuracy predicts ACCURACY about knowledge and justification.

We now consider three arguments for ACCURACY.

3   Arguments for Accuracy

A number of existing theories take something like ACCURACY for granted (at least when restricted to knowledge in specific domains). Recent work on perceptual knowledge adopts it as a starting point in theorizing about inexactness (Stalnaker (2006, 2015); Williamson (2013); Goodman (2013); Cohen and Comesaña (2013)). And it is also tacit in work relating normality and knowledge (Greco (2014); Goodman and Salow (2018); Beddor and Pavese (2018); Loets (forthcoming)).[4] While ACCURACY is sometimes assumed implicitly, it is rarely defended explicitly. This section develops three arguments for it. None are conclusive, but each reveals something interesting about the principle.

3.1   Anti-Skepticism

A preliminary argument for ACCURACY is that it systematizes our judgments about the skeptical predicament and extends these judgments to various levels of skeptical threat.

Anti-skepticism says that there is an asymmetry between the good case, where evidence is accurate, and the bad case, where it isn't. Despite being indiscriminable by appearance, some propositions unknowable in the bad case are knowable in the good case. The difference in accuracy across the two cases generates a difference in what can be known.

Not all bad cases are equally bad, however; there is an epistemic difference between a slow clock and a stopped clock. The basic anti-skeptical position is silent about this difference. It says nothing about how knowledge varies between mediocre and catastrophic cases. It is also silent on how much can be known in bad cases in general.

ACCURACY fills this gap, supplementing the basic anti-skeptical position in two ways. First, both strong and weak formulations of ACCURACY strengthen

---

[4]In recent work, Carter and Goldstein (2021) and Goodman and Salow (manuscript) draw connections between these literatures and question their shared presuppositions.

anti-skepticism by clarifying the epistemic difference between the good and bad cases. Anti-Skepticism simply says that there is something you cannot know in the bad case which you can know in the good case. But this allows that the two cases may be incomparable, in that there are things known in the bad case which are unknown in the good case. Our various accuracy principles constrain this relationship in different ways. STRONG ACCURACY adds to anti-skepticism that one is in at least as good an epistemic position in the good case as in the bad case. Anything which is knowable in the latter is knowable in the former. WEAK ACCURACY imposes the weaker requirement that there is some good case in which one is in a strictly better position than every bad case. And VERY WEAK ACCURACY implies that there is nothing that can be known in a bad case which cannot be known in some good case.

Second, depending on the strength of precisification adopted, ACCURACY can extend anti-skepticism to different gradations of bad case. STRONG ACCURACY says that the magnitude of inaccuracy constrains the magnitude of ignorance. One cannot know more in a catastrophic case than in a mediocre case. In contrast, in the mediocre case, there may be propositions knowable which cannot be known in the catastrophe. WEAK ACCURACY requires that there is some case in which you can know at least as much as in both the mediocre case and the catastrophe. However, it does not impose any constraint directly on the difference between what can be known at each. Finally, VERY WEAK ACCURACY says nothing about the relationship between the mediocre and catastrophic cases at all.

3.2   Independence

A second argument for ACCURACY relies on the premise that accuracy is in a certain sense independent from other factors relevant to knowledge. STRONG ACCURACY says that inaccuracy is a barrier to knowledge. You can't know p at w unless p can be known at any world where your evidence is at least accurate. The primary objection to STRONG ACCURACY above was that it is not the only barrier to knowing. Other factors may intervene to block knowledge without implying a decrease in accuracy. Where such a factor is absent at w yet present at a more accurate v, p may be unknown at v despite being known at w.

One kind of barrier is lack of justification: to be in a position to know p, p must be appropriately justified by your evidence. Another kind of barrier is falsity: regardless of how strong your evidence, you are not in a position to know p if p is false. And even among propositions which are justified and true, gettierization is a third kind of barrier: being in a position to know p requires more than a lucky connection between your evidence and the world.

VERY WEAK ACCURACY follows from the idea that any barriers to knowledge beyond accuracy are independent of accuracy. That is, if the absence of barriers to knowing some propositions is compatible with the evidence, then the absence of those barriers is compatible with the evidence having any higher level of

accuracy.

Why accept this? First, consider justification. On the assumption that justification (of the kind required for being in a position to know) supervenes on evidence, if a set of propositions is justified in some world in $E(w)$, it is justified in every world in $E(w)$. So trivially, that set of propositions being justified will be compatible with any level of accuracy across $E(w)$. Next, consider truth. It is plausible that there may be some propositions whose truth sets a limit to the evidence's accuracy. In particular, where the evidence justifies $\neg p$, it is plausible that p's truth will be a contributing factor to its inaccuracy. However, where p is justified by the evidence, it is hard to see how p's truth could require that evidence to possess a certain level of inaccuracy.[5]

Finally, consider gettierization. In many standard Gettier cases, gettierization can vary without any change in accuracy. The well-functioning clock and the stopped clock showing the same time are equally accurate. Not all Gettier cases are like this—sometimes, inaccuracy may be partially constitutive of gettierization. For example, Williamson (2013) shows that, in the presence of margin for error constraints, inaccuracy will sometimes give rise to gettierization by itself. Yet there are no cases in which the converse holds.

Suppose that the absence of justification, falsity and gettierization are independent of accuracy in the sense above.[6] Then as long as these, along with inaccuracy, exhaust the barriers to knowledge, VERY WEAK ACCURACY follows.

VERY WEAK ACCURACY says that for any $v \in E(w)$ such that $v \leq w$, there is some $z \in E(w)$ at which at least as much can be known. So consider the set of propositions in $K(v)$. Each of these propositions is justified, true and ungettierized at v. So, by the assumption that those barriers to knowledge are independent of accuracy, there is a world at least as accurate as w at which all of these propositions are justified, true and ungettierized. Call this world z. By the transitivity of accuracy, we know that z is at least as accurate as v. So, as long as these three exhaust the barriers to knowledge, any proposition which can be known at v can be known at z.

---

[5] A potential class of counter-examples involve cases in which the evidence justifies the proposition that some proposition justified by the evidence is false. One case of this kind is the preface (Makinson (1965)). As we will discuss later (§5), preface-like cases present a challenge to ACCURACY. We will take seriously the possibility that problems concerning the preface provide a reason to give up (some versions of) ACCURACY. However, we will also see that there are a number of ways of resisting this type of counter-example. One option we discuss is to deny DIRECTEDNESS. If DIRECTEDNESS is denied, we can maintain that, even in preface-like cases, for any justified proposition there is some maximally accurate world at which that proposition is true. We simply allow that maximally accurate worlds may be incomparable.

[6] What other barriers to knowledge could there be? One possibility is that justified true lottery propositions (propositions that, for some ticket in a fair lottery, that ticket will lose) are unknowable despite being ungettierized (cf. Hawthorne (2003, 9), Pritchard (2008, 4)). Recognizing an additional barrier which prevents knowledge of lottery propositions would not present a significant problem for the present argument; holding evidence fixed, lottery propositions appear unknowable independent of accuracy. However, absent an exhaustive list of barriers, the present argument will have to remain less than fully decisive.

The argument above falls short of establishing WEAK ACCURACY. To extend the argument to WEAK ACCURACY, we need a stronger form of independence: if the absence of barriers to knowing the propositions in X_i is compatible with possessing some evidence, for each X_i in a series $X_1, ..., X_n$, then the absence of barriers to knowing propositions in $\bigcup_{i=1}^{n} X_i$ is compatible with that evidence being arbitrarily accurate. While there may not be obvious counter-examples to this principle, the considerations discussed above fall short of supporting it.

3.3   Normality

Our last argument for ACCURACY appeals to the connection between knowledge, normality, and accuracy. A growing body of work appeals to normality in stating conditions on knowledge (Greco (2014); Stalnaker (2015); Goodman and Salow (2018, manuscript); Beddor and Pavese (2018); Carter (2019); Littlejohn and Dutant (2020); Carter and Goldstein (2021); Goldstein and Hawthorne (forthcoming); Loets (forthcoming)). Its core idea is that worlds can be compared according to the normality of the agent's epistemic situation. While details differ, each of these accounts is committed to the following necessary condition on knowledge:

KNOWLEDGE REQUIRES NORMALITY
p can be known at w only if p can be known at any world at least as normal as w.

Inaccuracy makes an epistemic situation abnormal. Some knowledge-theoretic accounts propose to identify normality with the accuracy of evidence (Goodman and Salow (manuscript), cf. Carter and Goldstein (2021); in some places, this identification may restricted, so that it applies only to the accuracy of, e.g., perceptual evidence). Given this strong assumption, KNOWLEDGE REQUIRES NORMALITY immediately entails STRONG ACCURACY.

However, inaccuracy may not be the only contributing factor to abnormality. Inebriation, insomnia, and insanity can also make your epistemic situation less normal, without changing the accuracy of your evidence. Rather than identify normality with accuracy, a more ecumenical approach takes accuracy to constrain normality. The following necessary condition offers one way of articulating this idea:

NORMALITY REQUIRES ACCURACY
v is at least as normal as w only if the evidence at v is at least as accurate as at w. [7]

---

[7]Note that NORMALITY REQUIRES ACCURACY is is a necessary but not sufficient condition for one world being at least as normal as another. This is crucial. Gettierization need not be accompanied by inaccuracy. If the converse were also endorsed, this would imply that some worlds in which gettierization is present are as normal as worlds in which it is absent. Yet, in combination with KNOWLEDGE REQUIRES NORMALITY this would impose implausible limitations on what could be known at worlds in which gettierization is absent.

Knowledge Requires Normality and Normality Requires Accuracy imply Very Weak Accuracy, as long as it is assumed that the normality ordering is directed among worlds with the same evidence. For suppose that $z \in E(w)$ and $z \leq w$. While w and z need not be comparable for normality, there will be some $v \in E(w)$ such that z is at least as normal as both. By Normality Requires Accuracy, it follows that $z \leq v$. Knowledge Requires Normality then guarantees that $K(z) \subseteq K(v)$, which is what Very Weak Accuracy requires.

We've considered a few arguments for Accuracy. With some *prima facie* motivation for the principle in place, we now develop several theories of what it takes for the evidence to be accurate.

4 Theories of Accuracy

What Accuracy says depends on how we measure accuracy. This section surveys a number of approaches and shows how they interact with theories of evidence. These approaches divide into two kinds: propositional and probabilistic. The former represent evidence propositionally: the better a set of propositions approximate the way things are, the more accurate the evidence they represent. The latter represent evidence probabilistically: the smaller the distance between the evidential probabilities and the way things are, the more accurate the evidence.

4.1 Propositional Theories

Propositional theories measure the accuracy of evidence by the degree to which a set of relevant propositions approximate the truth.[8] This idea has three components: first, that propositions approximate the truth to varying degrees. Second, that the aspects of evidence relevant to accuracy can be represented by a set of propositions. And third, that the degree to which a set of propositions approximate the truth depends on how closely each proposition approximates the truth. We consider each in turn.

The literature on truthlikeness (Tichý 1976; Hilpinen 1976; Oddie 1986; Niiniluoto 1987, 1998, 2020) offers a framework for thinking about approximation. Its guiding idea is that how far a proposition p is from accurately characterizing a world w can be measured in terms of a real value, $d(p, w)$. It is generally assumed that d has metric structure, so that degrees of approximation behave like distances. For simplicity, we will take its range to be a finite closed interval, $[0, n_{max}]$. The smaller the distance between a proposition and world, the more

---

[8] It is crucial to note, here, that we use 'accuracy' as a term for the pre-theoretic notion formally represented by our $\leq$-ordering, rather than in the technical sense in which it is sometimes employed in the literature on rational degrees of belief (as in, e.g., Joyce (1998); Leitgeb and Pettigrew (2010b); Pettigrew (2016)). As Oddie (2013) has demonstrated, the latter notion is importantly distinct from the notion of truthlikeness/approximation.

closely the former approximates the latter; so p approximates w at least as well as q approximates v iff $d(p, w) \leq d(q, v)$. We further assume that all and only the true propositions at a world perfectly approximate the truth at that world; so $d(p, w) = 0$ iff p is true at w.[9],[10]

This leaves open many theories of approximation. The simplest theory says that every falsehood is maximally inaccurate, so that $d(p, w) = n_{max}$ iff p is false at w. To see this theory in action, return to the class register from above. Where p is relevant to the accuracy of some evidence, we will say that p represents that evidence. The content of the register could be represented by a the members of a set comprising, for each student, either the proposition that they were present or that they were absent. The simple theory says that the extent to which a proposition in this set approximates the truth depends entirely on its truth value. This seems appropriate—to evaluate the accuracy of the register, we don't need to know anything other than which students were present and which were absent.

While the simple theory of approximation provides an adequate treatment of the class register, it does less well in more sophisticated cases. Consider a thermometer which reports the temperature in a room to be 50°F. We can represent the content of the thermometer's report with the proposition true iff the temperature in the room is in fact 50°F. Yet, the extent to which this proposition approximates the truth depends on more than its truth value. Intuitively, the thermometer's report is a better approximation of a world at which the temperature is 55° than a world at which the temperature is 60°. In order to capture this observation, however, we'll need a measure of approximation on which there can be differences in how well false propositions approximate a world.

This observation is not dependent on the quantitative structure of the case. Intuitively, our original example of the map is a case in which the extent to which a report approximates the truth can vary across worlds at which it is less than wholly accurate. Generalizing, as Hilpinen (1976) observes, we can think of degrees of approximation in terms of a metric similarity structure over worlds (cf. Lewis (1973); Spohn (2012)). For example, a natural proposal which meets our requirements is that the degree to which a proposition approximates the truth at w is proportional to the similarity between w and the most similar worlds at which it is true.

With some grip on propositional approximation, we turn to the question of which propositions represent an agent's evidence. Our guiding idea is that the overall accuracy of an agent's evidence is some function of how closely each of the propositions representing it approximates the truth. We think there are

---

[9]This property is a feature of many accounts of truthlikeness, including for example the minimality measure (Weston 1992, Teller 2001).

[10]We are not the first to propose understanding accuracy in terms of approximate truth. Williamson (2013)'s models of inexact knowledge can be understood as implementing a similar idea.

at least two good options here. First, the propositions representing an agent's evidence could be those which comprise part of it. Second, the propositions representing an agent's evidence could be those which are supported by it.

Propositional theories of evidence identify an agent's evidence with a set of propositions. On non-factive propositional theories, an agent's evidence may contain falsehoods (Schroeder 2008; Goldman 2009; Fantl and McGrath 2009; Rizzieri 2011; Arnold 2013). Accordingly, for proponents of such theories, it is natural to take the accuracy of one's evidence to match the degree to which the propositions it comprises approximates the truth.

This approach doesn't make sense for factive and non-propositional theories of evidence. On factive propositional theories, a proposition is part of an agent's evidence only if it is true (Williamson 2000; Bird 2004; Hyman 2006; Littlejohn 2012, 2013). But, as we proposed above, every truth approximates the truth perfectly. So the accuracy of the evidence cannot in this setting be measured by how well the evidence itself approximates the truth. On non-propositional theories of evidence, an agent's evidence can be comprised (either wholly or partially) of items which are not themselves capable of truth or falsehood (Pollock 1974; Moser 1989; Pollock and Gillies 2000; Davidson 2001; Huemer 2006; Conee and Feldman 2008). Yet it is unclear how to apply the notion of approximating truth to things which are not themselves capable of truth or falsehood.

Instead, in factive and non-propositional frameworks, the accuracy of evidence can be measured by considering some set of propositions which it supports. As long as evidence may support a proposition without guaranteeing its truth, this allows avoids the problems above.[11]

A simple option would be to identify the relevant relation of support with justification. It is not clear, however, that all propositions justified by an agent's evidence are relevant to its accuracy. To see the point, consider an urn which, for all you know antecedently, contains between 0% and 100% red balls (with the remainder black). Suppose 100 balls are drawn from an urn with replacement. If the first 100 draws each appeared red, then for any $1 \leq n \leq 100$, the proposition that the nth ball drawn was red is justified by your evidence. Plausibly, though, this is not all that your evidence justifies: you also gain justification about future draws. For example, the proposition that the 101st ball drawn will be red is also justified by your evidence. There is, however, a potential asymmetry between

---

[11] What propositions are relevant to the accuracy of some evidence can depend on global features of that evidence. For example, suppose that an agent knows her watch runs 15 minutes fast. Then, where the watch displays 6pm, the proposition relevant to the accuracy of her evidence (taken in its entirety) will be the proposition that the time is 6.15pm (not the proposition that it is 6pm). When dealing with cases involving known biases, it is important to distinguish the accuracy of a source of evidence from the accuracy of the evidence it produces.

The two can also come apart in cases where a source is believed to be inaccurate (though not in virtue of any regular bias). For instance, if an agent has evidence that a map only imprecisely represents the geography of a region, the accuracy of the evidence they acquire from consulting it may differ from the accuracy of the map itself. We are grateful to both referees for *Mind* on this point.

how these propositions bear on the accuracy of your evidence. In a world in which the 100th ball drawn was black, your evidence will be less accurate than it is in a world in which all 100 balls drawn were red. By contrast, it is unclear that in a world where the 101st ball drawn is black your evidence is any less accurate than it would be in a world in which the 101st ball drawn was red. One diagnosis of this asymmetry would be to ascribe it to a difference in the kind of justification the propositions possess. Perhaps non-inductively justified propositions are relevant to evidential accuracy in a way inductively justified propositions are not. If that is right, a natural move is to restrict the relevant notion of support to a sub-set of the propositions justified by evidence, perhaps those which are directly or immediately justified.

On the other hand, it is natural to say that one's evidence is highly misleading if as a matter of fact most balls are black and most draws after 100 will be black, despite the first 100 balls appearing red. Generalizing, the falsity of inductively justified claims affects the extent to which the evidence is misleading. Insofar as the inaccuracy of evidence is simply a matter of the evidence being misleading, then, it seems that any proposition justified by the evidence is potentially relevant.[12]

In what follows we let E_w be the set of propositions representing the evidence at w. This allows us to abstract away questions about the precise relation between evidence and the propositions it is represented by. We assume that whenever w and v have the same evidence, $\mathsf{E}_w = \mathsf{E}_v$. If propositions which represent evidence are those supported by—rather than part of—it, this amounts to the evidentialist assumption that any two worlds with the same evidence support the same propositions. We suppose that evidential accuracy at w depends exclusively on E_w. In these terms, the propositional theory of accuracy says that the closer E_w approximates the world, the more you get to know.

We now turn to our final question. How does the degree to which a set of propositions approximates the truth depend on how each of its members individually approximates the truth? Different answers to this question produce accuracy orderings with different structural properties. For instance, we'll see that it bears directly on the question of whether the ordering is connected or directed.

Suppose that, for any w, E_w is consistent. Then the simplest option is to measure accuracy in terms of how closely its closure, $\bigcap \mathsf{E}_w$, approximates the truth.

CONJUNCTION   If w and v have the same evidence, then $w \leq v$ iff $d(\bigcap \mathsf{E}_w, w) \geq d(\bigcap \mathsf{E}_w, v)$.

---

[12]Advocates of STRONG ACCURACY have a further reason to take evidence to be represented by propositions it inductively justifies. Given STRONG ACCURACY, if p is not entailed by the propositions representing the evidence at w, then p cannot be known at any $v \in E(w)$. Accordingly, on pain of inductive skepticism, (the closure of) the set of propositions representing an agent's evidence must be at least as strong as (some non-empty subset of) the propositions it inductively justifies.

For any set of propositions E_w, and any v with the same evidence as w, it is possible to compare how closely the conjunction $\bigcap \mathsf{E}_w$ approximates the truth at w and v. So, given CONJUNCTION, the accuracy order will be connected over any set of worlds with the same evidence.

If the propositions representing an agent's evidence can be jointly inconsistent, CONJUNCTION is inappropriate. The accuracy of one's evidence will just amount to the distance between the contradiction and a world. It is plausible that, at every world, the contradiction is maximally far away from approximating the truth. But, intuitively, the accuracy of a body of evidence may vary even if the propositions it comprises or justifies are not consistent.

Even if the relevant propositions are assumed to be consistent, CONJUNCTION may still be inappropriate. The extent to which a set of propositions approximates the truth can depend, not only on what it entails, but also on how it is structured. Consider two textbooks. Textbook A contains ninety-nine true claims, $t_1, ..., t_{99}$ and one false claim, f. Textbook B contains the false claim, f, plus, for each t_i, the claim that t_i is materially equivalent to f. Plausibly, Textbook A approximates the truth more closely than Textbook B. Yet the sets of propositions recorded in each will have the same closure.

To avoid these challenges, one option is to sum the accuracy of the relevant propositions (cf. Tichý (1974) Tichý (1976), Oddie (1986), and Oddie (2013)).

SUMMATION   If w and v have the same evidence, then $w \leq v$ iff $\sum_{p \in \mathsf{E}_w} d(p, w) \geq \sum_{p \in \mathsf{E}_w} d(p, v)$.

SUMMATION implies that the accuracy order is connected across worlds with the same evidence. One immediate worry is that different propositions may contribute to the accuracy of an agent's evidence to different degrees (Joyce (2005)). Put another way, among the set of propositions relevant to the accuracy of an agent's evidence, some propositions may be more relevant than others. A simple solution is for the proponent of SUMMATION to introduce a weighting function. Under this proposal, in calculating the accuracy of evidence, the distance between each proposition and the truth is multiplied by a factor representing its relative relevance (Oddie (1986)).[13]

A deeper issue arises from cases where two propositions address radically different subject matters. Here, it may be hard to compare the extent to which they contribute the overall accuracy of evidence which both represent. One response would be to take such cases to involve incomparabilities. Perhaps the contributions of my childhood memories and my present experiences to the accuracy of my total evidence are incommensurable. Yet this response is unavailable to the proponent of SUMMATION.[14] While they can assign different weights to different propositions, they are committed to holding that the weights of different

---

[13]See Dorst 2019 and Easwaran 2016 for use of SUMMATION to measure the accuracy of belief sets.

[14]This is equally an issue for the proponent of CONJUNCTION; cf. footnote 16.

propositions can be added together and compared.[15]

To model incomparabilities, a third conception of approximation universally quantifies over relevant propositions: v is as accurate as w iff every relevant proposition approximates the truth at least as closely at v as at w.

SUPERVALUATION   If w and v have the same evidence, then $w \leq v$ iff for every p in E_w: $d(p, w) \geq d(p, v)$.

This proposal differs from both of the previous proposals in permitting failures of CONNECTEDNESS. Indeed, where the propositions representing the evidence are jointly inconsistent, DIRECTEDNESS will also fail. That is, for some pairs of worlds, there may be no world at which the evidence is at least as accurate as it is at each. In this way, the structural conditions on accuracy depend on the nature of approximation and the consistency of the propositions which represent the evidence.[16]

Our different proposals yield different predictions, even in cases where the propositions representing the evidence are consistent. Suppose, for a simple example, that E_w comprises, for each student marked present on the register at w, the proposition that that student was present, and for each student marked absent, the proposition that that student was absent. Let the distance between two worlds be measured by the number of students present at one but absent at the other. We can then get a simple account of how well a proposition approximates a world in terms of the distance from the latter to the nearest world at which the former is true.[17]

---

[15]Another problem arises if E_w can contain infinitely many propositions. Where the evidence is represented by infinitely many propositions, SUMMATION will not be appropriate as a way of inducing an ordering over worlds. For discussion of how to generalize Tichý (1974); Tichý (1976, 1978) and Oddie (1986, 2013)'s proposal to infinite domains, see Kieseppä (1996b,a).

[16]Intermediate positions between these proposals are also possible. Faced with apparent incomparabilities, the set of propositions representing the evidence could be partitioned into cells according to subject matter. One option, combining elements of CONJUNCTION and SUPERVALUATION, would be to take w to be at least as accurate as v (assuming $v \in E(w)$) iff the closure of each cell approximates w at least as well as it approximates v. Equally, one could respond by positing a set of weightings of propositions (rather than a single weighting). Another option, combining elements of SUMMATION and SUPERVALUATION, would then be to take w to be at least as accurate as v iff the weighted sum of how far the representing propositions are from approximating the truth is at least as great on every weighting for v as it is for w. While we think these intermediate positions have much going for them, they do not generate significantly different structural properties for the accuracy ordering. In particular, absent further constraints both will allow for failures of CONNECTEDNESS and DIRECTEDNESS.

[17]Where the register carries more information, the set propositions which represent it may need to be configured differently. Imagine, for example, that the students are divided equally into two groups. Suppose that the same number of students are marked absent from each group. As long as errors among students marked present and students marked absent are both distributed equally among both groups, it seems you should be able to know that the rate of absenteeism does not vary substantially between groups. Yet, at least assuming STRONG ACCURACY, it follows that the proposition that the same number of students are absent from each group will need to be among those representing the evidence provided by register.

This does not strike us as implausible. This proposition will, it seems, be either included in or directly justified by the evidence of an agent who consults the register. Accordingly, on the

14

Given this set up, CONJUNCTION and SUMMATION make the same prediction: where $v \in E(w)$, w is at least as accurate as v iff there are at least as many errors on the register at v as at w. SUPERVALUATION makes a different prediction: w is at least as accurate as v iff every error on the register at w is an error at v.[18]

Importantly, note that SUPERVALUATION (and SUMMATION) make accuracy sensitive to the structure of the set of propositions representing the evidence. Two sets with the same closure can generate distinct accuracy orders.[19] There is more than one way of this set in our example. One alternative, as suggested above, is to take, for each student, either the proposition that they were present or that they were absent (according to what the register records). A second alternative is to take a pair of propositions: the conjunction of the propositions, for each student marked present, that they were present, and the conjunction of the propositions, for each student marked absent, that they were absent.

Both carry, in some sense, all the information relevant to the accuracy of the register.[20] However, under SUPERVALUATION, the accuracy orders they produce will differ. Consider two worlds at which there is exactly one error on the register: at each, a different student marked present was absent. Given the first way of representing the evidence, SUPERVALUATION will classify the two worlds as incomparable. At each, some proposition representing the evidence is more accurate at the other. In contrast, given the second way of representing it, the two worlds will be classified as equally accurate.[21] Each of the two relevant propositions will have the same level of accuracy at each.[22]

---

various accounts of representation suggested above, there would be no obstacle to including it in E_w.

[18] What about a more complicated case in which, for example, the register also contains the information that it contains at least one error? A simple way of extending the model would be to introduce a second distance measure over worlds, on which the distance between two worlds is measured by the difference in the number of errors on the register at each. How well a proposition approximates a world can then be measured by the aggregate distance from the latter to the nearest world at which the former is true (where aggregate distance is just the unweighted sum of the two distance measures). Obviously, more complicated modifications are also possible.

[19] Notably, under SUPERVALUATION, the more propositions that represent a body of evidence, the fewer worlds will be comparable. In the limiting case, where the set of propositions is closed under single premise entailment, evidence will be maximally inaccurate at a world unless every proposition representing it is true at the world. To see why, suppose that for some $p \in \mathsf{E}_w$, $w \notin p$. Assume $v \in E(w)$. Since E_w is upward closed, $p \cup \{v\} \in \mathsf{E}_w$. But $d(p \cup \{v\}, v) > d(p \cup \{v\}, w)$. So $w \not\succeq v$. Since v was arbitrary, it follows that any world which shares the same evidence as w is either incomparable or strictly more accurate. We are grateful to a referee at *Mind* on this point.

[20] Whether this corresponds to the propositions someone who consults it has as evidence or the propositions their evidence supports.

[21] At least, assuming a simple metric over propositions.

[22] Similar remarks apply to SUMMATION. Consider the set comprising each of the propositions in the first set, along with the first proposition in the second set. Obviously, this set has the same closure as each. However, it will generate different predictions. On a simple metric, either of the former two ways of characterizing what is relevant predict that the register will be equally accurate at worlds at which exactly one student marked present is absent and at worlds at which exactly one student marked absent is present. In contrast, given the way of

Some ways of representing evidence will be able to be ruled out as insufficiently natural. There is a sense in which the same information is carried by the set which contains, for each student marked present, the proposition that they were present along with the conjunction of the propositions, for each student marked absent, that they were absent. However, this alternative should, we will assume, be excluded as unacceptably *ad hoc*.

Yet in many cases (as above, for example) there will be more than one non-*ad hoc* way of representing the same evidence. In such cases, it may be indeterminate what propositions determine the accuracy order. As a result, how worlds are ordered for accuracy will be indeterminate too. Nevertheless, (some implementation of) ACCURACY may still be (determinately) true, as long as it holds under any resolution of the indeterminacy in the series.

4.2 Probabilistic Theories

Probabilistic theories characterize accuracy in terms of how likely an agent's evidence makes various hypotheses. Here, we follow Williamson 2000 in modeling this kind of likelihood with an evidential probability function Pr_w (cf. Kyburg (1971); Moser (1988)). We assume that, at any worlds which agree on the agent's evidence, her evidential probability function will be the same. Probabilisitic theories then hold that, where $v \in E(w)$, $w \leq v$ iff Pr_w is as accurate at w as at v.[23]

Once we have a notion of evidential probability, we need a way of employing it to measure the accuracy of an agent's evidence. To do this, we can think about how close the evidential probabilities come to getting it right about the actual world. Such measures have been studied at length in connection to rational degrees of belief (Joyce (1998); ?); Leitgeb and Pettigrew (2010b); Pettigrew (2016)) and outright belief (Easwaran (2016); Dorst (2019)). The guiding idea is to understand the proximity of Pr_w to a world in terms of a measure $d(Pr, p, w)$ of how close $Pr_w(p)$ is to the truth value of p at w.

A widespread assumption is that how close a probability function comes to being right about a world can be measured by the sum of how close it is to being right about each proposition in its domain (Pettigrew (2016)). That is, where Π is some way of partitioning logical space into various propositions:

SUMMATION   If $v \in E(w)$, then $w \leq v$ iff $\sum_{p \in \Pi} : d(Pr_w, p, w) \geq \sum_{p \in \Pi} : d(Pr_v, p, v)$.

Given SUMMATION, the accuracy order will be connected over worlds with the same evidence.

characterizing what is relevant proposed above, worlds at which a present student is marked absent will be strictly less accurate than worlds at which an absent student is marked present.

[23]Our discussion is also compatible with other approaches on which an agent's evidence can be represented probabilistically, such as Morrison (2016)'s account of perceptual experience.

SUMMATION makes different predictions, depending on the notion of distance. One prominent approach appeals to squared distance. That is, where I_w is the indicator function mapping each proposition to its truth value at w, $d(w, Pr, p) = |I_w(p) - Pr(p)|^2$. Combined with SUMMATION, this gives us the Brier score (Brier 1950).[24]

This theory connects accuracy-first epistemology to the theory of knowledge. Beddor (2020) and others suggest that a rational agent sets her degrees of belief to the evidential probabilities. ACCURACY then implies that whenever w and v have the same evidence, an agent's rational credences at v are more accurate than at w iff she knows more at v than w.

SUMMATION can be understood in another way. Given the Brier score, it says that the accuracy of the evidence at a world is the likelihood of that world on the evidence.[25]

WORLD PROBABILITY   If $v \in E(w)$ have the same evidence, then $w \leq v$ iff $Pr_w(w) \leq Pr_w(v)$.

SUMMATION assumes that an agent's evidence can be represented by a unique probability distribution. For many, this assumption may seem implausibly strong (cf. Douven (2009); Titelbaum (2010); Kelly (2013); Schoenfield (2014); Callahan (forthcoming)). However, this assumption can be weakened while retaining the idea that accuracy can be measured by the weight evidence assigns to different hypotheses. To do so, we can rely on the notion of imprecise probability. The idea is that in some cases, the evidence does not favor every proposition to a precise degree. Instead, the evidence assigns each proposition a range of probabilities. Imagine you are looking at a wall painted a color somewhere between red and orange. Your evidence may not assign a determinate probability to the claim that the wall is red. Rather, the probability that the wall is red on the evidence might be some range, say $[.4, .6]$.

Following Levi (1974, 1980); Fraassen (1980); van Fraassen (1984); Walley (1991), we can characterize imprecise evidential probabilities in terms of a representor: a set of precise probability functions. We can then define the accuracy of the imprecise evidential probabilities by universally quantifying over each precise probability function. Where P_w is a set of precise probability functions representing an agent's evidence at w:

SUPERVALUATION   If $v \in E(w)$, then $w \leq v$ iff for every $Pr \in P_w$: $\sum_{p \in \Pi} d(Pr, p, w) \geq \sum_{p \in \Pi} d(Pr, p, v)$.

SUPERVALUATION allows for incomparabilities. For example, imagine we assign evidential probabilities to (the singletons of) three worlds: w, v, and u. Imagine the imprecise evidential probabilities are represented by (the convex closure of)

---

[24]An alternative is to identify proximity with absolute distance (Maher (2002)).

[25]Thanks to Ben Levinstein for proving this result. See Goldstein and Hawthorne forthcoming and Goodman and Salow (2021) for further discussion of how to define knowledge-like operators in terms of such a likelihood ordering over worlds.

two probability measures, Pr_1 and Pr_2. Both measures agree that w is most likely (assigning it .4). But they disagree about v and u. Pr_1 says that v is .35 likely, and u is .25 likely; Pr_2 says that u is .35 likely and v is .25 likely. Each measure in the convex closure of $\{Pr_1, Pr_2\}$ is comes closer to getting things right at w than at any other world. So the evidence is more accurate at w than at v or u. However, the accuracy of the evidence is incomparable between v and u. Some measures in the representor come closer to getting it right at v than at u; others come closer to getting things it at u than at v. Generalizing, this view allows for failures of CONNECTEDNESS whenever two measures in the representor disagree on which of two worlds is more likely.

Summarizing, we have considered five different theories of accuracy which are associated with different structural conditions. CONJUNCTION (propositional) and SUMMATION (propositional and probabilistic) based theories yield an order which is connected. Among the propositional theories, the former but not the latter yields a trivial order when the set of relevant propositions is inconsistent. In contrast, SUPERVALUATION (propositional and probabilistic) based theories can yield an order which is not connected. The propositional theory will yield a directed order as long as the set of relevant propositions is consistent. The probabilistic theory will yield a directed order as long as all of the measures in the representor agree on the most probable world.

5 Consequences of Accuracy

We now show that ACCURACY has consequences for the theory of justification and knowledge. The exact shape of these consequences depends on the theory of accuracy we adopt. In this way, the study of accuracy promises rewards for traditional epistemology.

5.1 Justification

Consider the following instance of the preface puzzle.

The Preface:
Alex is a historian. She is just finishing a book about Napoleon I. The main body of the book contains 999 carefully researched claims about his life and times. For each claim, Alex has gathered several pieces of evidence. Reflecting on the fallibility of historical investigation, however, Alex recognizes that even the most carefully researched books tend to contain some errors. She adds a final claim as a preface of her book: 'Each claim in the body of this book is carefully researched, but at least one is false'. (Makinson 1965)

A common verdict about The Preface is that Alex is justified in believing each

of the claims in the preface and body of her book. This requires rejecting CONSISTENCY:

CONSISTENCY   If you are justified in believing each of $p_1, ..., p_n$, then $\{p_1, ..., p_n\}$ is consistent.

Consistency says that any finite set of justified propositions must be jointly consistent. We will show that, in the presence of relatively weak assumptions, STRONG and WEAK ACCURACY imply CONSISTENCY. Accordingly, there is at least some tension between the stronger forms of ACCURACY and standard treatments of preface puzzles.

The first assumption is that if you can epistemically rule out the possibility that you know that p, believing that p would be unjustified. That is, to be justified in believing p one must be in a state epistemically indiscriminable from being in a position to know p.[26]

CAUTION   You aren't justified in believing what you can know you can't know.

Versions of CAUTION have been embraced by Lenzen (1978); Stalnaker (2006); Williamson (2013); Rosenkranz (2018) and Carter and Goldstein (2021) among others. CAUTION is especially natural for those who accept a norm of belief which requires one to be in a state as least as strong as knowledge.[27] If you are in a position to know that you fail to satisfy the norm governing some kind of action, then presumably you would be unjustified in performing that action.

However, those who endorse a norm of belief weaker than knowledge also have reason to entertain CAUTION. In many cases, the only way of coming to know that you do not know p will be by coming to know either that p is false, that p is unjustified, or that your evidence regarding p is not appropriately connected to the matter of whether p. Yet knowledge that your epistemic position regarding p is defective in one of these ways is, arguably, sufficient to defeat justification for believing p.

The second assumption is that the accuracy ordering is directed. That is, for any worlds which share the same evidence it is possible to find a world with that evidence at which it is at least as accurate as it is at both. As we have seen, this will hold on many (though not all) of the ways of characterizing accuracy we considered.

Finally, as an idealization, we assume that the agent's evidence is transparent to her.[28]

TRANSPARENCY   You can know you possess the evidence you possess.

---

[26]Let $J(w)$ be the set of propositions justified at w. Then CAUTION says that if $p \in J(w)$, then $\exists v \in \bigcap K(w) : p \in K(v)$.

[27]See Williamson (2000); Adler (2003); Sutton (2005, 2007); ?; and ? among others.

[28]TRANSPARENCY says that for any w, $E(w) \in K(w)$

Before presenting it in full, we sketch our argument informally. DIRECTEDNESS implies that for any finite set of worlds, X, which share the same evidence, it is possible to find a world with the same evidence which is at least accurate as all of them. By WEAK ACCURACY, there is a world at least as accurate as the latter world at which you can know at least as much as you can at each of the worlds in X. By factivity, the propositions which can be known at this world must be consistent. So, the union of the propositions which can be known at worlds in X must be consistent too. It follows that, for any finite set of worlds which share the same evidence, the propositions which can be known at those worlds must be consistent. But, by CAUTION, you are justified in believing each member of a set of propositions only if, for each proposition in that set, there is a world compatible with what you can know at which it can be known. By the assumption that evidence is transparent, all worlds compatible with what you can know must share the same evidence. Yet, putting these observations together, it follows that for any finite set of propositions, you are justified in believing each member of that set only if the set is consistent.

Fact 1. WEAK ACCURACY, DIRECTEDNESS, TRANSPARENCY and CAUTION imply CONSISTENCY.

Proof: Let $J(w)$ be the set of propositions justified at w. We assume that knowledge is factive, so that $w \in \bigcap K(w)$. Suppose for reductio that for some series $p_1, ..., p_n$ of jointly inconsistent propositions, $\{p_1, ..., p_n\} \subseteq J(w)$. First, CAUTION implies that for each $p_i \in \{p_1, ..., p_n\}$, there is some world $w_i \in \bigcap K(w)$ such that $p_i \in K(w_i)$. By the assumption that evidence is transparent, it follows $w_1, ..., w_n \in E(w)$. By DIRECTEDNESS, we can infer that there is some world $w_{n+1} \in E(w)$ such that for each w_i: $w_i \leq w_{n+1}$. So, by WEAK ACCURACY, there is a world $w_{n+2} \in E(w)$ where $\{p_1, ...p_n\} \subseteq K(w_{n+2})$. By factivity, it follows that $w_{n+2} \in \bigcap\{p_1, ..., p_n\}$. But $p_1, ..., p_n$ are jointly inconsistent, so $\bigcap\{p_1, ..., p_n\} = \emptyset$. Contradiction.

Since STRONG ACCURACY implies WEAK ACCURACY, it follows immediately from **Fact 1** that STRONG ACCURACY, DIRECTEDNESS, TRANSPARENCY and CAUTION imply CONSISTENCY as well. No corresponding result holds for VERY WEAK ACCURACY, however. The key is that the weaker principle permits us to infer only, for each w_i, the existence of a world at least as accurate as w_{n+1} at which at least as much can be known as at w_i. It does not permit us to infer the existence of a world at least as accurate as w_{n+1} at which at least as much can be know as at each w_i.

One response to **Fact 1** would be to reject CAUTION. After all, there is a class of cases where CAUTION is controversial. You cannot know, on the basis of your poor odds alone, that you will lose a fair lottery. In fact, given sufficient reflection, you can know that you cannot know this. Yet, according to some, you are nevertheless justified in believing that you will lose (e.g., Kyburg (1961); Foley (1993); Christensen (2005); and Sturgeon (2008)).

Giving up CAUTION in general is insufficient. It is not enough that there merely be some counter-instances to CAUTION. Every failure of CONSISTENCY must be a failure of CAUTION. Holding fixed WEAK ACCURACY and DIRECTEDNESS, our argument establishes that if your evidence justifies an inconsistent set of propositions, then one of the propositions is known to be unknown. This strikes us as implausible.

The class of cases which motivate relinquishing CAUTION all involve statistical evidence of some kind. Some of these cases are also cases of CONSISTENCY failure. If you are justified in believing, of each ticket, that it will lose, then the lottery will be a case in which a failure of CAUTION accompanies a failure of CONSISTENCY. Crucially, however, CONSISTENCY failures also arise in cases involving no statistical evidence, like the Preface.[29] Yet, the Preface is not a good candidate for CAUTION failure. Alex is not in a position to know, of any of the 1,000 claims in her book, that it is unknown. While she knows that she does not know every claim, she does not know which claims she does not know. If so, CAUTION cannot be what is wrong in cases of this kind.[30]

Instead of rejecting CAUTION, one might instead reject TRANSPARENCY. TRANSPARENCY is at best an idealization. Creatures like us, endowed with imperfect introspection, frequently fail to be in a position to know what evidence we have.

As with CAUTION, however, any such response must be general. Our results establish that there cannot be any cases in which each of the three premises holds and yet an inconsistent set of propositions is justified. But the Preface need not be a case of introspective failure: we can imagine a preface case where the agent's evidence is transparent to her without feeling any pressure to insist that what she is justified in believing must be consistent.[31]

A third reponse rejects DIRECTEDNESS. DIRECTEDNESS is not guaranteed to hold under every theory in §4. In particular, consider the result of combining propositional SUPERVALUATION with an account on which the propositions representing evidence are not required to be consistent. This theory is compatible with failures of DIRECTEDNESS. For suppose that the set of relevant propositions is inconsistent.[32] Then at any pair of worlds making distinct maximal consistent

---

[29]One might worry that the grounds for Alex's claim in her preface are statistical, in a way which constitutes a barrier to knowledge. We can side-step this kind of concern by stipulating that the claim is based on the testimony of a reliable (though fallible) copy-editor instead.

[30]Indeed, even if CAUTION fails in our earlier preface example, we can construct a variant in which CAUTION is locally indisputable. Consider the set of all propositions which the agent justifiably believes and which, for all the agent knows, the agent knows. Reflecting on her fallibility, the agent seems justified in believing that at least one of these many claims is false. In this case, there must be a CONSISTENCY failure with respect to this set. But, by hypothesis, CAUTION holds locally.

[31]An alternative form of the argument can be run without assuming transparency of evidence, by replacing CAUTION with the principle that if you are justified in believing p at w, then there is a v in $E(w)$ where you know p. See Bird (2007) and Ichikawa (2014) for a defense of principles of this general form.

[32]Whether the preface is such a case is unclear. It will depend on which claims in the book

subsets of that set true, there will be no world with the same evidence which is at least as accurate as both. Under SUPERVALUATION, such a world would need to make every proposition in each maximal consistent subset true. Yet the two sets are maximal and distinct. In this way, adopting a model of accuracy along these lines can allow its proponent to accommodate the possibility of inonsistent sets of justified propositions without rejecting WEAK ACCURACY.

By contrast, several other theories of accuracy we considered did require DIRECTEDNESS. For example, according to WORLD PROBABILITY the accuracy of the evidence at a world is proportionate to the evidential probability of that world, then the evidence is most accurate at whichever world has the highest evidential probability. Proponents of such a theory must choose between WEAK ACCURACY, CONSISTENCY, and CAUTION. Things go similarly for propositional theories of accuracy which appeal to CONJUNCTION or SUMMATION.

The final response would be to reject WEAK ACCURACY (and, hence, STRONG ACCURACY). As we saw above, this approach would be compatible with retaining VERY WEAK ACCURACY, as an attempt to capture the informal idea with which we started. As we noted in §2, VERY WEAK ACCURACY is very weak. It is compatible with there being two propositions, each of which can be known when one's evidence is inaccurate, but which could not be known simultaneously, regardless of how accurate one's evidence were. One lesson of the Preface may be that this is the most we can hope for.

We've seen that some, but not all, of our accuracy principles are in tension with structural properties widely attributed of justification. In the next section, we'll see that all of these principles lead to surprising consequences for the theory of knowledge.

5.2   Knowledge

ACCURACY also has consequences for the theory of knowledge. It generates a type of epistemic blindspot, implying that agents cannot know that their evidence is inaccurate.

Suppose that w and v agree on the evidence and that the evidence is more accurate at v than w. Suppose for reductio that at w the agent knows her evidence is less accurate than at v. VERY WEAK ACCURACY implies that there is a world z that is at least as accurate as v and where the agent knows at least as much as at w. This contradicts our assumption that at w the agent knows her evidence is less accurate than at v, since this would imply that the agent at

---

are relevant to the accuracy of the Preface agent's evidence. On one way of characterizing the Preface case we consider, only claims in the body of the book are directly justified. Accordingly, if only directly justified propositions represent evidence, the set of propositions representing the agent's evidence will be consistent. It is possible that, under this kind of approach, whether accuracy is directed in a preface-style case may vary according to details of how it is spelled out. We briefly consider these issues, in relation to what can be known in the Preface, in §5.2.

z knows she is not in z.

Generalizing from this case, VERY WEAK ACCURACY implies that at any world w and any world v more accurate and with the same evidence, it is compatible with what the agent knows at w that her evidence is at least as accurate as it is v. Where $\mathsf{Acc}(w) = \{v \in E(w) \mid w \leq v\}$ is the set of worlds in $E(w)$ that are at least as accurate as w, we can put this as follows:

WEAK BLISSFUL IGNORANCE $\quad \forall v \in \mathsf{Acc}(w) : \bigcap(\mathsf{Acc}(v) \cup K(w)) \neq \emptyset$.

WEAK BLISSFUL IGNORANCE implies that if there are any worlds at which an agent's evidence is maximally accurate, then she cannot rule out that she is in such a world.

STRONG ACCURACY implies something stronger. Suppose again that w and v agree on the evidence and that the evidence is at least as accurate at v as at w. STRONG ACCURACY implies that what can be known at v includes what can be known at w. But factivity implies that what can be known at v is consistent with being in v. The result is that at w an agent cannot know she is not in v. More generally, STRONG ACCURACY implies a kind of obliviousness to actual inaccuracy. At any world w, any world v where the agent's evidence is at least as accurate as at w is epistemically possible.

STRONG BLISSFUL IGNORANCE $\quad \mathsf{Acc}(w) \subseteq \bigcap K(w)$.

STRONG BLISSFUL IGNORANCE is stronger than WEAK BLISSFUL IGNORANCE. It does not merely say that, for every way of being more accurate, it is always epistemically possible that one's evidence is at least as accurate as that. It says that every way of being more accurate while having the same evidence is epistemic possible.[33]

WEAK BLISSFUL IGNORANCE is a correlate of the kind of enrichment of anti-skepticism which holds that one's epistemic position is better in good case than it is in the bad case. If one knows strictly more in the good case than the bad case then in the bad case, one cannot know that one is not in the good case. STRONG BLISSFUL IGNORANCE extends this view to say that every bad case stands in a similar relation to any better case.

What our two blissful ignorance principles say exactly depends on what accuracy is. For example, if WORLD PROBABILITY is correct, then each says something about what the agent can know about the evidential probabilities. STRONG BLISSFUL IGNORANCE says that at any world w, any world as likely as w according to the evidential probabilities is compatible with the agent's knowledge (cf. Goldstein and Hawthorne (forthcoming)). WEAK BLISSFUL IGNORANCE

---

[33]STRONG BLISSFUL IGNORANCE and KK imply STRONG ACCURACY within a normal modal logic for knowledge. Suppose that w and v have the same evidence and that v is as accurate as w. Suppose some p is known at w and is unknown at v. Then KK implies that v is not epistemically accessible from w (since it is known at w that p is known). But this contradicts STRONG BLISSFUL IGNORANCE.

says that at any world w, if v is at least as likely as w, then it is compatible with the agent's knowledge that they are at a world whose evidential probability is at least as high as v.

By contrast, consider our propositional theories of accuracy. Under SUPERVAL-UATION, the evidence is at least as accurate at w as at v iff every proposition representing it which is true at v is true at w. In this case, STRONG BLISSFUL IGNORANCE says that you can't know, of any proposition representing your evidence, that that proposition is false. WEAK BLISSFUL IGNORANCE says that, for each maximal consistent subset of the propositions representing your evidence, you can't know that some member of that set is false.

Is either variant of blissful ignorance tenable? Proponents of the enriched anti-skepticism discussed in §3.1 certainly think so. As Williamson (2000) puts it "part of the badness of the bad case is that one cannot know just how bad one's case is" (165). However, preface cases have the potential to pose some trouble for this view. In particular, depending on the characterization of evidence, both blissful ignorance principles will be incompatible with the possibility of an agent possessing knowledge in the Preface.

Say that an agent has preface knowledge iff she knows that some claim in the body of the book is false. Whether blissful ignorance principles are compatible with preface knowledge depends on the relative accuracy of 'preface worlds' (where some claim in the body of the book is false), and 'body worlds' (where every claim in the body of the book is true). STRONG BLISSFUL IGNORANCE rules out preface knowledge on the assumption that, for any preface world, there is some body world which is at least as accurate. WEAK BLISSFUL IGNORANCE rules out preface knowledge on the assumption that, for every preface world, there is a more accurate body world at which all of the worlds at least as accurate are body worlds.

The relative accuracy of body and preface worlds depends on the underlying conception of accuracy and evidence. In probabilistic theories, this boils down to the relative evidential probability of preface and body worlds. On propositional theories, what will matter is whether the propositions representing the agent's evidence include the claims in the book or only those in its main body.[34] As above, the tenability of ACCURACY turns out to depend on broad issues to do

---

[34]As noted in footnote 32, the fact that the preface claim is justified by an agent's evidence does not by itself settle whether it represents that evidence. If only directly justified propositions are relevant and the preface is not directly justified, then WEAK BLISSFUL IGNORANCE may rule out preface knowledge. The same holds if only propositions which are evidence represent it and only claims in the body of the book meet the standard of evidence.

It appears coherent to accept that the preface claim may be known in preface worlds even if one denies that it is directly justified or part of evidence. Similarly, it appears coherent to entertain the possibility of acquiring preface knowledge on the basis of evidence which makes each individual preface world less likely than some body world. After all, assuming body claims are probabilistically independent, this will be the case wherever each body claim has a probability greater than .5. Yet, for the reasons noted above, both these positions will be incompatible with ACCURACY.

with one's theory of evidence and accuracy.

6   Conclusion

ACCURACY says that what an agent knows is constrained by the accuracy of her evidence. But ACCURACY does not by itself offer a complete theory of knowledge. One question for future research is whether it is possible to offer an accuracy-first reduction of knowledge to evidential accuracy. Here, one option is to define knowledge directly in terms of the accuracy ordering. For example (following similar ideas about normality in Beddor and Pavese 2018 and Goodman and Salow 2018), a natural proposal would be that an agent is in a position to know p at w iff p is entailed by $\mathsf{Acc}(w)$.[35] This is a full-fledged accuracy-first theory of knowledge that implies STRONG ACCURACY. Open questions for future research include how and whether this kind of proposal can be varied to produce interesting accuracy-first alternatives, including those that validate WEAK or VERY WEAK ACCURACY without STRONG ACCURACY.

---

[35]See Goldstein and Hawthorne (forthcoming); Goodman and Salow (manuscript) and Carter (forthcoming) for a development of this proposal where accuracy is understood in terms of WORLD PROBABILITY.

References

Adler, J. (2003), *Belief's Own Ethics*, MIT Press.

Arnold, A. (2013), 'Some evidence is false', *Australasian Journal of Philosophy* **91**(1), 165–172.

Beddor, B. (2020), 'New work for certainty', *Philosophers' Imprint* **20**(8).

Beddor, B. and Pavese, C. (2018), 'Modal virtue epistemology', *Philosophy and Phenomenological Research* .

Bird, A. (2004), 'Is evidence non-inferential?', *Philosophical Quarterly* **54**(215), 252–265.

Bird, A. (2007), 'Justified judging', *Philosophy and Phenomenological Research* **74**(1), 81–110.

Brier, G. (1950), 'Verification of forecasts expressed in terms of probability', *Monthly Weather Review* **78**, 1–3.

Callahan, L. F. (forthcoming), 'Epistemic existentialism', *Episteme* pp. 1–16.

Carter, S. (2019), 'Higher order ignorance inside the margins', *Philosophical Studies* pp. 1–18.

Carter, S. (forthcoming), 'Degrees of Assertability', *Philosophy and Phenomenological Research* .

Carter, S. and Goldstein, S. (2021), 'The normality of error', *Philosophical Studies* .

Chalmers, D. J. (2006), Perception and the fall from eden, *in* T. S. Gendler and J. Hawthorne, eds, 'Perceptual Experience', Oxford University Press, pp. 49–125.

Christensen, D. (2005), *Putting Logic in its Place*, Oxford University Press.

Cohen, S. and Comesaña, J. (2013), 'Williamson on gettier cases and epistemic logic', *Inquiry: An Interdisciplinary Journal of Philosophy* **56**(1), 15–29.

Conee, E. and Feldman, R. (2008), Evidence, *in* Q. Smith, ed., 'Epistemology: New Essays', Oxford University Press.

Davidson, D. (2001), *Subjective, Intersubjective, Objective: Philosophical Essays Volume 3*, Clarendon Press.

Dorst, K. (2019), 'Lockeans maximize expected accuracy', *Mind* **128**(509), 175–211.

Douven, I. (2009), 'Uniqueness revisited', *American Philosophical Quarterly* **46**(4), 347–361.

Easwaran, K. (2016), 'Dr. truthlove or: How i learned to stop worrying and love bayesian probabilities', *Noûs* **50**(4), 816–853.

Fantl, J. and McGrath, M. (2009), *Knowledge in an Uncertain World*, Oxford University Press.

Foley, R. (1993), *Working Without a Net*, Oxford University Press.

Fraassen, B. C. V. (1980), 'Rational belief and probability kinematics', *Philosophy of Science* **47**(2), 165–187.

Goldman, A. (2009), Williamson on knowledge and evidence, *in* P. Greenough, D. Pritchard and T. Williamson, eds, 'Williamson on Knowledge', Oxford University Press, pp. 73–91.

Goldstein, S. and Hawthorne, J. (forthcoming), 'Knowledge from multiple experiences', *Philosophical Studies* pp. 1–32.

Goodman, J. (2013), 'Inexact knowledge without improbable knowing', *Inquiry: An Interdisciplinary Journal of Philosophy* **56**(1), 30–53.

Goodman, J. and Salow, B. (2018), 'Taking a chance on kk', *Philosophical Studies* **175**(1), 183–196.

Goodman, J. and Salow, B. (2021), 'Knowledge from Probability', *Electronic Proceedings in Theoretical Computer Science* **335**, 171–186.

Goodman, J. and Salow, B. (manuscript), 'Epistemology normalized'.

Greco, D. (2014), 'Could kk be ok?', *Journal of Philosophy* **111**(4), 169–197.

Hawthorne, J. (2003), *Knowledge and Lotteries*, Oxford University Press.

Hawthorne, J. (Forthcoming), The epistemic use of 'ought', *in* L. Walters and J. Hawthorne, eds, 'Conditionals, Paradoxes, and Probability: Themes from the Philosophy of Dorothy Edgington', Oxford University Press.

Heylen, J. (2016), 'Being in a position to know and closure', *Thought: A Journal of Philosophy* **5**(1), 63–67.

Hilpinen, R. (1976), Approximate truth and truthlikeness, *in* M. Przełecki, K. Szaniawski and R. W'Ojcicki, eds, 'Formal Methods in the Methodology of the Empirical Sciences', Reidel, pp. 19–42.

Huemer, M. (2006), 'Phenomenal conservativism and the internalist intuition', *American Philosophical Quarterly* .

Hyman, J. (2006), 'Knowledge and evidence', *Mind* **115**(460), 891–916.

Ichikawa, J. J. (2014), 'Justification is potential knowledge', *Canadian Journal of Philosophy* **44**(2), 184–206.

Joyce, J. M. (1998), 'A nonpragmatic vindication of probabilism', *Philosophy of Science* **65**(4), 575–603.

Joyce, J. M. (2005), 'How degrees of belief reflect evidence', *Philosophical Perspectives* **19**(1), 153–179.

Kelly, T. (2013), Evidence can be permissive, *in* M. Steup and J. Turri, eds, 'Contemporary Debates in Epistemology', Blackwell, p. 298.

Kieseppä, I. A. (1996*a*), 'Truthlikeness for hypotheses expressed in terms of N quantitative variables', *Journal of Philosophical Logic* **25**(2), 109–134.

Kieseppä, I. A. (1996*b*), *Truthlikeness for Multidimensional, Quantitative Cognitive Problems*, Synthese Library, Springer Netherlands.

Kyburg, H. E. (1961), *Probability and the Logic of Rational Belief*, Wesleyan University Press.

Kyburg, H. E. (1971), 'Epistemological probability', *Synthese* **23**(2-3), 309–326.

Leitgeb, H. and Pettigrew, R. (2010*a*), 'An objective justification of bayesianism i: Measuring inaccuracy', *Philosophy of Science* **77**(2), 201–235.

Leitgeb, H. and Pettigrew, R. (2010*b*), 'An objective justification of bayesianism ii: The consequences of minimizing inaccuracy', *Philosophy of Science* **77**(2), 236–272.

Lenzen, W. (1978), 'Recent work in epistemic logic', *Acta Philosophical Fennica* **30**.

Levi, I. (1974), 'On indeterminate probabilities', *Journal of Philosophy* **71**(13), 391–418.

Levi, I. (1980), *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*, MIT Press.

Lewis, D. K. (1973), *Counterfactuals*, Blackwell.

Littlejohn, C. (2012), *Justification and the Truth Connection*, Cambridge University Press.

Littlejohn, C. (2013), 'No evidence is false', *Acta Analytica* **28**(2), 145–159.

Littlejohn, C. and Dutant, J. (2020), 'Justification, knowledge, and normality', *Philosophical Studies* **177**(6), 1593–1609.

Loets, A. (forthcoming), 'Choice points for a theory of normality', *Mind* .

Maher, P. (2002), 'Joyce's argument for probabilism', *Philosophy of Science* **69**(1), 73–81.

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

Morrison, J. (2016), 'Perceptual confidence', *Analytic Philosophy* **57**(1), 15–48.

Moser, P. K. (1988), 'The foundations of epistemological probability', *Erkenntnis* **28**(2), 231–251.

Moser, P. K. (1989), *Knowledge and Evidence*, Cambridge University Press.

Niiniluoto, I. (1987), *Truthlikeness*, Dordrecht: Reidel.

Niiniluoto, I. (1998), 'Verisimilitude: The third period', *British Journal for the Philosophy of Science* **49**(1), 1–29.

Niiniluoto, I. (2020), 'Truthlikeness: Old and new debates', *Synthese* **197**(4), 1581–1599.

Oddie, G. (1986), *Likeness to Truth*, number 1, Reidel.

Oddie, G. (2013), 'The content, consequence and likeness approaches to verisimilitude: Compatibility, trivialization, and underdetermination', *Synthese* **190**(9), 1647–1687.

Pettigrew, R. (2016), *Accuracy and the Laws of Credence*, Oxford University Press.

Pollock, J. (1974), *Knowledge and Justification*, Princeton University Press.

Pollock, J. and Gillies, A. (2000), 'Belief revision and epistemology', *Synthese* **122**(1-2), 69–92.

Pritchard, D. (2008), Sensitivity, safety, and anti-luck epistemology, *in* J. Greco, ed., 'The Oxford Handbook of Skepticism', Oxford University Press.

Rizzieri, A. (2011), 'Evidence does not equal knowledge', *Philosophical Studies* **153**(2), 235–242.

Rosenkranz, S. (2018), 'The structure of justification', *Mind* **127**(506), 309–338.

Schoenfield, M. (2014), 'Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief', *Noûs* **48**(2), 193–218.

Schroeder, M. (2008), 'Having reasons', *Philosophical Studies* **139**(1), 57–71.

Siegel, S. (2010), *The Contents of Visual Experience*, Oxford University Press USA.

Siewert, C. (1998), *The Significance of Consciousness*, Princeton University Press.

Spohn, W. (2012), *The Laws of Belief: Ranking Theory and its Philosophical Applications*, Oxford University Press.

Stalnaker, R. (2006), 'On logics of knowledge and belief', *Philosophical Studies* **128**(1), 169–199.

Stalnaker, R. (2015), Luminosity and the kk thesis, *in* S. Goldberg, ed., 'Externalism, Self-Knowledge, and Skepticism', Cambridge University Press.

Sturgeon, S. (2008), 'Reason and the grain of belief', *Noûs* **42**(1), 359–396.

Sutton, J. (2005), 'Stick to what you know', *Nous* **39**(3), 359–396.

Sutton, J. (2007), *Without Justification*, MIT Press.

Teller, P. (2001), 'Twilight of the perfect model model', *Erkenntnis* **55**(3), 393–415.

Tichý, P. (1974), 'On Popper's definitions of verisimilitude', *British Journal for the Philosophy of Science* **25**(2), 155–160.

Tichý, P. (1976), 'Verisimilitude redefined', *British Journal for the Philosophy of Science* **27**(1), 25–42.

Tichý, P. (1978), 'Verisimilitude revisited', *Synthese* **38**(2), 175–196.

Titelbaum, M. G. (2010), 'Not enough there there evidence, reasons, and language independence', *Philosophical Perspectives* **24**(1), 477–528.

van Fraassen, B. C. (1984), 'Belief and the will', *Journal of Philosophy* **81**(5), 235–256.

Walley, P. (1991), *Statistical Reasoning with Imprecise Probabilities, Monographs on Statistics and Applied Probability*, Vol. 42, Chapman and Hall.

Weston, T. (1992), 'Approximate truth and scientific realism', *Philosophy of Science* **59**(1), 53–74.

Williamson, T. (2000), *Knowledge and its Limits*, Oxford University Press.

Williamson, T. (2013), 'Gettier cases in epistemic logic', *Inquiry: An Interdisciplinary Journal of Philosophy* **56**(1), 1–14.