Degrees of confidence are not subjective probabilities[1]

Jeremy GOODMAN — *Johns Hopkins University*

Abstract. We assume that *confident* reports have their standard degree-based truth conditions: to be *confident* that p is to have a degree of confidence that p which is at least as high as some contextually determined threshold; to be *more confident* that p than that q is to have a degree of confidence that p that is higher than one's degree of confidence that q; and so on. But what are degrees of confidence? The standard answer is that agents' degrees of confidence are simply their Bayesian subjective probabilities: for example, how confident one is that it's raining = how likely one thinks it is that it's raining. We raise a number of challenges for this Bayesian account, and propose an alternative. This new account supports a pragmatic explanation of the apparent equivalence of degrees of confidence and subjective probabilities, and offers a more integrated picture of how different doxastic attitudes fit together.

Keywords: confidence, gradable doxastic attitudes, subjective probability, plausibility orders.

1. Introduction

Philosophers, economists, and psychologists often find it productive to model agents' opinions as *subjective probabilities* (also called *degrees of belief* or *credences*). Intuitively, an agent's subjective probabilities measure how likely they think it is that various propositions are true. Ideally, these subjective probabilities should satisfy the axioms of probability theory and guide the agent's behavior (for example, by determining which bets are fair by the agent's lights).

Despite the intuitive gloss just given, any reductive analysis of subjective probabilities as beliefs about how likely things are faces significant challenges.[2] As a result, many authors have instead sought to operationalize subjective probability in other ways (Ramsey, 1931; de Finetti, 1937; Jeffrey, 1965), or been content to treat the notion as theoretical posit that earns its keep through its explanatory power (Eriksson and Hájek, 2007).[3]

But perhaps the dominant approach to subjective probabilities is to identify them with degrees of *confidence*. Here are some representative quotations:

> [Subjective] probability measures the confidence that a particular individual has in the truth of a particular proposition" (Savage, 1954: 3)
>
> [T]he levels of confidence you might have in various propositions [are] your degrees of belief in them. (Foley, 1993: 140)

[1]For their helpful feedback thanks to Kyle Blumberg, Cian Dorr, Kevin Dorst, Peter Fritz, John Hawthorne, Dan Hoek, Ben Holguín, Harvey Lederman, Matt Mandelkern, Jessica Moss, Kyle Rawlins, Bernhard Salow, Alexis Wellwood, Tim Williamson, and audiences at SuB, JHU, and the New York Philosophy of Language Workshop.

[2]Challenges include: (i) worries that such analyses over-intellectualize subjective probability, (ii) worries about circularity for those who wish to analyze belief (Genin, 2019) and/or probability operators like *likely* (Rothschild, 2012; Yalcin, 2012; Moss, 2018) in terms of subjective probabilities, and (iii) impossibility results showing that no such analysis is possible in an orthodox (static) semantic framework (Russell and Hawthorne, 2016).

[3]Others still have questioned whether the notion is even in good standing (Holton, 2014; Williamson, ming).

 In: Baumann, Geraldine, Daniel Gutzmann, Jonas Koopman, Kristina Liefke, Agata Renans, and Tatjana Scheffler (eds.) 2024. Proceedings of Sinn und Bedeutung 28. Bochum: Ruhr-University Bochum, 329-344.

> A person's credence in a proposition X is her level of confidence in its truth. (Joyce, 2009: 263)

> Credences are numerical degrees of confidence. (Titelbaum, 2019: 1)

Call this the *Bayesian account*. Is it right? Is the attitude expressed by *confident* reports in ordinary English really the same as these theorists' notion of subjective probability?

This paper argues that it is not. Or, more cautiously, it argues that the primary use of *confident* in ordinary English is one for which degrees of confidence are not subjective probabilities. Obviously many academics writing about subjective uncertainty have a practice of using graded *confident* reports as synonymous with subjective-probability jargon, and this practice comes naturally and is easily inculcated. The extent to which this 'Bayesian use' occurs spontaneously in naïve speakers, and whether it constitutes an ambiguity in *confident*, is an important question, but one which we must leave for future work.

In this connection, though, a word of caution before proceeding. There are some people for whom the Bayesian use of *confident* has become second-nature. For them, *confident* reports immediately prompt explicitly probabilistic thinking. If you are such a reader, and you balk at some of the (in)felicity indications in what follows, you are not alone. But (except where otherwise noted) most speakers consulted shared the indicated judgments.

Here is the plan. After reviewing degree semantics for *confident* reports (§2), we explain and motivate the orthodox Bayesian account of degrees of confidence (§3). We then identify three ways in which degrees of confidence fail to behave like subjective probabilities (§§4-6), propose a new account that explains these data (§§7-9), and conclude with some further theoretical considerations in its favor (§10). Some generalizations of the account are explored in appendices.

To preview, here are two distinctive features of the account developed below. First, the attitude of having (at least) a given degree of confidence in a proposition is modeled by universally quantifying over a non-empty set of accessible worlds. This is the standard way of modeling other doxastic and epistemic attitudes, but it is incompatible with degrees of confidence being subjective probabilities. Second, though, subjective probabilities still play an important role: they are necessary but not sufficient for having the corresponding degrees of confidence, and they coincide with agents' degrees of confidence in the cases we are typically concerned with.

2. Degree semantics for *confident* reports

We begin by rehearsing the standard degree-based truth conditions for *confident*-reports; see Cariani et al. (ming) for discussion of how these can be derived compositionally.

For any agent A, world w, and proposition p, let $conf(A,w,p)$ be the degree of confidence which, in w, is A's degree of confidence that p. This is undefined if, in w, A has no degree of confidence that p (i.e., if the question *how confident is A that p* has a false presupposition at w). We assume that degrees of confidence are totally ordered by a relation $>$.

The truth conditions for comparative confident reports are then the obvious ones:

$[\![A \text{ is more confident that } \varphi \text{ than that } \psi]\!](w) = 1 \text{ iff } conf([\![A]\!], w, [\![\varphi]\!]) > conf([\![A]\!], w, [\![\psi]\!])$

$[\![A \text{ is just as confident that } \varphi \text{ as that } \psi]\!](w) = 1 \text{ iff } conf([\![A]\!], w, [\![\varphi]\!]) = conf([\![A]\!], w, [\![\psi]\!])$

Likewise for explicit talk about degrees of confidence:

$[\![\text{is how confident } A \text{ is that } \varphi]\!](d)(w) = 1 \text{ iff } conf([\![A]\!], w, [\![\varphi]\!]) = d$

A positive (i.e., non-comparative) *confident* report is true just in case the agent's degree of confidence in the relevant proposition meets a contextually determined threshold θ_c^{conf}:

$[\![A \text{ is confident that } \varphi]\!]^c(w) = 1 \text{ iff } conf([\![A]\!]^c, w, [\![\varphi]\!]^c) \geq \theta_c^{\text{conf}}$

3. Motivating the Bayesian account

Let us assume that the notion of subjective probability is in good standing and that, at the present level of idealization, we can model agents' subjective probabilities by a function Pr from agents and worlds to probability distributions. That is, for any agent A and world w, $Pr_{A,w}(\cdot)$ is a probability distribution, and $Pr_{A,w}(p)$ is the agent's subjective probability in w that p is true. As usual, a proposition p is identified with the sets of worlds in which it is true.

We can now formalize the orthodox Bayesian account of degrees of confidence, which identifies them with subjective probabilities:[4]

BAYESIANISM
$conf = Pr$

In what follows we make the standard assumption that *think . . . likely* reports express agents' subjective probabilities (notwithstanding the complications mentioned in note 2). That is:

$[\![A \text{ thinks that } \varphi \text{ is more likely than } \psi]\!](w) = 1 \text{ iff } Pr_{[\![A]\!],w}([\![\varphi]\!]) > Pr_{[\![A]\!],w}([\![\psi]\!])$

$[\![\text{is how likely } A \text{ thinks it is that } \varphi]\!](d)(w) = 1 \text{ iff } Pr_{[\![A]\!],w}([\![\varphi]\!]) = d$

This allows us to motivate BAYESIANISM by observing that, in many cases, graded *confident* reports seem to pattern with *think . . . likely* reports.

(1) *Context*: Alice asks Bob what the capital of Spain is. He says he doesn't know, but he thinks that it's either Madrid or Barcelona. Alice then asks:

a. How confident are you that it's either Madrid or Barcelona?
b. How likely do you think it is that it's either Madrid or Barcelona?

It is very difficult to hear a difference between (1a) and (1b) in this context. More generally:

ANSWER EQUIVALENCE
When a speaker offers their opinion in answer to a question (even if it is just their best guess), their conversation partners do not distinguish *how confident* the speaker is in their answer from *how likely* the speaker *thinks* it is that their answer is true.

[4] More precisely: $conf(A, w, \cdot)$ and $Pr_{A,w}(\cdot)$ are the same (partial) function from $\mathscr{P}(W)$ to $[0, 1]$ (for all worlds w and agents A), and degrees of confidence are real numbers in the unit interval under the usual ordering.

Any account of degrees of confidence must explain this pattern. BAYESIANISM does so in the simplest and most straightforward way.

4. First challenge: doxastic constraints

BAYESIANISM predicts that agents have degrees of confidence in every proposition for which they have subjective probabilities. But this seems incorrect. Consider:

(2) *Context*: Petra, Quinn, and Rita are about to race. Bob bets on Quinn. Alice then tells him that Carl, a seasoned track coach, thinks that Petra will win. Bob dejectedly asks:

a. How confident is he that Petra will win?
b. How likely does he think it is that Quinn will win?
c. #How confident is he that Quinn will win?

Unlike (2b), (2c) seems to have a false presupposition (that Carl has some degree of confidence that Quinn will win). BAYESIANISM wrongly predicts that (2b) and (2c) should be equivalent.

The fact that (2a) is licensed in this context, but (2c) is not, motivates two general principles:

THINKING TRUE
Thinking that a proposition is true entails having some degree of confidence in it.

THINKING FALSE
Thinking that a proposition is false entails not having any degree of confidence in it.

These two principles are silent about cases where agents neither think that p nor think that not-p. What should we say about such cases? The two most natural proposals are:

DISBELIEF EXCEPTION
Agents have degrees of confidence in all propositions in which they have some subjective probability, except for propositions that they think are false.

THINKING REQUIREMENT
Agents have degrees of confidence only in propositions that they think are true.

The remainder of this section argues that the second proposal is preferable to first.

The simplest motivation for the DISBELIEF EXCEPTION is that it is a minimal departure from BAYESIANISM. To that extent, we think it is well motivated only insofar as it can be combined with the following principle (which is also suggested by ANSWER EQUIVALENCE):

BAYESIAN WHEN DEFINED
$conf(A,w,p) = Pr_{A,w}(p)$ whenever $conf(A,w,p)$ is defined.

But there are strong reasons to reject the combination of DISBELIEF EXCEPTION and BAYESIAN WHEN DEFINED. This is because together they conflict with the following three principles:[5]

[5] *Proof*: By NON-EXTREMITY, it is possible to be confident that p and have a non-maximal degree of confidence d in p. By BAYESIAN WHEN DEFINED, d is a non-maximal subjective probability. By NON-ENTAILMENT, it is possible to have subjective probability at least d in a proposition q without thinking that q. Consider such an agent

CONFIDENCE ENTAILMENT
Being confident that p entails thinking that p.

NON-EXTREMITY
The threshold θ_c^{conf} for being *confident* is non-maximal in many ordinary contexts c.

NON-ENTAILMENT
There is no non-maximal degree of subjective probability that entails thinking that p.

These principles can be motivated as follows:

- CONFIDENCE ENTAILMENT is intuitively plausible, and also explains the infelicity of claims like *A is confident that p, but I wouldn't go so far as to say that A thinks that p*.[6]
- NON-EXTREMITY is the natural explanation of the felicity of sentences like:

  (3) She's confident that Djokovic will win Wimbledon, and she's even more confident that the winner will be either Djokovic or Alcaraz. (See also Cariani et al. (ming).)

- NON-ENTAILMENT is defended at length in Holguín (2022) and Williamson (ming). Among other considerations, they observe that, from the fact that a lottery entrant knows how likely they are to lose, we cannot conclude that they think that they will lose.

Let us now turn to THINKING REQUIREMENT. It derives some support from the fact that it is the simplest alternative to DISBELIEF EXCEPTION (which we have just argued against). It is further motivated by the fact that, given CONFIDENCE ENTAILMENT, it holds for all degrees of confidence d that are the threshold θ_c^{conf} for the positive form of *confident* in some context c.[7] And there are many such degrees of confidence, given NON-EXTREMITY, since the positive forms of non-maximal gradable adjectives are generally context-sensitive in this way. For these reasons (among others) we will assume THINKING REQUIREMENT in what follows.

Now observe that, when combined with ANSWER EQUIVALENCE, THINKING REQUIREMENT has the following striking implication:

EXTREME WEAKNESS
There is no non-minimal degree of subjective probability that entails not thinking that p.

This is because an agent may assign low subjective probability to their best guess in answer to a question. So they have the corresponding (low) degree of confidence this answer, by ANSWER EQUIVALENCE, despite thinking that it is true, by THINKING REQUIREMENT. By considering sufficiently fine-grained questions we can make this degree of confidence arbitrarily low.

This picture of how thinking, guessing, and subjective probabilities are related is defended by Holguín (2022) on independent grounds. To illustrate, in response to the question *who will win Wimbledon*, someone who thinks that Djokovic has no more than a 25% chance of winning

who also doesn't think that not-q. By DISBELIEF EXCEPTION and BAYESIAN WHEN DEFINED, they are confident to degree d in q, and hence they are confident that p, contradicting CONFIDENCE ENTAILMENT.

[6]Following Hawthorne et al. (2016) and Rothschild (2020), we use *I wouldn't go so far as to say that they think* rather than explicit negation to control for the fact that *think* neg-raises (i.e., *A doesn't think that p* tends to be interpreted as equivalent to *A thinks that not-p*).

[7]We assume that *think* is not context sensitive in a way that depends on how confident it takes to count as *confident*.

might truly answer *I'm not sure, but I think that Djokovic will win* – provided, that is, that there isn't anyone else who they think is more likely to win.[8]

Let's take stock. We first argued against BAYESIANISM, on the grounds that it conflicts with THINKING FALSE. We then argued against a natural fallback position, the combination of BAYESIAN WHEN DEFINED and DISBELIEF EXCEPTION, on the grounds that it conflicts with CONFIDENCE ENTAILMENT, NON-EXTREMITY, and NON-ENTAILMENT. We then motivated THINKING REQUIREMENT on two grounds: (i) it is the most natural alternative to DISBELIEF EXCEPTION, and (ii) it has many true instances (given CONFIDENCE ENTAILMENT and the relevant context sensitivity of the positive form of *confident*).

Before moving on, it will be convenient to have a name for the following biconditional, which is equivalent to the conjunction of THINKING TRUE and THINKING REQUIREMENT:

> COMMITMENT
> An agent has some degree of confidence in a proposition if and only if they think that the proposition is true.

5. Second challenge: closure without extremity

This section raises a further challenge for BAYESIANISM, which also puts pressure on BAYESIAN WHEN DEFINED. It concerns the closure of *confident* under conjunction.

To illustrate the issue, consider examples like the following:

(4) *Context*: Juan thinks that it is 90% likely that he weights at least 75kg, and thinks that it's 90% likely that he weighs at most 77kg.
   a. #Juan is confident that he weighs at least 75kg and confident that we weighs at most 77kg, but he isn't confident that he weighs between 75 and 77kg.

BAYESIANISM predicts that (4a) is true in contexts c where $.8 < \theta_c^{\text{conf}} \leq .9$. And there are presumably such contexts, given NON-EXTREMITY.[9] Yet (4a) sounds like a contradiction, or at the very least seems to attribute a quite strange state of mind which Juan presumably isn't in.

The example is also a challenge for BAYESIAN WHEN DEFINED. Suppose, as seems natural, that Juan thinks that he weighs between 75 and 77kg, and hence both thinks that he weights at least 75kg and thinks that he weighs at most 77kg. So Juan has some degree of confidence in each of these propositions, by THINKING TRUE. BAYESIAN WHEN DEFINED then wrongly predicts that (4a) should have a true reading in some contexts (where $.8 < \theta_c^{\text{conf}} \leq .9$).

This argument does not assume that *confident* is always closed under conjunction. Still, it will be helpful in what follows to isolate a family of general closure principles:

[8] The fact that such uses of *I think* function as hedged assertions is sometimes offered as a reason to doubt that they are literal first-personal attitude ascriptions; see Clarke (2024). Against this suggestion, note (i) that such uses licenses subsequent third-person reports (e.g., *she thinks that Djokovic will win*), and (ii) their hedging function is naturally explained, as a scalar implicature, by them being genuine first-person attitude ascriptions.

[9] If .9 feels implausibly low, or .8 implausible high, for θ_c^{conf} in an ordinary contexts, we can choose different numbers: the example depends only on $0 < \theta_c^{\text{conf}} < 1$, which is secured by NON-EXTREMITY.

CLOSURE$_{\text{think}}$
If A thinks that p and thinks that q, then A thinks that *p and q*.

CLOSURE$_{\text{conf}}$
If A is confident that p and confident that q, then A is confident that *p and q*.

CLOSURE$_{\text{deg}}$
If A is confident to degree at least d that p and confident to degree at least d that q, then A is confident to degree at least d that *p and q*.

CLOSURE$_{\text{think}}$ is intuitively compelling, and is defended by Holguín (2022). It is also validated by orthodox Hintikka semantics for *think* (see §7). CLOSURE$_{\text{conf}}$ is similarly compelling, and it will be validated by the account of degrees of confidence developed below (since, as advertised, this account yields a version of Hintikka semantics for *confident*).

There are of course well known challenges to CLOSURE$_{\text{think/conf}}$, such as those deriving from Makinson's (1965) influential preface paradox. In appendix B we show how the account of degrees of confidence developed below can be non-disruptively modified to accommodate such (purported) closure failures. But for now, to reiterate: the challenge that (4a) poses for BAYESIANISM does not assume the truth of any general closure principle.

Let us now turn to CLOSURE$_{\text{deg}}$. It can be motivated by CLOSURE$_{\text{conf}}$ in the same way that THINKING REQUIREMENT was motivated by CONFIDENCE ENTAILMENT: CLOSURE$_{\text{conf}}$ is valid only if CLOSURE$_{\text{deg}}$ holds for all degrees of confidence d that could be θ_c^{CONF} in some context c, in which case CLOSURE$_{\text{deg}}$ plausibly holds for all degrees of confidence whatsoever.

CLOSURE$_{\text{deg}}$ entails the validity of further principles involving graded *confident* reports which do not explicitly mention degrees of confidence, such as the following:[10]

CLOSURE$_{=}$
If A is just as confident that p as they are that q, then A is just as confident that *p and q* as they are that p.

It hence predicts the infelicity of sentences like:

(5) ??Deb is just as confident that she will show up tomorrow as she is that she will show up the next day, but she is less confident that she will show up on both days.

Speakers' judgments about such sentences vary, and while the view developed in the main text does validate CLOSURE$_{\text{deg}}$, the more general account explored in appendices B-C does not.

6. Third challenge: second guesses

This section raises a different challenge for BAYESIAN WHEN DEFINED, which does not turn the closure of *confident* under conjunction. However, the relevant judgments are subtle and not shared by all speakers consulted, and so serve more to illustrate some distinctive empirical predictions of the account developed below than to provide a persuasive argument for it.

Consider the following variant of (3):

[10] This entailment assumes that agents are never more confident in a conjunction than they are in its conjuncts.

(6) *Context*: Joey is prognosticating about who will win an upcoming squash tournament.
a. Parke: Who do you think will win?
b. Joey: Nouran Gohar.
c. Parke: How confident are you that she will win?
d. Joey: Fairly confident, and I'm even more confident that the winner will be either Gohar or El Sherbini.

Many speakers take (6d) to imply that El Sherbini is the player who Joey thinks is second most likely to win (after Gohar). While this could have a pragmatic explanation, it is worth isolating a general principle which would license a purely semantic explanation of such inferences:

> SECOND GUESSING
> If an agent thinks that p is the true complete answer to a question they are considering, and they are even more confident of the disjunction p *or* q (where q is another complete answer to this question), then q is the answer that they think is next most likely after p.

Note finally that BAYESIAN WHEN DEFINED does not merely fail to support the inference from (6d) to El Sherbini being Joey's second guess. It also seems to predict, incorrectly, that the second conjunct of (6d) should sound redundant (provided it is common ground that El Sherbini has a non-trivial chance of winning), since it should be obvious that Joey has higher subjective probability that either Gohar or El Sherbini will win than he has that Gohar will win.

7. Part I: Hintikka semantics for *confident* and *think*

This section outlines a simple account of degrees of confidence, and shows how it can be used to validate both COMMITMENT from §4 and the family of CLOSURE principles discussed in §5. Assessing the other principles discussed above will require bringing in subjective probabilities (see §8) and question sensitivity (see §9), which will also help in motivating and fixing the intended interpretation of the formalism introduced here.

The basic idea is to use Hintikka semantics for degrees of confidence. Informally, agent A is confident to degree at least d in the proposition p if and only if p is true in all of A's confident-to-degree-at-least-d worlds.

More precisely, for any agent A and degree of confidence $d \in [0,1]$, $C^d_A(w)$ is the set of worlds compatible with everything that A is, in w, confident of to degree at least d. As such, we require that $C^d_A(w) \subseteq C^{d'}_A(w)$ whenever $d > d'$ (since being confident to degree at least d entails by being confident to degree at least d' for any $d' > d$). It follows that $C^0_A(w) = \bigcap\{C^d_A(w) : d \in [0,1]\}$.

We can now define the function $conf$ used to specify the truth conditions for *confident* reports:[11]

> DEGREES OF CONFIDENCE
> $conf(A,w,p) = max\{d : C^d_A(w) \subseteq p\}$; this is undefined if $C^0_A(w) \not\subseteq p$

We adopt standard Hintikkan truth conditions for *think*. Where $D_A(w)$ is the set of 'doxastically accessible' worlds (those compatible with everything that, in w, A thinks is true) we have:

[11] This definition assumes that $\{d : C^d_A(w) \subseteq p\}$ contains its own supremum whenever it is non-empty, which is a consequence of the probabilistic constraints introduced in §8.

$[\![A \text{ thinks that } \varphi]\!](w) = 1$ iff $D_{[\![A]\!]}(w) \subseteq [\![\varphi]\!]$

Finally, we relate *think* to *confident* by identifying the set of doxastically accessible worlds with the set of worlds compatible with everything that the agent has any degree of confidence in:

DOXASTIC ACCESSIBILITY
$D_A(w) = C^0_A(w)$

The orthodox degree semantics for *confident* reports in §2, together with Hintikka semantics for *think*, then validates COMMITMENT as well as all of CLOSURE$_{\text{think/conf/deg/=}}$.

8. Part II: Probability + Plausibility ⇒ Confidence

This section explains how degrees of confidence relate to subjective probabilities.

Having rejected BAYESIANISM, we cannot characterize degrees of confidence from subjective probabilities alone. Any such characterization will require introducing additional structure.

To this end, we appeal to a relation of 'comparative plausibility' between worlds. Let $v \succeq_{A,w} u$ mean that, in w, A finds v at least as plausible as u. This is a technical notion, which will be analyzed further in §9 (where $\succeq$ is defined from subjective probabilities and a contextually determined question). For now, we require only that $\succeq_{A,w}$ be a well-founded total preorder on the set of possible worlds W. Following Lewis (1973), we can repackage the information encoded in $\succeq_{A,w}$ using a set $\$_{A,w}$ of 'spheres', defined as follows:

$\$_{A,w} = \{\{u : u \succeq_{A,w} v\} : v \in W\}$

We are now in a position to given an explicit definition of C, the basic notion in terms of which degrees of confidence are defined:

DEGREES FROM PROBABILITY
$C^d_A(w) = \bigcap\{p \in \$_{A,w} : Pr_{A,w}(p) \geq d\}$

Given DEGREES OF CONFIDENCE, this entails that an agent's degree of confidence in p is the highest subjective probability that the agent assigns to any p-entailing sphere:

CONFIDENCE FROM PROBABILITY SPHERES
$conf(A, w, p) = max\{Pr_{A,w}(q) : q \subseteq p \text{ and } q \in \$_{A,w}\}$

Moreover, given DOXASTIC ACCESSIBILITY, it entails that the doxastically accessible worlds are all and only the worlds that the agent takes to be maximally plausible:

PLAUSIBLE OPINIONS
$D_A(w) = \{v : \forall u(u \not\succ_{A,w} v)\}$ (where $v \succ_{A,w} u$ iff $v \succeq_{A,w} u$ and $u \not\succeq_{A,w} v$)

These direct characterizations of $conf$ from Pr and $\succeq$ (via \$) and of D from $\succeq$ will prove useful when considering generalizations of the present proposal, as discussed in the appendices.

One nice consequence of DEGREES FROM PROBABILITY is that how confident an agent is that a given proposition is true is never less than how likely the agent thinks it is that the proposition is true. That is, we validate the following weakening of BAYESIAN WHEN DEFINED:

LOWER BOUND
$conf(A,w,p) \leq Pr_{A,w}(p)$ whenever $conf(A,w,p)$ is defined.

9. Part III: Probability + Partitions ⇒ Plausibility

This section gives a probabilistic analysis of comparative plausibility. This analysis essentially appeals to a contextually determined *question*. The resulting question-sensitivity of *confident* reports is not unprecedented (Yalcin, 2018; Hoek, ming), and it will allow us to account for ANSWER EQUIVAELNCE despite having rejected BAYESIANISM.

We model the relevant question Q as a partition of W, and write $[w]_Q$ for the cell of Q containing w. Comparative plausibility is defined in terms of the comparative probability of these cells.[12]

PLAUSIBILITY FROM PROBABILITY
$v \succeq^Q_{A,w} u$ iff $Pr_{A,w}([v]_Q) \geq Pr_{A,w}([u]_Q)$

The question Q is a parameter supplied by context. How it is determined is an urgent issue, since we do not assume that Q is always the same as the question intuitively under discussion.[13] That said, following Holguín (2022), we think the following is typically the case:

PRAGMATIC HYPOTHESIS
Q is the contextually relevant question for interpreting *think* and *confident* reports about an agent who is offering their opinion in answer to Q.

Armed with this hypothesis, we can offer a pragmatic account of ANSWER EQUIVALENCE. This is because our subjective probabilities systematically constrain the opinions we give when answering questions. For example, when discussing the question *who will win the tournament*, Joey will only guess that Gohar will win if there is no other player who he thinks is more likely to win. Building on observations like this, Holguín (2022) and Dorst and Mandelkern (2022) defend something very close to the following generalization[14]:

QUESTIONS AND OPINIONS
When, in w, a speaker A offers their opinion p in answer to a question Q (even if it's just their best guess), $p \in \$^Q_{A,w}$ (where $\succeq^Q$ is defined as above, and $\$^Q$ is defined from $\succeq^Q$).

Finally, note that DEGREES OF CONFIDENCE and DEGREES FROM PROBABILITY entail:

COINCIDENCE
$conf(A,w,p) = Pr_{A,w}(p)$ for all $p \in \$_{A,w}$

Taken together, PRAGMATIC HYPOTHESIS, QUESTIONS AND OPINIONS and COINCIDENCE entail ANSWER EQUIVALENCE (as well as SECOND GUESSING).

By way of illustration, consider the following example:

[12] This definition is borrowed from recent work on knowledge and belief in formal epistemology (Levi, 1967; Lin and Kelly, 2012; Leitgeb, 2017; Goodman and Salow, 2021, 2023; Holguín, 2022; Hong, 2023).

[13] Goodman and Salow (2021, 2023) give cases where the question intuitively under investigation (e.g., whether a coin is fair) is more coarse-grained than the question (e.g., how many time will the coin land heads) relevant to determining comparative plausibility for the purposes of giving adequate truth conditions for attitude reports.

[14] They also allow guesses to break ties in plausibility; see appendix A for how this can be accommodated.

(7) *Context*: Ann is about to flip a fair coin as many times as it takes until it lands heads. We ask her opinion about how many flips this will be. Consider some potential answers:
a. Just one.
b. At most five.
c. ??Not two.

Let us assume that Ann's subjective probabilities match the known chances: $Pr_{A,w}(\mathbf{n}) = .5^n$, where **n** is the proposition that the coin will take *n* flips to land heads. Where $Q = \{\mathbf{n} : n > 0\}$ (i.e., the question she is asked), it follows that $\$^Q_{A,w} = \{\bigcup\{\mathbf{m} : m \leq n\} : n > 0\}$. So *Ann thinks that the coin will be flipped just one time* is predicted to be true in this context (although in appendix A we show how this prediction can be relaxed). So (7a)-(7c) are then all propositions that Ann thinks are true. QUESTIONS AND OPINIONS is consistent both with her answering (7a) and with her answering (7b), since both $\mathbf{1} \in \$^Q_{A,w}$ and $\bigcup\{\mathbf{m} : m \leq 5\} \in \$^Q_{A,w}$. For these answers, Ann's subjective probabilities equal her degrees of confidence (in context), in keeping with ANSWER EQUIVALENCE. By contrast, her answering (7c) is ruled out by QUESTIONS AND OPINIONS (even though it is a proposition that she thinks is true), since $W \backslash \mathbf{2} \notin \$^Q_{A,w}$. So although Ann's subjective probability (.75) and her degree of confidence (.5) in this proposition come apart, this is no counterexample to ANSWER EQUIVALENCE.

10. Conclusion

Our account of degrees of confidence has three core features: (i) linking *confident* and *think* via COMMITMENT, (ii) rejecting BAYESIAN WHEN DEFINED in favor of the weaker LOWER BOUND, and (iii) appealing to worlds' comparative plausibility to explain how agents' degrees of confidence can sometimes be lower than their subjective probabilities (albeit elusively so). These features are all retained in the more sophisticated proposals explored in the appendices.

We conclude by highlighting a broader theoretical motivation for the present approach – namely, that it coheres with recent work on the interconnections between various doxastic and epistemic attitudes and with work on how these attitudes are shaped by comparative plausibility.

Goodman and Holguín (2023) argue that something like our account of degrees of confidence is needed to accommodate the normative connections between knowing, being sure, and assertion. In brief, they argue that *be sure* expresses the least degree of confidence *d* which is both subject to a knowledge norm (one should be confident to degree *d* only of things that one knows) and normative for assertion (one should assert only things that one confident of to degree at least *d*), and then argue that no degree of subjective probability satisfies this theoretical role.

Comparative plausibility, as defined in §9, also figures centrally in recent theories of knowledge (Goodman and Salow, 2021, 2023; Hong, 2023), assertion (Mandelkern and Dorst, 2022), and 'full' or 'outright' belief (Levi, 1967; Lin and Kelly, 2012), often combined with probability thresholds along the lines proposed here (Leitgeb, 2017; Cantwell and Rott, 2019; Goodman and Salow, ming).[15] Moreover, the general idea that an ordering of worlds is key to theorizing about the structure of knowledge and rational belief is a recurring theme in both traditional and

[15]As for the attitude expressed by *believe* in ordinary English, Hawthorne et al. (2016) and Rothschild (2020) argue that it is synonymous with *think* in the constructions considered here.

formal epistemology; see Goodman and Salow (2023, ming) and references therein.

These precedents are significant. One might have thought that, whatever its merits, our account of degrees of confidence is less parsimonious and more complicated than BAYESIANISM, since it uses more primitive notions and features more complicated definitions. We think this charge is mistaken, since the notions and ideas appealed to here have independent motivation in the surrounding literature. If anything, it is a demerit of BAYESIANISM that it treats *confident* so differently from typical treatments of other doxastic and epistemic attitudes.

Appendix

Three appendices show how our account of degrees of confidence can be generalized to apply to agents whose doxastic attitudes have a more complicated structure than we have so far allowed.

A. *Thinking* outside the box

So far we have assumed that the worlds doxastically accessible for an agent are all and only the worlds that are maximally plausible for the agent. Here we show how to generalize our account of degrees of confidence to agents who violate this PLAUSIBLE OPINIONS idealization.

We start by taking the set $D_A(w)$ of worlds doxastically accessible for A (in w) as primitive. While much of the exciting work in the study of doxastic attitudes lies in exploring further constraints on D (cf. Holguín's (2022) 'cogency' constraint that $D_A(w)$ be non-empty and closed under $\succ_{A,w}$) we impose no such constrains here in the interest of generality.

Next we use D to define a new system of spheres $\$^*$ from the old one $\$$:

$$\$^*_{A,w} = \{D_A(w)\} \cup \{D_A(w) \cup p : p \in \$_{A,w}\}$$

(Observe that $\$^* = \$$ whenever D is defined from $\succeq$ and Pr via DOXASTIC ACCESSIBILITY, DEGREES OF CONFIDENCE, and DEGREES FROM PROBABILITY.)

For any φ, let φ^* be the result of replacing all occurrences of $\$$ in φ with $\$^*$. The new account of degrees of confidence is DEGREES FROM PROBABILITY*. It preserves all of the earlier account's important features. In particular:

1. COMMITMENT, CLOSURE$_{\text{think/conf/deg/=}}$, and LOWER BOUND continue to hold.
2. COINCIDENCE* holds. As a result, ANSWER EQUIVALENCE and SECOND GUESSES are still predicted given QUESTIONS AND OPINIONS* (which is the link between thinking and guessing that is actually defended in Holguín (2022); cf. note 14).

B. Relaxing CLOSURE

As mentioned in §5, the preface paradox raises an important challenge for the various CLOSURE principles discussed above. To take Makinson's (1965) original example, it seem possible that every claim in a book be one that its author is confident is true without the conjunction of all of the claims in the book being something that the author is confident is true. This appendix shows how to modify our account of degrees of confidence to accommodate such (purported) failures of CLOSURE$_{\text{conf}}$. Crucially, and unlike BAYESIANISM, doing so is compatible with endorsing the pattern of judgments about (4a) used to motivate it and similar instances of CLOSURE$_{\text{conf}}$.

The basic idea is that preface-style cases involve incomparabilities in comparative plausibility: cases where neither $v \succeq_{A,w} u$ nor $u \succeq_{A,w} v$. For example, where p and q are two claims in A's book that concern different questions, v could be a world where p is true but q is false, and u a world where q is true but p is false. One way of developing this idea would be to derive comparative plausibility from a set of questions $\mathscr{Q}$, so that $v \succeq^{\mathscr{Q}}_{A,w} u$ iff $v \succeq^{Q}_{A,w} u$ for all $Q \in \mathscr{Q}$, with each $\succeq^{Q}$ derived from Pr according to PLAUSIBILITY FROM PROBABILITY as before; see Goodman and Salow (2023). In the interest of simplicity, here we will simply show how failures of CLOSURE$_{\text{conf}}$ can arise organically and non-disruptively by allowing incomparabilities in comparative plausibility – i.e., by relaxing the requirement that $\succeq_{A,w}$ be a *total* pre-order.

Here is an implementation of this idea that departs as little as possible from the account in the main text. The treatment of *think* is unchanged, and PLAUSIBLE OPINIONS continues to hold. We make two changes. First, we adopt a more ecumenical definition of $\$_{A,w}$:

$$\$_{A,w} = \{p : \{v : \forall u(u \not\succ_{A,w} v)\} \subseteq p \text{ and } \forall v \in W, \forall u \in p(v \succeq_{A,w} u \Rightarrow v \in p)\}$$

This definition agrees with the previous one whenever $\succeq_{A,w}$ is a total pre-order. Intuitively, a sphere is now any set of worlds that (i) is entailed by the strongest proposition that the agent thinks is true (i.e., it contains all of the worlds that the agent takes to be maximally plausible), and (ii) contains every world that is at least as plausible as any other it contains.

Second, we adopt CONFIDENCE FROM PROBABILITY SPHERES as a definition of degrees of confidence, bypassing the intermediary C. Let us survey some features of the resulting account:

1. COMMITMENT, LOWER BOUND, and COINCIDENCE continue to hold.
2. CLOSURE$_{\text{think}}$ continues to hold.

   Note: given EXTREME WEAKENESS, it is not obvious that the same considerations that motivate rejecting CLOSURE$_{\text{conf}}$ must extend to CLOSURE$_{\text{think}}$.
3. CLOSURE$_{\text{conf/deg/=}}$ can now fail; an example is given in a footnote.[16]

How widespread are failures of CLOSURE$_{\text{CONF/DEG}}$? Examples like (4a) show that they are less common than we would expect from BAYESIAN WHEN DEFINED. Moreover, such cases suggest that, when considering a question like *how much one weighs*, CLOSURE$_{\text{CONF/DEG}}$ holds for propositions about this question. Let us briefly explore this idea. First, some definitions:

- Say that p *concerns* Q iff $p = \bigcup X$ for some $X \subseteq Q$.
- For any partition Q of W and $p, q \in Q$, say that $p \succeq_Q q$ iff $\forall w \in p \exists v \in q(w \succeq v)$.
- Say that $\succeq$ *editorializes* Q iff the preorder $\succeq_Q$ on Q is total (i.e., no incomparabilities).

[16] $W = \{w, v, u\}$, $w \succ v$, $w \succ u$, $v \not\succeq u$, $u \not\succeq v$, $Pr(\{w\}) = .8$, $Pr(\{v\}) = Pr(\{u\}) = .1$. Then $conf(\{w,v\}) = conf(\{w,u\}) = .9 > .8 = conf(\{w\})$.

The present account validates a qualified version of CLOSURE$_{\text{conf}}$ (and likewise of CLOSURE$_{\text{deg}}$):[17]

> RESTRICTED CLOSURE$_{\text{conf}}$
> Suppose $\succeq_{A,w}$ editorializes Q and both p and q concern Q. Then, in w, if A is confident that p and confident that q, then A is confident that *p and q*.

Finally, we continue to predict ANSWER EQUIVALENCE given QUESTIONS AND OPINIONS provided we accept a suitably modified pragmatic hypothesis:

> PRAGMATIC HYPOTHESIS′
> The contextually relevant plausibility order $\succeq$ when describing an agent who is offering their opinion in answer to Q is such that, for all w, v, $[w]_Q \succeq_Q [v]_Q$ iff $w \succeq^Q v$.

This says that the plausibility order, when lifted to cells of Q, is congruent with the order defined by PLAUSIBILITY FROM PROBABILITY. It holds given natural assumptions (see footnote).[18]

C. Combining both generalizations

This appendix shows how to combine the ideas from the previous two appendices: modeling agents who violate both PLAUSIBLE OPINIONS and CLOSURE$_{\text{conf/deg/=}}$. Again in the interest of generality, we also allow for failures of CLOSURE$_{\text{think}}$. This requires moving beyond Hintikka semantics for *think*. Instead of there being a single strongest proposition $D_A(w)$ that (in w) A thinks is true, we appeal to a set $\mathscr{D}_A(w)$ of all of the strongest propositions that (in w) A thinks are true. Formally, $\mathscr{D}_A(w) \subseteq \mathscr{P}(W)$ such that $p \not\subseteq q$ for any distinct $p, q \in \mathscr{D}_A(w)$. Then:

$$[\![A \text{ thinks that } \varphi]\!](w) = 1 \text{ iff } p \subseteq [\![\varphi]\!] \text{ for some } p \in \mathscr{D}_{[\![A]\!]}(w).$$

Again, it may be natural to impose further constraints on $\mathscr{D}$ (e.g., that $p \in \mathscr{D}_A(w)$ only if p is non-empty and closed under $\succ_{A,w}$). But for generality we again impose no such constraints.

We now modify the definition of \$ from the previous appendix by replacing $\{v : \forall u(u \not\succ_{A,w} v)\}$ (i.e. $D_A(w)$) with an existentially bound variable ranging over members of $\mathscr{D}_A(w)$:

$$\$_{A,w} = \{p : \exists q \in \mathscr{D}_A(w)(q \subseteq p) \text{ and } \forall v \in W, \forall u \in p(v \succeq_{A,w} u \Rightarrow v \in p)\}$$

As before, we then adopt CONFIDENCE FROM PROBABILITY SPHERES as a definition of *conf*.

The main change with this account is that CLOSURE$_{\text{think}}$ can now fail, and as a result so can RESTRICTED CLOSURE$_{\text{conf}}$. However, the slightly weaker principle remains valid:

> VERY RESTRICTED CLOSURE$_{\text{conf}}$
> Suppose $\succeq_{A,w}$ editorializes Q and both p and q concern Q. Then, in w, if A is confident that p and confident that q and thinks that p and q, then A is confident that *p and q*.

[17] Proof sketch: the key fact is that, if $\succeq$ editorializes Q and p and q both concern Q, then either every p-entailing sphere is q-entailing or vice versa.

[18] Suppose that $\succeq$ is determined as suggested earlier: by universally generalizing over the preorders which are probabilistically generated from the members of a contextually determined set of questions $\mathscr{Q}$. If $\mathscr{Q}$ includes the salient question Q, then we can show that PRAGMATIC HYPOTHESIS′ holds so long as Q is logically independent of the question that results from combining all of the remaining questions in $\mathscr{Q}$.

In other respects the resulting account leaves the situation largely unchanged as far as the principles discussed above are concerned.

We conclude by observing an interesting edge case that arises when comparative plausibility maximally incomparable, in the sense that $v \not\succeq_{A,w} u$ for any $v \neq u$. BAYESIAN WHEN DEFINED then holds, since every set of worlds trivially contains every world at least as plausible as any other world that it contains. This raises the intriguing prospect that BAYESIAN WHEN DEFINED may be true in some contexts after all, without requiring any ambiguity in *confident*.

References

Cantwell, J. and H. Rott (2019). Probability, coherent belief and coherent belief changes. *Annals of Mathematics and Artificial Intelligence 87*(3), 259–91.

Cariani, F., P. Santorio, and A. Wellwood (forthcoming). Confidence reports. *Semantics & Pragmatics*.

Clarke, R. (2024). Strong belief is ordinary. *Episteme 21*(3), 773–793.

de Finetti, B. (1937). La prévision: ses lois logiques, ses sources subjectives. *Annales de l'institut Henri Poincaré 7*, 1–68.

Dorst, K. and M. Mandelkern (2022). Good guesses. *Philosophy and Phenomenological Research 105*(3), 581–618.

Eriksson, L. and A. Hájek (2007). What are degrees of belief. *Studia Logica 86*(2), 185–215.

Foley, R. (1993). *Working Without a Net: A Study of Egocentric Epistemology*. Oxford: Oxford University Press.

Genin, K. (2019). Full & partial belief. In R. Pettigrew and J. Weisberg (Eds.), *The Open Handbook of Formal Epistemology*, pp. 437–498. PhilPapers Foundation.

Goodman, J. and B. Holguín (2023). Thinking and being sure. *Philosophy and Phenomenological Research 106*(3), 634–54.

Goodman, J. and B. Salow (2021). Knowledge from probability. In J. Y. Halpern and A. Perea (Eds.), *Theoretical Aspects of Rationality and Knowledge 2021 (TARK 2021)*, pp. 171–186.

Goodman, J. and B. Salow (2023). Epistemology normalized. *Philosophical Review 132*(1), 89–145.

Goodman, J. and B. Salow (forthcoming). Belief revision normalized. *Journal of Philosophical Logic*.

Hawthorne, J., D. Rothschild, and L. Spectre (2016). Belief is weak. *Philosophical Studies 173*(5), 1393–404.

Hoek, D. (forthcoming). Minimal rationality and the web of questions. In D. Kindermann, P. van Elswyk, and A. Egan (Eds.), *Unstructured Content*. Oxford University Press.

Holguín, B. (2022). Thinking, guessing and believing. *Philosophers' Imprint 22*(6), 1–34.

Holton, R. (2014). Intention as a model for belief. In M. Vargas and G. Yaffe (Eds.), *Rational and Social Agency: Essays on the Philosophy of Michael Bratman*. Oxford University Press.

Hong, F. (2023). *Uncertain Knowledge*. Ph. D. thesis, University of Southern California.

Jeffrey, R. C. (1965). *The Logic of Decision*. New York, NY, USA: University of Chicago Press.

Joyce, J. (2009). Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In F. Huber and C. Schmidt-Petri (Eds.), *Degrees of Belief*, pp. 263–297. Srpinger.

Leitgeb, H. (2017). *The Stability of Belief: How Rational Belief Coheres with Probability*. Oxford: Oxford University Press.

Levi, I. (1967). *Gambling with Truth: An Essay on Induction and the Aims of Science*. London, England: MIT Press.

Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.

Lin, H. and K. Kelly (2012). Propositional reasoning that tracks probabilistic reasoning. *Journal of Philosophical Logic 41*(6), 957–81.

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

Mandelkern, M. and K. Dorst (2022). Assertion is weak. *Philosophers' Imprint 22*(19), 1–20.

Moss, S. (2018). *Probabilistic Knowledge*. Oxford: Oxford University Press.

Ramsey, F. (1926 [1931]). Truth and probability. In R. Braithwaite (Ed.), *Foundations of Mathematics and Other Logical Essays*, pp. 156–98. London: Kegan, Paul, Trench, Trubner & Co. Ltd.

Rothschild, D. (2012). Expressing credences. *Proceedings of the Aristotelian Society 112*, 99–114.

Rothschild, D. (2020). What it takes to believe. *Philosophical Studies 177*, 1345–62.

Russell, J. S. and J. Hawthorne (2016). General dynamic triviality theorems. *Philosophical Review 125*(3), 307–339.

Savage, L. J. (1954). *The Foundations of Statistics*. New York: Wiley.

Titelbaum, M. (2019). Precise credences. In R. Pettigrew and J. Weisberg (Eds.), *The Open Handbook of Formal Epistemology*, pp. 1–55. PhilPapers Foundation.

Williamson, T. (forthcoming). Knowledge, credence and the strength of belief. In A. Flowerree and B. Reed (Eds.), *Expansive Epistemology: Norms, Action, and the Social World*. London: Routledge.

Yalcin, S. (2012). Bayesian expressivism. *Proceedings of the Aristotelian Society 112*, 123–160.

Yalcin, S. (2018). Belief as question sensitive. *Philosophy and Phenomenological Research 97*(1), 23–47.