Evidence and Justification in Groups with Conflicting Background Beliefs

Kent W. Staley

Abstract: Some prominent accounts of scientific evidence treat evidence as an unrelativized concept. But whether belief in a hypothesis is justified seems relative to the epistemic situation of the believer. The issue becomes yet more complicated in the context of group epistemic agents, for then one confronts the problem of relativizing to an epistemic situation that may include conflicting beliefs. As a step toward resolution of these difficulties, an ideal of justification is here proposed that incorporates both an unrelativized evidence requirement and the requirement of the security of the evidence on which a conclusion from data is based. The latter requirement incorporates the consideration of epistemic modal statements.

1. Introduction

Among the various accounts of scientific evidence or confirmation, an important conceptual distinction can be drawn between those that make the application of such concepts *relative* (to a set of background beliefs, to a body of information, or to an epistemic situation) and those that attempt to articulate *unrelativized* conceptions of evidence or confirmation.

Unrelativized concepts of evidence seem to be open to the following objection: If *e* being evidence for *h* is not relative to the beliefs or epistemic situation K_S of a person *S*, then, unless either *e* or *h* is somehow about K_S, *e* being evidence for *h* cannot in any way depend on K_S.[1] Assuming, however, that *e* being evidence for *h* suffices to provide *S* a good (or at least some) reason to believe *h* (or assert *h* as true), it would seem that this determination *should* depend on K_S. For whether *S* is epistemically justified in believing or asserting[2] *h* on the basis of *e* should depend not only on whether *e* is true (and other facts about *e*'s relationship to *h* that are independent of K_S), but also on facts about *S*, particularly those facts about *S* implicated in K_S.

On the other hand, relativization brings its own difficulties. Here I focus on just one of these that arises in the consideration of what I have elsewhere called *evidential collaborations*. These are collaborations that exist in order to contribute to an understanding of some aspect of the world through their coordinated research activities, and that make collective assertions of evidence claims in partial fulfillment of this constitutive aim.[3]

The members of such groups sometimes hold incompatible individual beliefs. Moreover, the beliefs implicated in such internal disagreements may concern matters bearing directly on the evidential status of particular experimental findings with regard to substantive scientific claims under investigation by the group. This poses a significant challenge to any attempt to understand evidence claims made by collaborations as relativized to some background beliefs or epistemic situation, as the latter will in such cases contain a contradiction. Call this the *problem of incompatible backgrounds.*

The extent of the problem of incompatible backgrounds is difficult to assess. It is in the nature of evidential collaborations to craft their public assertions so as to present an image of consensus. Internal disagreements that lie beneath the surface of such consensus documents as conference papers and published research reports can usually only be unearthed by asking members of the group pointed and not always welcome questions. Setting aside the extent of the problem as an open empirical question, it is clear that incompatible backgrounds do *sometimes* exist, and it is at least plausible that this is especially likely at the leading edges of research in any given area.[4]

The difficulties I have described facing unrelativized and relativized evidence concepts are far from conclusive. Consider Peter Achinstein's response to the criticism that such a concept does not make the status of an individual's beliefs regarding a hypothesis appropriately sensitive to other beliefs she holds. Achinstein distinguishes two evidence concepts. First, there is an unrelativized evidence concept such that if *e* is evidence for *h*, then it constitutes a reason to believe *h* in an absolute and abstract sense, independently of any epistemic situation. A second evidence concept, according to Achinstein, captures the idea of epistemically justified belief in *h* on the basis of *e*, and this concept *is* relativized to epistemic situations (Achinstein 2001; Staley forthcoming).

Defenders of relativized notions of evidence might respond to the problem of incompatible backgrounds with various strategies. The simplest might be to insist that groups with such unresolved disagreements cannot have beliefs that are justified by evidence – at least not beliefs for which the matters under dispute are relevant to their evidential support. Thus consensus conclusions based on the "pragmatic" setting aside of conflicting background beliefs would be regarded as epistemically defective. (I have argued

against such a view in Staley 2007.) Another approach would be to restrict the relativization to the group's collective beliefs – i.e., those beliefs that can be attributed to the group as a whole. Although this move does not guarantee consistency (groups, no less than individuals, might hold inconsistent beliefs), it would at least solve the problem of relativizing evidence claims to incompatible beliefs held by disjoint subgroups or distinct individuals within the group. The beliefs in the relativization base would be purged somehow of such inconsistencies.

The present discussion pursues another approach that seeks to preserve elements of both the unrelativized and the relativized conceptions of evidence – not, as does Achinstein, by articulating two separate concepts, and not by excluding or purging incompatible beliefs, but by articulating a single ideal of justification that includes both relativized and unrelativized components. For an agent *S*'s assertion of a hypothesis *h* on the basis of evidence *e* to be justified, according to the present proposal, *e* must in an unrelativized sense constitute good evidence for *h*, but also *S* must undertake to justify the inference from *e* to *h*, so as to rule out possible scenarios under which *e* would fail to be good evidence for *h*. Whether the latter requirement is met *is* relative to background beliefs or epistemic situation, and I discuss some advantages that this approach enjoys when it comes to handling the problem of incompatible background beliefs within an evidential collaboration.

My discussion proceeds as follows: In section two I discuss objective, unrelativized theories of evidence. In section three I discuss the justification of inferences through the *securing of evidence* and propose an ideal of justified inference based on the framework

presented in the previous two sections. In section four I apply this framework to the problem of incompatible backgrounds.

2. Objective Theories of Evidence

Here I briefly discuss two prominent approaches to the evidential evaluation of scientific claims that purport to be objective in a robust sense. Both define evidence in a manner that is not relativized to background beliefs or epistemic situations. First I will discuss Achinstein's probabilistic-explanatory hybrid account of evidence. Then I turn to Deborah Mayo's error-statistical account of evidence.

2.1 Achinstein's Probabilistic-Explanatory Hybrid

Achinstein conceives of some kinds of statements that scientists make about evidential relationships between facts and hypotheses as both: (1) *objective* in the sense that statements of those kinds are true or false independently of what anyone believes about the hypotheses and facts in question; and (2) *empirical* in the sense that ascertaining the truth or falsehood of such statements is (at least sometimes) a matter for investigation by means of experiment and observation rather than a priori calculation.

These features are incorporated into the concepts of "potential evidence" and "veridical evidence." He proposes the following as necessary and sufficient conditions for the former:

PE (Potential Evidence): *e* is potential evidence that *h*, given *b*, if and only if:

1. p(there is an explanatory connection between *h* and *e*/*e* and *b*) > ½
2. *e* and *b* are true
3. *e* does not entail *h*. (2001, 170)

Here *e* refers to some fact, *h* is a hypothesis, and *b* is background information. The probability statement in the first condition should be understood in terms of objective epistemic probability, according to which a statement of the form "p(h/e) = *r*" should be interpreted as asserting that "the degree of reasonableness of believing *h*, on the assumption of *e*, is *r*." (106)

Veridical evidence is defined by adding to the conditions specified in *PE* the further requirement that *h* is true.

Both potential evidence and veridical evidence share a feature that is central to Achinstein's concerns: If *e* is evidence that *h* in either the potential or veridical sense, then *e* is a *good reason to believe* that *h* is true. Moreover, it is a good reason to believe *h* in a sense that is completely independent of any epistemic situation.

2.2 Mayo's Error-Statistical Account

According to Deborah Mayo's error-statistical account, good evidence results from the use of testing procedures that have certain good characteristics when applied to hypotheses of interest. Tests that possess such characteristics, which can be described in terms of error probabilities, enable investigators to learn from data because of their probative value with regard to the hypotheses being investigated. More specifically, "Data $\boldsymbol{x}_0$ in test *T* provide good evidence for inferring *H* (just) to the extent that *H passes severely* with $\boldsymbol{x}_0$" (Mayo and Spanos 2006, 328; emphasis added).

The notion of severely passing a test can be schematized as follows:

ES (Error Statistics): H passes a severe test *T* with $\boldsymbol{x}_0$ if

ES1 $\boldsymbol{x}_0$ fit *H*;

ES2 with very low probability, test *T* would have produced a result that fits *H* as well as (or better than) $\boldsymbol{x}_0$ do, if *H* were false and some alternative to *H* were true.

While various measures of fit might be employed, at a minimum, for $\boldsymbol{x}_0$ to fit *H*, $\boldsymbol{x}_0$ must not be improbable under *H* by comparison with competing hypotheses.

The probabilistic framework supporting these criteria articulates the relevant probabilities (particularly the error rates reflected in the severe test requirement) in frequentist terms. That is to say that the relevant probabilities are to be construed as objective facts about the relative frequency with which certain kinds of events occur in specified (actual or hypothetical) replications of the experimental procedures followed, under the assumption of the relevant hypotheses.

Note here that the error-statistical framework is not concerned with relations among beliefs. Any reference to a severe test relates a testing procedure, a body of data, and a hypothesis. The *ES* account explicitly distances itself from making the status of an inference depend on the beliefs – real or imagined – of investigators (Mayo 1997), thus marking a strong contrast with personalist Bayesian accounts (e.g., Howson and Urbach 2006).

Although Mayo's and Achinstein's accounts differ in important ways over how to understand evidence (see Achinstein 2009 and subsequent commentary for an interesting exchange), they agree insofar as they endorse concepts of evidence that are defined in terms of conditions that obtain or fail to obtain objectively, i.e., in a manner that is independent of actual or hypothetical beliefs or epistemic situations.[5] What is perhaps less obvious is that both accounts recognize, at least implicitly, that this very

feature entails that the notion of epistemic justification is distinct from the notion of evidence thus theorized.

Achinstein confronts this point directly. He distinguishes between *e* being a good reason to believe *h* and *e* justifying a person in a particular epistemic situation in believing *h*. He denies that the latter has any very important role to play in science. What scientists seek, according to him, are facts that constitute in an absolute sense good reasons to believe that their hypotheses are true. Nonetheless, he acknowledges that there is an evidence concept that is concerned with epistemic justification, which he calls ES-evidence, such that if *e* is ES-evidence that *h* relative to an epistemic situation K_S, then anyone in situation K_S would be justified in believing *h*. (However, the ES-evidence concept does not really provide an account of such epistemic justification, since Achinstein defines it simply in terms of being *justified* in believing that *e* is potential evidence that *h* given a certain epistemic situation.)

Mayo's approach is less explicit regarding this contrast. First, it is important to note that the error-statistical account is concerned with the justification of inferences from data. Mayo has written that the central focus of her account is to "provide a way to determine the evidence that a set of data x_0 supplies for making warranted inferences about the process giving rise to x_0." (Mayo and Spanos 2006, 327)[6] There is reason to believe that the evidence concept articulated in conditions *ES1* and *ES2* above does not tell the full story about justification in the error-statistical approach.

To see this point it helps to look more closely not simply at the schematized notion of evidence defined by *ES1* and *ES2*, but at the error-statistical methodology that Mayo has articulated in support of her account of evidence. That methodology, developed so

as to enable the investigator to pursue evidence that meets the requirements of *ES1* and *ES2*, emphasizes that scientific investigations must "severely probe" for error in the drawing of inferences from data. In the field or laboratory, this amounts to the need to engage in a wide variety of activities aimed at checking for errors in assumptions about instrumentation, the control of confounding variables, the nature of the data-generating process under investigation, auxiliary theoretical assumptions, ceteris paribus factors, etc.

Moreover, in cases of explicitly statistical inference, in which error-probabilities are evaluated based on an assumed statistical model, Mayo and Spanos have developed and insisted upon the deployment of tests directed at the assumed model itself, to determine whether it is adequate for the primary inference being drawn (Spanos 1999; Mayo and Spanos 2004; Staley 2010).

This supporting methodological framework suggests that Mayo also recognizes that, for purposes of drawing *justified* inferences, the bare satisfaction of the conditions set forth in *ES1* and *ES2* is insufficient. One must also be in a good position (i.e., an epistemic situation) to validate that one has taken the appropriate measures to secure the satisfaction of these conditions – and hence the question of justification does, in some sense, depend on one's epistemic situation.

3. Justifying Inferences by Securing Evidence

Here I want to take up this limitation of unrelativized notions of evidence and suggest a framework by means of which we might supplement these accounts so as to arrive at a richer ideal of justification that builds upon unrelativized evidence but extends to

acknowledge the relevance of an investigator's epistemic situation, whether conceived individually or collectively.

It will quickly become apparent that the framework I propose for thinking about justification in the context of science bears little resemblance to existing accounts of justification within the literature from general epistemology in the analytic tradition. This stems from a difference over philosophical methodology. Rather than specifying necessary and sufficient conditions for the truth of "S is justified in believing *p*," I aim to specify a conceptual framework that serves to explicate *justificatory practices* in the sciences. Such a framework centers on an *ideal* of justification, at which justificatory practices aim (Staley and Cobb forthcoming), and thus approaches justification by asking what it is that such practices accomplish, and how.

The decision by an evidential collaboration to present a conclusion from the data they have gathered rests upon the group's collective conviction that they are prepared to justify their inference in response to whatever challenges they expect to encounter. Their confidence will result from their having already posed many such challenges to themselves.

New challenges will emerge from the community of researchers with which they communicate. Such challenges take many forms, depending on the nature of the experiment and conclusions: Are there biases in the sampling procedure? Have confounding variables been taken into account? To what extent have alternative explanations been considered? Are estimates of background reliable? Have instruments been adequately shielded, calibrated, and maintained?

To a large extent, such challenges present possible scenarios in which the

experimenters have gone wrong in drawing the conclusions that they draw. Such challenges are not posed arbitrarily. For example, being logically possible does not constitute a challenge that the experimenter is responsible for addressing. Rather, such scenarios are judged significant by those in a certain kind of epistemic situation, incorporating relevant disciplinary knowledge; an appropriate response needs to provide a basis for concluding that the scenario in question is not actually the case.

The steps that investigators take, then, to justify an inference from data succeed just insofar as they provide grounds for setting aside worries about possible states of affairs that would, if actual, undermine that inference by falsifying a premise or background assumption on which it rests. In what follows, I will refer to such possible states of affairs as "error scenarios."

I propose thinking of practices of justifying an inference as the *securing* of that inference against such error scenarios. To explain what I mean, I will first introduce the notion of security as it applies to propositions generally, then extend the notion to inferences and evidence claims specifically.

The first concept we need is that of a *scenario*. David Chalmers has developed this notion as the basis of a modal semantics; his rough characterization will here suffice: a scenario is a "maximally specific way things might be" (Chalmers forthcoming). I will also speak of scenarios in terms of *epistemic possibility*, which can be thought of as the modality employed in such expressions as "For all I know, magnetic monopoles might be created at the Large Hadron Collider" and "For all I know, that sandwich might still be in Baltimore."[7] If there is, relative to one's epistemic situation, an epistemically possible scenario in which a proposition *P* is true, that means that, for all one knows, a

complete and maximally specific description of the world entails P.[8] (In what follows the word "possible" and its cognates should be understood to refer to the modality of epistemic possibility unless otherwise specified.)

Suppose, then, that relative to a certain epistemic situation K, there is a range of scenarios that are epistemically possible, and call that range Ω_0 (one might think of it as a set of scenarios). If proposition P is true in every scenario in the range Ω_0, then P is *fully secure* relative to K. If P is true across some more limited portion Ω_1 of Ω_0 (one might think of Ω_1 as a subset of Ω_0), then P is *secure throughout* Ω_1.

To put this notion more intuitively, then, a proposition is secure for an epistemic agent just insofar as, whatever might be the case for all that the agent knows, that proposition remains true. Although thus defined, security applies to any proposition, the application of interest here is evidence claims and inferences. Specifically, an *inference* from fact e to hypothesis h is secure relative to K insofar as the *proposition* "e is good evidence for h" is secure relative to K.

Note that the notion of a fully secure inference functions as an ideal for use in articulating an account of justification. Moreover, this account does not suppose that investigators can or should attempt to determine some *degree of security* of any of their inferences.

Rather, the value of the concept of security lies in its capacity to conceptualize methods of justification encountered in scientific practice in a systematic way. Indeed, the methodologically significant notion is not security *per se*, but the *securing* of inferences, i.e., the use of methods that serve to increase the *relative* security of an evidence claim, either by expanding the range, in a fixed space of possible scenarios,

across which that claim is true, or by decreasing the range of possible scenarios in which the evidence claim would be false.

With this terminology in hand, I propose the following as an ideal of justification:

JE (Justified Evidence): An assertion of *h* as a conclusion inferred from observed fact(s) *e* is fully justified relative to epistemic situation *K* if:

(1) *e* is good evidence for *h*; and

(2) the proposition "*e* is good evidence for *h*" is secure throughout all scenarios that are epistemically possible relative to *K*.[9]

This account articulates a notion of full justification as an ideal. The point is that methods of justification serve two distinct purposes. First, they aim (fallibly) to create conditions that will render (1) true for the inference at which the investigators arrive. Second, they aim to facilitate the pursuit of (2) by providing investigators with the resources to respond to the challenge of possible error scenarios and, thus, serve to secure the inference proposed. Though full security may remain an unachieved ideal, the increase in relative security puts investigators in a better epistemic situation, and it is in this sense that methods aimed at securing evidence claims provide justification.

To clarify the present proposal, some general comments are in order:

(I) Two general strategies for the securing of evidence pervade experimental science: In a *weakening* strategy one replaces a conclusion *h* with a weaker conclusion *h′* that is true across a broader range of possible scenarios. A *strengthening* strategy calls for changing one's epistemic situation *K* into a stronger situation *K′* such that error scenarios possible relative to *K* are not possible relative to *K′*.

(II) Strengthening strategies only work, of course, if it is possible to gain the kind of

knowledge that renders previously possible scenarios no longer possible. Thoroughgoing skepticism would therefore pose a difficulty for the present account. Beyond assuming that substantive knowledge about the world is possible, no particular account of knowledge is assumed or needed here, as many strategies for securing evidence can easily be shown to be effective at increasing the relative security of a claim for any plausible account of knowledge.

(III) Just how such strengthening knowledge is gained is also not specified in this model. In particular, it is not assumed that only apodictic reasoning yields knowledge enabling one to rule out scenarios as possible. On the contrary, one plausible understanding of error-statistical reasoning is that it secures substantive conclusions against alternative hypothetical possibilities by a form of probabilistic reasoning. Because the reasoning is probabilistic, those alternatives remain in some sense possible, even when the arguments against them on error-statistical grounds are very strong. *If* error-statistical reasoning can yield knowledge that errors against which the hypothesis has been severely tested are ruled out (as Mayo insists it does), then the sense in which those hypothetical errors remain possible must be regarded as distinct from epistemic possibility.

(IV) That said, note that condition (1) in *JE* is noncommittal regarding the evidence concept in question. Versions of *JE* adapted to Mayo's error statistics (Staley and Cobb 2010) and Achinstein's account of evidence (Staley forthcoming) have been articulated elsewhere.

4. Justified Evidence and Group Knowledge

The ideal of justification articulated in the previous section can, I propose, be fruitfully applied in contexts where evidence claims are advanced by evidential collaborations whose members hold conflicting beliefs regarding matters relevant to the evidence claims they assert.

The first thing to note is that the truth of epistemic modal statements depends in some way on a state of information or knowledge. Just *how* they so depend, and on *whose* states of information or knowledge is a matter of considerable dispute in the literature. Roughly put, solipsistic contextualists consider the truth of such statements to depend on what is known by the person making the assertion; non-solipsistic contextualists (DeRose 1991) make the truth conditions of such statements depend on features of the context of assertion that go beyond what the person asserting knows, such as information that is in some sense available (Hacking 1967) or on a group's distributed knowledge (Teller 1972); and solipsistic relativists consider the truth of such statements to vary according to the context of the person assessing their truth (MacFarlane forthcoming).

I propose to show the usefulness of the *JE* ideal in dealing with the problem of incompatible backgrounds in a way that bypasses such disputes over semantics. Although the details of applying the framework (consideration of what makes an epistemic modal statement true, for example) might depend on whether we attend to the context of assertion or of assessment, for example, the methodological payoff will require attending to both.

To keep things simple, we will suppose the following partially unrealistic situation. A group *G* of experimenters has been seeking evidence of a theoretically postulated

particle *X*, the decay of which has a known signature. They have attempted to produce *X*'s by producing them in collisions at a particle accelerator. The decay signature of *X* can be mimicked by other particles *A*, *B*, and *C*. Let us suppose that the expected background due to *A*s and *B*s has been calculated, and that the group *G* has found a statistically significant excess of decay signatures beyond the *A*- and *B*-backgrounds. *C*s cannot, let us suppose, be created in particle collisions of the type produced in *G*'s experiment, so they have not included *C*-decays in their background calculation.

The *G* group concludes that they have evidence for *X* particles being produced, and they send physicist *P*, who led the group that did the background calculations, to present these findings at a conference. A few other members of the *G* group are present in the audience as *P* gives his talk, including physicist *Q*, who oversaw detector construction. The talk goes well, but *P* becomes a bit flustered when an audience member with expertise in the physics of *C* particles raises the following objection: "You do not expect *C*s to be produced at your accelerator, but what about *C*s in cosmic rays? They are pretty common, and a back-of-the-envelope calculation suggests they might account for your excess." After a few long seconds of forehead-furrowing and paper-shuffling, *P* grumbles, "I hadn't ever really thought about *C*s in cosmic rays. I have to admit that the excess signature events might be decays of cosmic ray *C*s." At this point, *Q*, who has been squirming in her seat, speaks up: "No, they cannot be. We shielded the detector against *C*s from cosmic rays." Brightening, *P* turns to the questioner and says, "No, the excess cannot be due to cosmic ray *C*s; they must be from the decay of *X*s."

The questioner in the audience poses an error-scenario for the claim that *G*'s data constitutes evidence for the production of *X*s. Call it "the cosmic ray *C* scenario." For

purposes of defining the semantics of epistemic modals, one would want to attend to such issues as whether *P* spoke truly at first in asserting that the *C* scenario is possible, whether his subsequent denial of that possibility should be regarded as a retraction or a revision of his first statement, and so on. For our purposes, the significant point is that *Q*'s statement provided information explaining the steps that the group had taken to remove the cosmic ray *C* scenario from the realm of possibility, and thus the shielding of the detector, as described by *Q*, can be regarded as a successful justificatory step regarding the group's claim to have found evidence of *X*s.

Next I wish to argue that regardless of one's favored approach to the semantics of epistemic modals, thinking of such justifications in terms of the elimination of epistemic possibilities, backed by an objectively-grounded evidence relationship, captures the right sense in which such justifications are relative to background beliefs or epistemic situations and the sense in which their success or failure is an objective matter.

It will help to consider two possible variations on the story just told:

(1) Suppose that the *G* group in fact did *not* take any steps to shield their detector, and that, as the objector worried, *C* particles in cosmic rays really are common enough that they might account for the statistical excess of *X*-decay signature events. At *P*'s conference talk, *Q* says nothing but squirms uncomfortably in her seat during the discussion. However, unbeknownst to the *G* group, the subcontractor who put the roof on the structure housing the detector used a kind of material that just happens to absorb 98% of *C*s in cosmic rays (and a similar proportion of those coming from the other side of the earth are likewise absorbed by the earth itself).

(2) Suppose that, although the *G* group did install shielding around the detector to

block *C* particles from cosmic rays, their shielding was ineffective (perhaps they miscalculated the decay length of *C*s in the shielding material) and most *C*s due to cosmic rays passed right through it.

In variation (1), the cosmic ray *C* scenario remains uneliminated because, although there is a sense in which the cosmic ray *C* scenario is not possible, the information on the basis of which this claim would be correct is *ex hypothesi* not available to anyone in the group *G* or to anyone in the audience to which the evidence is presented. Recall that the point of the ideal *JE* is in part to capture the importance of an epistemic agent's being in a good position to defend her inferences.

The details of this story will vary according to one's views on the semantics. For a contextualist, the point is that in every relevant context of assertion in the story, the cosmic ray *C* scenario remains possible given the knowledge or information available to the person making the assertion, even if we extend the relevant knowledge base to that of the entire collaboration. For a relativist like MacFarlane (forthcoming), the point is that in every relevant context of assessment within the story, the cosmic ray *C* scenario remains possible. Of course, in the context of *our* assessment, we may consider the cosmic ray *C* scenario not to be possible. That is not a relevant scenario for the purposes of the story, however.

In variation (2), the effort by *Q* to justify the inference presented by *P* by describing the efforts to shield the detector fails. Although *Q believes* that the detector is adequately shielded against *C*s so that the *C* scenario is rendered impossible, she does not know that the detector is adequately shielded (for it is not). Hence the cosmic ray *C* scenario, given what she (or anyone else in the *G* group) does know, remains a possibility.

Finally, to see how the two components of the *JE* account (the evidence component and the security component) interact, consider a third variation that combines (1) and (2):

(3) As in (2), the *G* group attempted to shield the detector, but failed to do so adequately. But, as in (1), the choice of roofing material, unbeknownst to anyone, effectively shields the detector against *C* particles anyway.

Plausibly, the requirements for evidence, on either Achinstein's or Mayo's account, could be met under these conditions. *C* particles, after all, really are shielded in version (3), though not in the way that members of *G* believe. But the security requirement that I introduced is *not* met, because nothing that the members of *G* know entitles them to rule out the cosmic ray *C* scenario as a possibility.

Of course, in the original version of the case just described, the initial discrepancy between *P* and *Q* regarding the possibility of the cosmic ray *C* scenario was easily and quickly resolved simply by *Q*'s testimony regarding the shielding. Many cases of incompatible background beliefs will not yield so readily to quick resolution. In such cases, it may remain unclear to members of the collaboration whether relevant error scenarios have been eliminated. The point of the present framework, however, is not to remove the problem of intra-collaboration disagreement through the magic of conceptual analysis. Rather, I sought to show how a relativized notion of justification could be brought to bear in a way that could make sense of relativizing to the epistemic situations of group members holding incompatible beliefs.

On the present proposal, an error scenario *S* remains unsecured by a group if: (i) relative to the epistemic situations of some group members, *S* is possible and (ii) no

other member of the group has knowledge that suffices to show that *S* is not possible. Of course, in the methodological sense, to *secure* an inference against an error scenario *S* requires more than simple failure of condition (i) or (ii). In keeping with the emphasis on scientific practice mentioned above, a full account of the securing of inferences needs to go on to consider the kinds of steps that must be taken to *show how* and on *what basis* an error scenario can be ruled out. (It is for this reason that I have given only sufficient and not necessary conditions for *S* remaining unsecured.) But that would be the subject of another paper (see Staley 2010, for example).

Finally, let me note a major point of incompleteness in the present discussion. I have at times resorted to the locution "relevant error scenario," but have not addressed the question of relevance. To be sure, it is trivial to cook up error scenarios that are logically and perhaps even epistemically possible.[10] The question of how well justified any given inference is, based upon the *JE* ideal, will depend not only on the elimination of error scenarios, but on the relevance of those that have and have not been eliminated. The delineation of criteria for judging relevance remains an open question for this approach that goes beyond the scope of the present essay.

5. Conclusion

For too long (though it was not always so), it has seemed that epistemologists and philosophers of science have had little to say to each other (but see, e.g., Roush 2005). This has been a missed opportunity for both fields. Given the social nature of scientific inquiry, social epistemology seems like a natural place for the discourses of these two specializations to engage with one another. Here I have attempted to articulate a way in

which, by bringing together ideas from debates over the nature of scientific evidence and ideas about epistemic modality, one might begin to articulate a framework for understanding how evidential collaborations justify their inferences, even in the midst of substantive disagreement in their background beliefs.

Such justification could, on this account, be thought of as aimed at two desiderata: (1) satisfying the requirements for evidence in an unrelativized sense and (2) securing that evidence against possible error scenarios, where such possibility is relativized, but in a way for which the incompatibility of background beliefs among members of the collaboration pose no special difficulties.[11]

References

Achinstein, Peter. 2001. *The Book of Evidence*. New York: Oxford University Press.

Achinstein, Peter. 2009. "Mill's Sins or Mayo's Errors?" In D. Mayo and A. Spanos (eds.), *Error and Inference: Recent Exchanges on Experimental Reasoning, Reliability, and the Objectivity and Rationality of Science*, pp. 170–88. New York: Cambridge University Press.

Chalmers, David. Forthcoming. "The Nature of Epistemic Space." In A. Egan and B. Weatherson (eds.), *Epistemic Modality*. Oxford: Oxford University Press.

DeRose, Keith. 1991. "Epistemic Possibilities." *The Philosophical Review* 100: 581–605.

Hacking, Ian. 1967. "Possibility." *The Philosophical Review* 76: 143–68.

Hintikka, Jaakko. 1962. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Ithaca: Cornell University Press.

Howson, Colin and Peter Urbach. 2006. *Scientific Reasoning: The Bayesian Approach.* 3rd ed. Chicago: Open Court.

Kratzer, Angelika. 1977. "What 'Must' and 'Can' Must and Can Mean." *Linguistics and Philosophy* 1: 337–55.

MacFarlane, John. Forthcoming. "Epistemic Modals Are Assessment-Sensitive." In A. Egan and B. Weatherson (eds.), *Epistemic Modality*. Oxford: Oxford University Press.

Mayo, Deborah. 1996. *Error and the Growth of Experimental Knowledge*. Chicago: University of Chicago Press.

Mayo, Deborah. 1997. "Duhem's Problem, the Bayesian Way, and Error Statistics, or, 'What's Belief Got to Do with It?' " *Philosophy of Science* 64: 222–44.

Mayo, Deborah and Aris Spanos. 2004. "Methodology in Practice: Statistical Misspecification Testing." *Philosophy of Science* 71: 1007–25.

Mayo, Deborah and Aris Spanos. 2006. "Severe Testing as a Basic Concept in a Neyman-Pearson Philosophy of Induction." *The British Journal for the Philosophy of Science* 57: 323–57.

Rehg, William and Kent Staley. 2008. "The CDF Collaboration and Argumentation Theory: The Role of Process in Objective Knowledge." *Perspectives on Science* 16: 1–25.

Roush, Sherrilyn. 2005. *Tracking Truth: Knowledge, Evidence, and Science*. New York: Oxford University Press.

Salerno, Joe. 2009. "Epistemic Modals: Relativism and Ignorance." Talk given at Midwest Epistemology Workshop, Saint Louis University, September 25, 2009.

Spanos, Aris. 1999. *Probability Theory and Statistical Inference*. Cambridge: Cambridge University Press.

Staley, Kent. 2004a. *The Evidence for the Top Quark: Objectivity and Bias in Collaborative Experimentation.* New York: Cambridge University Press.

Staley, Kent. 2004b. "Robust Evidence and Secure Evidence Claims." *Philosophy of Science* 71: 467–88.

Staley, Kent. 2007. "Evidential Collaborations: Epistemic and Pragmatic Considerations in 'Group Belief.' " *Social Epistemology* 21: 321–36.

Staley, Kent. 2010. "Strategies for Securing Evidence through Model Criticism." Unpublished ms. Available online. http://sites.google.com/site/staleykw/home/papers/StratSecEv.pdf

Staley, Kent. Forthcoming. "Evidence and Justification." In G. Morgan (ed.), *Philosophy of Science Matters: The Philosophy of Peter Achinstein*. New York: Oxford University Press.

Staley, Kent and Aaron Cobb. Forthcoming. "Internalist and Externalist Aspects of Justification in Scientific Inquiry." *Synthese*. Available online. DOI: 10.1007/s11229-010-9754-y.

Teller, Paul. 1972. "Epistemic Possibility." *Philosophia* 2: 303–20.

Wray, K. Brad. 2007. "Who Has Scientific Knowledge?" *Social Epistemology* 21: 337–47.

---

[1] The notion of an epistemic situation as formulated by Achinstein is "an abstract type of situation" describing what "one knows or believes," what "one is not in a position to know or believe" and whether "one knows (or does not know) how to reason from the former to the hypothesis of interest" (2001, 20). Future occurrences of "K_S" should be understood to refer generically and indifferently to sets of background beliefs, or states of information, or epistemic situations, unless otherwise specified.

[2] Whether the focus of discussions of justification in science should be beliefs or assertions is a question not undertaken in the present essay. For a defense of focusing on the latter, see Staley and Cobb (forthcoming). Henceforth, I will drop the disjunctive "asserting or believing," when referring to the target of justification, but not because I am here deciding in favor of one over the other.

[3] In the present essay I will discuss such groups as constituting epistemic agents in their own right. Just what kinds of groups constitute collective epistemic agents is a matter of dispute (see, e.g., Wray 2007), but apart from outright epistemic individualists, the groups that I describe as evidential collaborations would seem to count as epistemic agents on anyone's account.

[4] I have written previously about a clear example of such a case from the history of experiments leading to the first claim of evidence for the existence of the top quark. Details aside, suffice it here to say that members of the Collider Detector at Fermilab (CDF) collaboration, which made the claim, disagreed about the extent to which possible sources of bias in the analysis of their data had been ruled out by appropriate methodological precautions. In spite of such disagreements, they were able to present cogent arguments in support of the consensus conclusion that their data constituted evidence for the production of top–anti top quark pairs (Staley 2004a, 2004b; Rehg and Staley 2008).

[5] Mayo's account has been accused of not living up to its claims of objectivity, on the grounds that error probabilities are derived from a statistical model, which in turn is chosen by the investigator, and on the basis of the relevance of "stopping rules" in the evaluation of evidence (Howson and Urbach 2006). A full discussion of these issues would take us far afield from present concerns. Suffice it here to note that: (1) statistical models

themselves are used to *evaluate* error probabilities, but that actual error-probabilities are determined by features of the test actually employed; (2) statistical models themselves are subject to testing; and (3) stopping rules are relevant not as subjective mental states, but as procedures materially implemented in the execution of test procedures (Mayo 1996; Staley 2004a, ch. 7; Mayo and Spanos 2006; Staley 2010).

[6] Here Mayo and Spanos use the term "warranted" as a synonym for justified, and not in the sense that the term is used by epistemologists, as denoting the property that, in addition to truth, qualifies a belief as knowledge (Mayo, personal communication).

[7] Hintikka's seminal (1962) takes expressions of the form "It is possible, for all that *S* knows, that *P*" to have the same meaning as "It does not follow from what *S* knows that not-*P*." Just how to formulate the semantics of such statements is, however, contested (see, e.g., DeRose 1991; Salerno 2009; Chalmers forthcoming; MacFarlane forthcoming). The central claims of the present proposal are independent of disputed issues regarding the semantics of epistemic possibility.

[8] Readers who are more comfortable with the locution "possible worlds" may substitute this phrase for "possible scenarios" (see, e.g., Kratzer 1977), keeping in mind that the modality at issue here is not the subjunctive or counterfactual one with which that wording is often associated.

[9] Because I do not wish to claim that *JE* is the only justificatory ideal that might be relevant to science, I offer (1) and (2) as sufficient but not necessary conditions.

[10] Just how trivially easy this is depends on how much substantive knowledge we have. Is it true that for all I know God created the universe about 6,000 years ago, but went to the trouble of creating such things as the cosmic microwave background and fossils of ten

different orders of trilobites? Not if I know that the universe is approximately 14 billions of years old and that the last trilobites succumbed to extinction about 250 million years ago.

[11] I would like to thank Aaron Cobb for many helpful conversations on these issues, and Joe Salerno for guidance regarding the literature on epistemic modals. Work on this project was supported by National Science Foundation Award SES-0750691.

Brief bio:

Kent Staley is associate professor of philosophy at Saint Louis University. His work is primarily concerned with issues pertaining to scientific evidence and experiment, as well as the history and philosophy of physics and philosophy of statistics. His articles have appeared in *Philosophy of Science*, *Synthese, Perspectives on Science, Social Epistemology,* and *Physics in Perspective*. His book *The Evidence for the Top Quark* (Cambridge University Press) was published in 2004.