A Conflict between Indexical Credal Transparency and Relevance Confirmation

Joel Pust*†

According to the probabilistic relevance account of confirmation, E confirms H relative to background knowledge K just in case P(H/K & E) > P(H/K). This requires an inequality between the rational degree of belief in H determined relative to two bodies of total knowledge that are such that one (K&E) includes the other (K) as a proper part. In this article, I argue that it is quite plausible that there are no two possible bodies of total knowledge for ideally rational agents meeting this requirement. Hence, the positive relevance account may have to be rejected.

1. Positive Relevance Confirmation and Total Evidence Sets. According to the well-known probabilistic relevance account of confirmation, E confirms H just in case P(H/E) > P(H) for some suitable probability function.[1] This account is justified by the intuitively plausible claim that E is evidence, in one relatively clear pretheoretic sense, for H just in case knowledge of E makes it rational to be more confident that H is true (Maher 1996). Of course, whether E is evidence for H often depends on what other evidence or background knowledge one possesses. So, the more precise view is one according to which E confirms H, given background knowledge K, just in case P(H/K & E) > P(H/K).

Received January 2020; revised July 2020.

*To contact the author, please write to: Department of Philosophy, University of Delaware; e-mail: jpust@udel.edu.

†For helpful comments on previous drafts of this article, I thank Julia Staffel, Konstantin Genin, Aaron Bronfman, Kevin Dorst, Paul Draper, Kaila Draper, two anonymous referees, and audience members at the 2017 Central Division Meeting of the American Philosophical Association and the 2019 Meeting of the Society for Exact Philosophy.

1. Some philosophers identify relevance confirmation with the holding of the inequality noted above on each of a set of probability functions (Maher 1996, 163). This wrinkle does not affect what follows.

Philosophy of Science, 88 (July 2021) pp. 385–397. 0031-8248/2021/8803-0001$10.00
Copyright 2021 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/713907 Published online by Cambridge University Press

There are, to be sure, a wide variety of possible views on the nature of the probability function appearing in the aforementioned inequality. Some objective Bayesians, such as Carnap (1962) and, in their own way, Skyrms (1986) and Hawthorne (2005), maintain that P(H/E) represents, in the first instance, the objective degree of support that E provides for H and, in the second instance, the degree of belief in H that is therefore rational for any agent whose total evidence is E.[2] They thereby ground the positive relevance account in the aforementioned intuition that E is evidence for H, given K, just in case an agent whose total evidence is K&E should be more confident that H than one whose total evidence is K alone.

As the problem of old evidence (Glymour 1980) shows that the probabilities in the positive relevance account cannot simply be an agent's credences—on pain of making it impossible for E to be evidence when E is known—many subjective or permissive Bayesians of various sorts take P(H/E) to represent the relevant agent's historical or counterfactual degree of belief in H, conditional on E, either when E was not known or if E were to be not known. They identify positive relevance confirmation of H, given K, by E, with an increase in the relevant agent's confidence in H (Howson and Urbach 2006) upon actually or counterfactually adding knowledge of E to K. Such Bayesians go on either to eschew any claim about the rationality of such an increased credence and thereby equate relevance confirmation by E with E's causal power to increase an agent's confidence in H or, if they endorse a diachronic rational requirement to conditionalize, to claim that such an increase is rational only relative to the agent's historical or counterfactual credences (which credences are likely constrained only by probabilistic coherence).[3]

On all these views, relevance confirmation involves a relation between two possible bodies of total knowledge or evidence, where one (K&E) includes the other (K) as a proper part. In this article, I show that no two possible bodies of total knowledge for an ideally rational agent can satisfy this condition if a plausible requirement on rational agents' self-knowledge is correct. Hence, we are forced to make an unpleasant choice between this

2. A similar sort of "objectivism" is also to be found in Horwich when he suggests that E confirms H if and only if "reason requires that with background assumptions B, anyone's conditional degree of belief in H given E, should be greater than his unconditional degree of belief in H" (1982, 52).

3. See Maher (1996, 151–52) for reasons to think that subjective Bayesians are not analyzing the confirmation relation but are instead analyzing what it is for a person to *take* E to be such that, if it were evidence, it would be evidence for H. (A similar charge is made by Horwich [1982, 51]). See Kaplan (1996, 51) for just such an avowedly deflationary account, although one that tries to extend its application by claiming that one is never entitled to certainty in a contingent proposition.

https://doi.org/10.1086/713907 Published online by Cambridge University Press

plausible requirement of self-knowledge and the plausible positive relevance account of evidence.

2. Rationality and Indexical Credal Omniscience. Temporally indexical opinions are those that cannot be properly characterized without the use of the temporal indexical "now" or its relatives. Personally indexical opinions are those that cannot be properly characterized without the use of the personal indexical "I" or its relatives. I simply assume here that there are indeed "essentially indexical" opinions of both sorts (Perry 1979). Temporally indexical *self*-knowledge is a special sort of temporally and personally indexical knowledge—knowledge of one's own current mental states *as* one's own current mental states. Especially significant in the current context is essentially temporally-cum-personally indexical certainty regarding one's current credences. Letting "$P_{s,t}(H) = x$" represent that agent s's credence in proposition H at time t is x and letting "$P_{s,t}(P_{my,now}(H) = x$" represent that s's credence at t for the claim that her own current credence in H is x (where the time and agent are both considered in an essentially indexical manner), consider the following requirement on ideally rational agents.

Indexical Credal Transparency: $P_{s,t}(H) = x$ only if $P_{s,t}(P_{my,now}(H) = x) = 1$.

Given probabilistic coherence at t (which I here assume is required by ideal rationality) and some credence in H, this entails its converse.[4]

Indexical Credal Infallibility: $P_{s,t}(P_{my,now}(H) = x) = 1$ only if $P_{s,t}(H) = x$.

I will refer to the conjunction of these doctrines as Indexical Credal Omniscience, but it is the former transparency doctrine that generates the problem of interest in this article.[5]

4. Note that I assume that ideally rational agents have some credence in each of the relevant propositions. Some, like Titelbaum (2013), depart from this idealization of Bayesian orthodoxy and hold that a rational agent cannot conflict with ideal rationality in failing to have any credence at all in a proposition. Still, Titelbaum holds that such agents can be required by ideal rationality to have a particular credence in a proposition if they have any credence in the proposition. He would likely accept the claim that if $P_{s,t}(H) = x$ then s is rationally committed to $(P_{my,now}(H) = x) = 1$ if s has any credence at all in $P_{my,now}((H) = x))$ at t. Such "rational commitments" will likely appear along with K or E&K on the right side of the relevant conditional probabilities in the positive relevance account, and so the problem developed in the next section will arise in a slightly different way.

5. The arguments of this article would go through given a weaker doctrine according to which ideal rationality requires merely certainty at each time as to whether one is certain at that time: $P_{s,t}(H) < 1$ only if $P_{s,t}(P_{my,now}(H) < 1) = 1$, and $P_{s,t}(H) = 1$ only if $P_{s,t}(P_{my,now}(H) = 1) = 1$. (Thanks to Kevin Dorst for pointing this out.)

https://doi.org/10.1086/713907 Published online by Cambridge University Press

Notice that Indexical Credal Transparency should be distinguished from a similar doctrine, Credal Transparency, according to which $P_{s,t}(H) = x$ only if $P_{s,t}(P_{s,t}(H) = x) = 1$. Credal Transparency does not require of an agent any essentially indexical grasp of herself and the current time. Hence, it would be satisfied by one's being certain that some agent who one was unaware was oneself at a time one was unaware was the current time had the credence in H that one in fact then had. I do not maintain that Credal Transparency is a requirement of ideal rationality. One might be ideally rational while lacking knowledge of one's identity and of the current time under suitable nonindexical modes of presentation. So, even if I am required by ideal rationality to be certain that my current credence (grasped in the essentially indexical manner) in a given proposition is what it is, I am not required to be certain that the author of this article has such a credence at 11:01 p.m. on April 3, 2018, even if I am the author and it is 11:01 p.m. on April 3, 2018.

There are a number of reasons to hold that an ideally rational agent must be indexically credally omniscient—certain, at a given time, of all and only the true claims about her (then) current credences (as her current credences). As I am concerned in this article merely to argue that such a doctrine is quite plausible and conflicts with the positive relevance account, I will here simply sketch some of the main reasons, leaving a detailed defense of the doctrine and responses to objections for another occasion.[6] It should, however, be noted that some of the defenses of Indexical Credal Transparency below are variations on the defenses Bayesians give of the synchronic norm of probabilistic coherence, and so those Bayesians who wish to reject credal omniscience face the dialectical challenge of doing so while retaining a cogent defense of probabilistic coherence.

The first reason for endorsing Indexical Credal Transparency is simply that the claim is intuitively plausible. Even if we mere mortals cannot attain it, such knowledge of one's own mind seems a requirement of ideal rationality. Indeed, Sobel (1987, 69) suggests that "the extent of a person's self-possession is . . . a partial determinant of his intellectual self." This line of thought is strengthened by the observation that rationality surely requires some degree of fit between an agent's first-order credences at a time and her higher-order personally-cum-temporally indexical credences regarding her credences at that time. No ideally rational agent could, it seems, have a wholly mistaken conception of her own first-order credences under the requisite mode of presentation. This being so, it seems plausible that ideal rationality requires perfect self-knowledge of the relevant sort. In this respect, such knowledge is quite unlike other contingent knowledge, as a lack of other contingent knowledge is compatible with perfect rationality.

6. In particular, a full defense would have to consider the probabilistic version of Williamson's (2008) argument against luminosity (here called "transparency").

https://doi.org/10.1086/713907 Published online by Cambridge University Press

A further reason to think the doctrine is intuitively plausible is the ease with which many discussions of "common knowledge" in game-theoretic settings assume that ideally rational agents are certain about their own credences. Indeed, such introspective omniscience seems required for many of the proofs of common knowledge theorizing, as it is required for an agent to make use of her knowledge that her own knowledge is also had by others in the game. While some of these results may be obtained without assuming Indexical Credal Transparency, it remains unclear what can be so obtained, and attempts to do so are clearly dealing with intuitively less than fully rational agents.

Another reason for embracing Indexical Credal Transparency is the fact that it seems presupposed by the intuitive verdict on so-called Thomason cases—counterexamples to the claim that a rational conditional probability is identical to a rational probability given knowledge of the condition. One example of such a case is the intuition that P(I am certain that the NSA has surveilled me/The NSA has surveilled me) is quite low but P(I am certain that the NSA has surveilled me) is very high (perhaps 1) if I am certain that the NSA has surveilled me. That intuition, if an intuition about rationally required credence rather than what credence would be typical, seems simply to presuppose that rational epistemic probability obeys Indexical Credal Transparency. Why else would certainty in the claim about the NSA rationally require a high probability in the claim regarding my own certainty about the NSA?

A more theoretical reason for thinking that Indexical Credal Transparency is a requirement of ideal rationality is the fact that there is a synchronic Dutch Book argument for such a requirement because a Dutch Book can be made against any agent failing Indexical Credal Transparency (but not against those probabilistically coherent agents possessed of it) by a bookie who knows only what the agent's credences currently are (Sobel 1987; Milne 1991). Such a bookie can guarantee a profit simply by offering a bet based on the agent's mistaken higher-order credence, and the agent is guaranteed, given the agent's first-order credence, to lose that bet. Hence, if, as many subjectivists maintain, synchronic Dutch Book arguments provide a justification for probabilism, so also do they provide a justification for a requirement of Indexical Credal Transparency.[7]

A final theoretical reason for endorsing Indexical Credal Transparency is to be found in the accuracy-based framework of epistemic utility theory (Joyce 1998; Pettigrew 2016). If one violates Indexical Credal Transparency, then one's credal state is a priori guaranteed to be less accurate than a

7. Christensen (2007) and Mahtani (2015) argue that the relevant Dutch Book argument is not one that reveals irrationality. I contest Mahtani's argument in my "Dutch Books and Logical Form" (Pust 2021).

https://doi.org/10.1086/713907 Published online by Cambridge University Press

credal state that is otherwise identical but satisfies Indexical Credal Transparency.[8] This shows that the accuracy domination induced by failures of Indexical Credal Transparency is, again, wholly internal to the system of credences and so constitutes a genuine failing of rationality rather than a mere failure to know some contingent truth or other.

Now, both the Dutch Book and accuracy arguments for Indexical Credal Transparency just sketched turn on holding fixed, as seems natural, the fact that an agent has the credences she does in assessing whether the agent's credal state sanctions a sure loss set of wagers or is accuracy dominated by another set of credences.[9] Caie (2013), however, argues, by appeal to a self-referential proposition, that both arguments for probabilism fail on this natural understanding. If his arguments are sound, then friends of probabilism (which is here assumed) must accept some restriction on the situations in which accuracy arguments and Dutch Book arguments can support probabilism, or repudiate such arguments altogether and find support for probabilism elsewhere, or reject the notion that we should hold fixed a credal state in assessing its accuracy or loss inducing character at each world. The first option would fit naturally with endorsing a similar limitation on Indexical Credal Transparency and, so long as the exceptions are suitably restricted, leave my argument largely unaffected. The second and third options would leave Indexical Credal Transparency without the support of accuracy arguments or Dutch Book arguments and reliant on other support. Those who take the third option hold that the relevant worlds for assessing whether a set of bets is guaranteed to be a losing set or assessing whether another set of credences dominates a given set include worlds where the agent's credences are different or, more peculiarly, where the agent at issue does not exist.

While a full assessment of these complex issues is quite clearly beyond the scope of this article, let me here simply note that the third option suffers significant implausibility. First, such accounts implausibly allow that an agent may be ideally rational even though she is certain (or close to certain) that she does not exist or that there are no credal states. Second, assessing the rationality of an agent's credal state in part by its inaccuracy at worlds where there are no agents and credal states or its sanctioning of wagers at worlds where no wagers are sanctioned by any credences seems to be of dubious coherence. While we can assess whether a set of wagers is a winning or losing set if settled relative to what is true at a world with no agents or credences and we can assess the distance from the truth of a function from propositions to

8. Although she rejects Credal Transparency (and, presumably, Indexical Credal Transparency) as a rational requirement, Carr (2017) perspicuously notes that it is a consequence of many ways of developing epistemic utility theory.

9. Thanks to an anonymous referee for noting the controversy surrounding this assumption and suggesting that I address it.

https://doi.org/10.1086/713907 Published online by Cambridge University Press

numbers in the [0, 1] interval at a world with no agents or credences, in both cases we are clearly not assessing any agent's credences at that world. It seems to me, however, that the core idea behind both Dutch Book and accuracy domination arguments is the idea that a set of credences is rationally defective when an agent is in a position to know a priori that it sanctions a set of losing bets or is more inaccurate than an alternative set of credences, no matter what the extracredal world is like.

Even if one endorses the third option above and rejects construals of Dutch Book and accuracy arguments on which they directly support indexical credal transparency, the ability of such arguments to support probabilism may, in fact, presuppose credal transparency. This is because accounts of the rationality of various credal states seem to presuppose that one has access to one's credal state. After all, exclusively a priori knowledge that a given credal state sanctions a set of wagers guaranteed to be a losing set is insufficient to yield the result that one is irrational if one occupies that state. One would need, in addition, reason to think that the loss sanctioning credal state was one's own and Indexical Credal Transparency would capture this sort of access. Similarly, knowledge that a given credal state is accuracy dominated is insufficient to yield the result that one is irrational in occupying that state. One would again need, in addition, reason to think that the dominated credal state was one's own, knowledge of exactly the sort mandated by Indexical Credal Transparency.[10]

It should be emphasized that the sort of introspective omniscience requirement supported by these various considerations is limited to the doubly indexical sort indicated earlier. What is most intuitively plausible is that ideal rationality requires that one is certain of one's own current opinions as one's own current opinions, regardless of whether one has any nonindexical knowledge of the current time or of one's own identity. What the Dutch Book and accuracy domination arguments reveal is a difficulty wholly internal to one's current credences, one the grasping of which requires of one no nonindexical grasp of the current time or of one's identity, both of which might be thought to require contingent knowledge of nonindexical nonmental facts, which knowledge is not entailed by ideal rationality. These considerations do not support a requirement of credal omniscience regarding one's nonindexically specified opinions, a diachronic requirement of self-knowledge, or a synchronic general requirement of knowledge regarding personally indexical truths.

3. The Conflict. As previously indicated, the positive relevance account would have us determine whether E confirms H, relative to K, by comparing

10. Here it is worth noting that Konek and Levinstein's recent (2019) account of epistemic utility theory, while not defending Indexical Credal Transparency on accuracy grounds, assumes its truth (secs. 4.2–4.4) in order to yield plausible verdicts about cases in which one knows the truth of a proposition is dependent on one's own credences.

https://doi.org/10.1086/713907 Published online by Cambridge University Press

an epistemic situation in which an agent's total knowledge is K to an epistemic situation in which her total knowledge is K&E. What is required, in other words, is that there are two sets of possible certainties for a rational agent that are such that one is a proper subset of the other.

If one were uncertain at t of E then one's credence at t in E would be some n less than 1. Given Indexical Credal Transparency, in such a situation, one would have a credence of 1 that one now has such a credence in E—$P_{s,t}(P_{my,now}(E) = n) = 1$. Consider, then, E's being certain for one, either at the same time, t (the synchronic case), or a subsequent one, $t+$ (the diachronic case). Take the synchronic case first. If E were instead certain at t, then one would, given Indexical Credal Transparency, instead be certain that one now has a credence of 1 in E, $P_{s,t}(P_{my,now}(E) = 1) = 1$. This yields a contradiction in the set of certainties as $P_{s,t}(P_{my,now}(E) = n) = 1$ and $P_{s,t}(P_{my,now}(E) = 1) = 1$ are inconsistent, given that $n \neq 1$. (The same result follows from the fact that, when E is not certain, $P_{s,t}(P_{my,now}(E) = 1) = 0$, and when E is certain, $P_{s,t}(P_{my,now}(E) = 1) = 1$, which contradict each other.) Hence, the left side of the inequality at issue will be undefined.

Consider, now, the diachronic case. If E were to be certain at some later time, $t+$, one would (then) know that one (then) has credence 1, $P_{s,t+}(P_{my,now}(E) = 1) = 1$. Whether the set of certainties at t and at $t+$ contradict each other depends on whether the proposition expressed at t by $P_{my,now}(E) = n$ contradicts that expressed at $t+$ by $P_{my,now}(E) = 1$, given that $n \neq 1$. If so, then the above problem remains. If not, it must be because the proposition expressed at a time by the indexical claim is not even graspable at distinct times.[11] In that case, there is no contradiction between the t and the $t+$ certainty sets, but it remains the case that the two diachronically distinct certainty sets cannot be such that one is a proper subset of the other.

I have framed the difficulty in terms of the *qualitative* relevance notion of confirmation, but we should note that it applies equally to all reasonable proposals regarding the proper *quantitative* measure of relevance confirmation, that is, the degree to which E confirms H. These include the probability difference measure (Huber 2008), according to which the measure is $P(H/K \& E) - P(H/K)$; the probability ratio measure (Milne 1996), according to which it is $P(H/K \& E) \div P(H/K)$; and the likelihood ratio measure (Zalabardo 2009), according to which it is $P(E/H \& K) \div P(E/\sim H \& K)$. The problem applies equally to such measures because they entail that the positive relevance account is the correct qualitative measure, merely imposing additional requirements to yield suitable quantitative measures.[12]

11. I argue (in Pust 2012) that this is a consequence of the three main views on temporally indexical credence.

12. Moreover, the problem also afflicts the qualitative Likelihood Principle (Edwards 1972), according to which E favors (in the likelihoodists' proprietary sense) H1 over

https://doi.org/10.1086/713907 Published online by Cambridge University Press

In essence, the problem is that, given the requirement of Indexical Credal Transparency, it is impossible for there to be two sets of certainties that are such that (*a*) they could each constitute the total certainty set of an ideally rational agent and (*b*) one is a proper subset of the other. Considered at two distinct times, an ideally rational agent must have at each time higher-order temporally indexical certainties that, depending on one's views regarding temporally indexical statements, were either not even graspable at the other time or contradict her certainties at the other time. Considered counterfactually or merely possibly at a single time, she must have higher-level certainties that contradict her actual higher-level certainties at that time. So, no two possible bodies of total knowledge for an ideally rational agent can be such that one is a proper subset of the other. Just such a possibility, however, is what is required for an application of the positive relevance account.

4. Revising the Positive Relevance Account? It may be thought that the problem just outlined is simply an implication of the now widely recognized fact that standard conditionalization principles cannot properly accommodate temporally self-locating or temporally indexical credences and that existing proposals for alternative diachronic principles can simply be adapted to solve the current problem. Consider an agent watching a clock that she is certain is accurate, from 09:00 to 09:01. When she learns that it is now 09:01, she might be thought to come to know a proposition that she was previously certain was false and to be certain that a previously certain proposition is false. If that is so, a straightforward application of the principle of conditionalization would imply, incorrectly, that all her subsequent credences should be undefined.

A number of proposals have been offered for how a rational agent's credences ought to evolve over time in light of her changing essentially indexical knowledge.[13] However, they are all unhelpful with the current problem as they all begin from the assumption that what one learns, in learning a temporally indexical proposition, is something in which one had either zero credence or no credence whatsoever at the prior moment. More precisely, in accepting that there is essentially temporally indexical knowledge, they hold that there are credal states that are essentially synchronic and such that they cannot be had by one and the same agent at distinct times or, alternatively, such that a rational agent can be certain that an indexical statement is false at one time and certain that it is true at a subsequent one (or vice versa). As I have shown, this, along with the Indexical Credal Transparency, implies that

H2 just in case $P(E/H1 \,\&\, K) > P(E/H2 \,\&\, K)$, provided that K includes all rationally required knowledge and E represents all the knowledge gained with a new certainty.

13. For a helpful taxonomy of these many proposals and critical discussion, see Titelbaum (2016).

https://doi.org/10.1086/713907 Published online by Cambridge University Press

it is simply impossible to capture an ideally rational agent's total knowledge at two distinct times in a manner that yields the result that the agent's total knowledge at one time is a proper subset of the agent's total knowledge at another time.[14]

However, guided by some of these accounts of rational diachronic change, one might attempt to preserve the core of the positive relevance account by simply excluding the problematic temporally indexical self-knowledge from its scope so that K&E could be a monotonic expansion of K. Letting 'E^*' and 'K^*' stand for E and K when we exclude all temporally-cum-personally indexical claims about the agent's credences from E and K, the suggestion would be that $P(H/E^* \,\&\, K^*) > P(H/K^*)$ can stand in the relevant relation even if no ideally rational agent could be bereft of additional temporally indexical knowledge of her own credences and so neither K^* nor $K^*\&E^*$ could constitute the entirety of an ideally rational agent's knowledge.

Unfortunately, this proposal unacceptably unmoors the positive relevance account from its intuitive justificatory rationale. That rationale, recall, is that E is evidence for H, given K, when knowledge of E makes it rational to be more confident of H. Even those subjectivists who eschew claims about the rationality of one's degree of confidence in H treat E as evidence for H when one's degree of confidence in H would be or was lower in the absence of knowledge of E, holding K constant. Hence, even if E^* and K^* have their content restricted in the manner under consideration, E^* (alone) cannot be evidence unless a positive answer can be given to the counterfactual question, "How confident of H would *s* be if *s* were ideally rational and certain of E^* but lacked accompanying indexical higher-order credence?" If ideal rationality does indeed require Indexical Credal Transparency, then, assuming that E^*'s being evidence cannot require a failure of ideal rationality, the counterfactual is a counterpossible and so, on standard semantics, vacuously or trivially true. Hence, this rejoinder fails.[15]

14. The same point regarding the impossibility of *diachronic* monotonic increases in evidence could be made given a rational requirement to be certain at every time of the distinct proposition one would at that time express with "I exist now" (Pust 2007). However, such a case will not suffice to preclude a wholly *synchronic* monotonic increase in total knowledge because the certainty at issue would remain unchanged by the addition of further nontemporally indexical certainties.

15. Proponents of the so-called grounding relation in contemporary metaphysics might claim that knowledge of E* (but not the higher-level knowledge of one's knowledge of E* that, if one is fully rational, necessarily accompanies E*) rationally *grounds* or makes reasonable a higher confidence in H. Grounding, it should be noted, is alleged to be a relation that allows one of two (or more) necessarily co-obtaining states of affairs to be the one that "makes it the case" that another of the necessarily co-obtaining states of affairs obtains. However, it has been plausibly argued that the sort of grounding at issue yields

https://doi.org/10.1086/713907 Published online by Cambridge University Press

Setting aside the problem with invoking counterpossibles, notice that this proposal would also give us, at best, an account of what it is for E^* to be evidence for H, relative to an epistemic situation that is not a possible situation of an ideally rational agent. The envisaged restriction on the content of K^* yields an account on which H has a certain probability relative to K^* and relative to $K^*\&E^*$ where K^* and $K^*\&E^*$ do not capture possible certainty sets of an ideally rational agent. At best, then, this proposal counts E^* as evidence relative to a total knowledge set that is not one of any possible ideally rational agent. In general, however, we would reject the suggestion that E is evidence for H, relative to K, just in case, given K' (K less some proposition P) and E' (E less some proposition Q), $P(H/K' \,\&\, E') > P(H/K')$.

Accepting that E simply cannot be cleaved from the accompanying temporally indexical self-knowledge required by ideal rationality, one might instead try to amend the relevance account so as to allow the comparison of epistemic situations that are not such that one is a proper part of the other but that are otherwise relevantly similar. So, where 'E^*' and 'K^*' are as before and 'I-' is the introspective knowledge when E^* is uncertain and 'I' the introspective knowledge when E^* is certain, it might be claimed that E^* is evidence for H, given K^*, when $P(H/K^* \,\&\, E^* \,\&\, I) > P(H/K^* \,\&\, I\text{-})$.

Unfortunately, such a maneuver would, again, lose touch with the intuitive justificatory rationale behind the positive relevance conception. That rationale, to repeat, is that E is evidence for H, relative to K, when knowledge of E makes it rational to be more confident in H or, for certain subjectivists, explains one's being more confident in H. This clearly requires that we compare a situation in which E is known with an epistemic situation that is, but for knowledge of E, identical. The envisaged proposal flouts this requirement. Indeed, this proposal would seem, instead, to be a way of comparing the *total confirmation* that two distinct total bodies of evidence afford a particular hypothesis. Such comparisons are, in my view, sensible and often known to be correct, but they cannot be said to be justified by the fundamental intuition behind the relevance conception of confirmation.

The proposal just outlined may, however, be adapted to give us a plausible error theory of the cases in which it intuitively seems that the positive relevance account is correct. According to this error theory, cases that have seemed aptly captured by the positive relevance account are really cases in which true comparative claims of total confirmation hold. In cases that seem to support the positive relevance account, one's epistemic situation has changed so as to make it reasonable to be more confident in some proposition, but such a change is

the nontriviality of certain counterpossibles (Wilson 2016), and so those of us who adhere to orthodoxy on counterpossibles will find an appeal to grounding to be no escape from the dilemma here presented.

https://doi.org/10.1086/713907 Published online by Cambridge University Press

not, as required by the positive relevance account, a simple monotonic increase in one's total evidence. This error theory explains why we might mistakenly think the positive relevance account plausible by locating a true claim (concerning rationally greater confidence after a change in one's epistemic situation) alongside a mistaken one (concerning a merely monotonic increase in certainties in the shift between those epistemic situations), where the mistake is an easy one to make because we often neglect to consider essentially indexical knowledge.

5. Conclusion. I conclude that the positive relevance account is at odds with what is quite plausibly the extent of an ideally rational agent's self-knowledge. Given Indexical Credal Transparency, cogent sense cannot be made of the notion that, with respect to a given body of total evidence, some part of that evidence justifies or explains an ideally rational agent's having a higher confidence in some hypothesis than does the remainder of that evidence. We must, therefore, confront the unpleasant dilemma of choosing whether to reject the positive relevance account or the requirement of Indexical Credal Transparency. In light of the reasons outlined above for endorsing Indexical Credal Transparency and the explanation of the mistaken appeal of the positive relevance account just outlined, I am inclined to reject the positive relevance account. However, those who would hold on to it must rebut the case for Indexical Credal Transparency while, ideally, preserving a suitable rationale for the synchronic and diachronic probabilistic constraints required for the positive relevance account.

REFERENCES

Caie, Michael. 2013. "Rational Probabilistic Incoherence." *Philosophical Review* 122:527–75.
Carnap, Rudolph. 1962. *The Logical Foundations of Probability*. 2nd ed. Chicago: University of Chicago Press.
Carr, Jennifer. 2017. "Epistemic Utility Theory and the Aim of Belief." *Philosophy and Phenomenological Research* 95:511–34.
Christensen, David. 2007. "Epistemic Self-Respect." *Proceedings of the Aristotelian Society* 107: 319–37.
Edwards, Anthony W. F. 1972. *Likelihood*. Cambridge: Cambridge University Press.
Glymour, Clark. 1980. *Theory and Evidence*. Princeton, NJ: Princeton University Press.
Hawthorne, James. 2005. "*Degree-of-Belief* and *Degree-of-Support*: Why Bayesians Need Both Notions." *Mind* 114:277–320.
Horwich, Paul. 1982. *Probability and Evidence*. New York: Cambridge University Press.
Howson, Colin, and Peter Urbach. 2006. *Scientific Reasoning: The Bayesian Approach*. New York: Open Court.
Huber, Franz. 2008. "Assessing Theories, Bayes Style." *Synthese* 161:89–118.
Joyce, James. 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65:575–603.
Kaplan, Mark. 1996. *Decision Theory as Philosophy*. New York: Cambridge University Press.
Konek, Jason, and Ben Levinstein. 2019. "Foundations of Epistemic Decision Theory." *Mind* 509:69–107.
Maher, Patrick. 1996. "Subjective and Objective Confirmation." *Philosophy of Science* 63:149–74.

https://doi.org/10.1086/713907 Published online by Cambridge University Press

Mahtani, Anna. 2015. "Dutch Books, Coherence and Logical Consistency." *Noûs* 49:522–37.
Milne, Peter. 1991. "A Dilemma for Subjective Bayesians—and How to Resolve It." *Philosophical Studies* 62:307–14.
———. 1996. "log[P(h/eb/P(h/b)] is the One True Measure of Confirmation." *Philosophy of Science* 63:21–46.
Perry, John. 1979. "The Problem of the Essential Indexical." *Noûs* 13:3–21.
Pettigrew, Richard. 2016. *Accuracy and the Laws of Chance*. New York: Oxford University Press.
Pust, Joel. 2007. "Cartesian Knowledge and Confirmation." *Journal of Philosophy* 104:269–89.
———. 2012. "Conditionalization and Essentially Indexical Credence." *Journal of Philosophy* 109:295–315.
———. 2021. "Dutch Books and Logical Form." *Philosophy of Science*, forthcoming.
Skyrms, Brian. 1986. *Choice and Chance*. 3rd ed. Belmont, CA: Wadsworth.
Sobel, Howard. 1987. "Self-Doubts and Dutch Strategies." *Australasian Journal of Philosophy* 65:56–81.
Titelbaum, Michael. 2013. *Quitting Certainties*. New York: Oxford University Press.
———. 2016. "Self-Locating Credences." In *The Oxford Handbook of Probability and Philosophy*, ed. Alan Hajek and Christopher Hitchcock. New York: Oxford University Press.
Williamson, Timothy. 2008. "Why Epistemology Can't Be Operationalized." In *Epistemology: New Philosophical Essays*, ed. Q. Smith. New York: Oxford University Press.
Wilson, Alastair. 2016. "Grounding Entails Counterpossible Non-triviality." *Philosophy and Phenomenological Research* 96:716–28.
Zalabardo, José. 2009. "An Argument for the Likelihood Ratio Measure of Confirmation." *Analysis* 69:630–35.

https://doi.org/10.1086/713907 Published online by Cambridge University Press