Internalist and Externalist Aspects of Justification in Scientific Inquiry

Kent Staley, Saint Louis University

Aaron Cobb, Auburn Montgomery University

1. Introduction

Contemporary epistemologists have devoted considerable attention to conceptual analyses of the nature of epistemic justification but there is great disagreement about whether the factors relevant to the justification of a person's belief must be internally accessible to that person (Alston 1989; Fumerton 1996; Kornblith 2001; Pryor 2001; BonJour and Sosa 2003; McGrew and McGrew 2006; Goldberg 2007; and Poston 2008). This debate between internalists, who endorse the access requirement, and externalists, who reject it, has been little discussed by philosophers of science.[1] Yet epistemic justification is a central concern in philosophy of science. In particular, the wide-ranging debates over evidence and confirmation seem to be concerned to a significant degree with the question of justifying conclusions from data. Theories of evidence can indeed be understood in part as attempts to explicate a concept of scientific justification. But how do such theories depict scientific justification? Do they employ an internalist or externalist notion of justification?

To facilitate an inquiry into these questions we reconsider the dichotomy between internalism and externalism from the perspective of justificatory practices in the sciences. In doing so, we find that the dichotomy as traditionally formulated does not adequately capture the nature of justification in scientific inquiry. We motivate our reformulation by attending to the socially-situated nature of practices of scientific justification. As part of this reformulation, we redirect the debate away from a concern with the justification of *beliefs* and toward the justification of assertions of experimental inferences. Such a redirection has a further basis in the nature of our

inquiry, which considers the question of justification from the perspective of objective accounts of scientific evidence.

More precisely, we are concerned with theories that treat evidential relationships as obtaining in a manner that is epistemically independent of the beliefs of particular individuals or groups. To that end, we examine this issue in the context of Deborah Mayo's error-statistical theory of evidence (Mayo 1996). We choose Mayo's account because it is a paradigmatically objective theory of evidence that is strongly informed by methodological practice. Our main argument, however, insofar as it does not depend strongly on the details of Mayo's account, would plausibly apply equally well to other objective theories such as likelihood accounts (Royall 1997; Lele 2004; Sober 2008), objective Bayesian theories (Jaynes 2003, Williamson 2008), or Peter Achinstein's explanatory-probabilistic hybrid (Achinstein 2001).

Our thesis is that, as understood from the standpoint of such an objective theory of evidence, justification in science has both externalist and internalist characteristics.

Our discussion proceeds as follows. In section two, we briefly review the terms of the contemporary debate between internalists and externalists in epistemology, and reframe that debate so as to make its terms applicable to scientific inquiry. In section three, we discuss Mayo's error-statistical theory of evidence and identify both externalist and internalist aspects of the concept of justification implicit in that account. Although error-statistical evidence is unrelativized, the justification of experimental conclusions does, we argue, depend on an epistemic situation. Section four introduces an epistemic notion – the security of an inference – that shares this

dependence on epistemic situation and illuminates the nature of the additional epistemic work that takes us beyond *de facto* evidence for a hypothesis to the justification of an inference to it. We articulate an *ideal* of justification that incorporates both an objective evidence component and a relativized security component. In section five, we return to the accessibility requirement that is at the heart of the internalist-externalist debate and present our conclusions. We find that security itself is not a strictly internalist notion, insofar as access, reformulated in terms of a community of inquirers, emerges as a necessary but not sufficient condition for securing experimental conclusions.

2. Reframing the Internalist-Externalist Debate

A full discussion of the general debate in contemporary epistemology between the proponents of internalism and externalism is beyond the scope of this paper, but a brief excursus into this territory is important for understanding the internalist and externalist aspects of justification implicit in an objective theory of evidence like Mayo's. Generally speaking, the debate between internalists and externalists concerns whether those factors justifying a person's belief must be cognitively accessible to the person.[2] Internalists argue that accessibility is both a necessary and sufficient condition for epistemic justification.[3]

The intuitive ground supporting internalism is the idea that providing a compelling answer to questions or worries about the epistemic status of a particular belief requires appealing to evidence or reasons that support the belief in question. Put more explicitly, internalism is the view that

> *Internalism*: a belief *b* is justified for a subject *S* at time *t* if and only if that which justifies *b* is cognitively accessible to *S*.[4]

Internalists generally understand cognitive accessibility as a relation holding between a subject *S* and what *S* can discover on reflection alone. James Pryor notes that internalists often understand the notion of accessibility

> in terms of the *route* by which one has access: one could understand it as meaning that one can *know by reflection alone* whether one is in one of the relevant states. (By 'reflection' I mean *a priori* reasoning, introspective awareness of one's own mental states, and one's memory of knowledge acquired in those ways…Most epistemologists understand the notion [of access] in [this way]. (James Pryor 2001, 103-104)

The basic idea is that the justifiers *j* for a belief *b* must be the sort of thing that could be an object of conscious awareness; if *j* cannot be an object of conscious awareness, *j* cannot serve as a justifier for *b*.

Externalists argue that cognitive accessibility is neither necessary nor sufficient for justification. It is not necessary because there are subjects (i.e., children or adults with relatively little cognitive sophistication) whose beliefs are justified even though access to the relevant justifiers for their beliefs may be impossible. Accessibility is not sufficient for epistemic justification because there is no guarantee that the information and evidence available to a subject is properly connected with the truth of the belief in question. Given the epistemic limitations of any subject, one cannot take that which is within one's cognitive grasp to exhaust the full range of what is relevant to the justification of a belief.

So, externalists want to develop an account of justification that accords with two basic intuitions. First, since there is good reason to believe that some epistemic subjects possess justified beliefs even though there is no reason to think they could have access to the relevant justifiers, an adequate account of epistemic justification must, at a minimum, show that it is possible for these subjects to possess justified beliefs. Second, there must

be a strong connection between justification and truth. The primary ground of this intuition is that the epistemic significance of whatever justifies a belief lies in its truth-conduciveness. (We use "truth-conduciveness" here as an umbrella term that applies to processes that have a tendency to produce true beliefs. Different accounts of justification will specify such a tendency in different terms.) So, in order to facilitate an explicit contrast with the internalist thesis articulated above, let externalism be the thesis that

> *Externalism*: a belief *b* is justified for a person *S* if and only if that which produces *b* is truth-conducive.

The internalist/externalist debate in contemporary epistemology is not concerned primarily with analyzing the nature of epistemic justification in the sciences. With some amendments, however, one can employ this framework schematically to clarify the nature and significance of particular justificatory principles and practices in the sciences.

Our first proposed modification requires a shift from the appraisal of beliefs to the appraisal of assertions as the proper object of epistemic evaluation. Whereas beliefs are private and individually held, at least in the paradigmatic cases, scientific knowledge is best regarded as a public and collective achievement. The activity of knowledge-production in the sciences generally occurs within a social structure and this involves acts of assertion by scientists in various forums (i.e., preprints, publications, presentations, decisions taken in collaboration meetings, etc.). In fact, one could argue that it is intrinsic to scientific knowledge not merely that the acquisition of it often requires groups of people but that one aim of the scientific enterprise is a particular kind of rationally persuasive communication in which reasons are presented to other members of the community that will serve to underwrite, within that community, the status of particular claims as knowledge. We do not deny that knowledge of scientific matters can be

ascribed to individual scientists. Rather, we are directing our attention to a distinct sense of scientific knowledge as publicly accessible content that arises from the socially organized efforts of individuals working in collaboration (cf. Kitcher 1993; Suppe 1993; Thagard 1997; Longino 2002; Wray 2002).

While beliefs are certainly relevant to actions particular scientists perform, including the activity of endorsing particular experimental conclusions, the assertion and endorsement of the relevant assertions in these forums can be distinguished from an individual scientist's beliefs about these assertions. Although not conclusive, these considerations suggest pursuing the idea that scientific knowledge should not be understood as essentially a species of belief.[5] At any rate, whether or not this is the case, one seeking to understand epistemic justification in the sciences and the practices that produce it would be better off looking to what is asserted in the appropriate social contexts than worrying about underlying beliefs, as it is through the interaction of communicative acts that the corpus of scientific knowledge is formed.[6]

The recognition of this intrinsic social structure of the sciences also suggests a further modification to the internalist/externalist dichotomy in contemporary epistemology. Since the objects of epistemic evaluation are the assertions made by scientists in particular professional forums and, often, these assertions are the product of a wide-ranging set of experiments conducted in collaboration with many other scientists, the notion of accessibility that divides internalists from externalists must be understood more broadly. Instead of thinking of accessibility as a relation holding between an individual subject *S* and what *S* can know through conscious reflection alone, accessibility must be relativized to the particular scientific community rather than individuals within the

community. We employ the term 'community' recognizing that there are a variety of communities in scientific inquiry and these communities can have distinct characteristics (i.e., they can be broad or small, loosely organized or tightly structured, etc.) and relations to other communities working on related questions.

To make this more precise, we can borrow the notion of an epistemic situation from Peter Achinstein's (2001) work on evidence. On Achinstein's analysis, an epistemic situation is

> an abstract type of situation in which, among other things, one knows or believes that certain propositions are true, one is not in a position to know or believe that others are, and one knows (or does not know) how to reason from the former to the hypothesis of interest, even if such a situation does not in fact obtain for any person. (Achinstein 2001, 20)

We propose that an epistemic situation describes the basic epistemological framework of the relevant scientific community working on a shared problem or question and advancing an experimental conclusion grounded on the basis of evidence produced in their research. Accessibility is a relation holding between a community as a whole and the data or evidence available within its epistemic situation.

We are now in the position to reframe the internalist/externalist dichotomy in terms relevant to epistemic justification in the sciences. In the remainder of this paper, we will employ the following definitions:

> *Internalism**: the assertion of an experimental conclusion (*h*) is justified relative to a particular epistemic situation (*K*) if and only if that which justifies *h* is accessible to those within *K*.
>
> *Externalism**: the assertion of an experimental conclusion (*h*) is justified if and only if that which justifies *h* is truth-conducive.

Our goal in the remaining sections is to show that understanding justification in scientific inquiry requires both internalist* and externalist* aspects—that is, an assertion *h* is

justified relative to *K* if and only if that which justifies *h* is both (i) accessible to those within *K* and (ii) truth-conducive.

3. The Error-Statistical Theory of Evidence

The internalist-externalist debate is concerned with the question of accessibility and not primarily with the manner in which scientists seek to justify conclusions from experimental data. We regard debates over evidence and confirmation in the philosophy of science as concerned with the latter question. But philosophical theories of evidence may be implicitly committed to either internalist* or externalist* views. This can be seen clearly in Deborah Mayo's error-statistical account. Mayo has herself 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).[7] An error-statistical theory of evidence is, then, concerned with justification. To see just how error-statistics characterizes the justification of inferences from data, we next review Mayo's account.

According to the error-statistical (*ES*) 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).

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

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 (Lele 2004), 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 (both the likelihoods implicated in fit assessment and 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, in keeping with our reformulation of the internalism-externalism debate in terms that eschew reference to beliefs, 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).

The *ES* account is meant to apply not only in cases where one quantitatively evaluates these probabilities using a statistical model of the data-generating process, but also in experimental settings that take a more casual or intuitive approach to

statistical analysis. As the former setting highlights important conceptual features of the error-statistical account, it is worth briefly noting the requirements of a more formal statistical assessment.

Any statistical inference in the *ES* approach will make use of a statistical model. As Cox notes, formal statistical inferences regard "the family of models as given and the objective as being to answer questions about the model in light of the data" (Cox 2006, 3).[8] Such a model will represent the data as the outcome of repeated trials resulting in values assigned to a random variable. It will specify a test statistic, defined in terms of the data, to be used in determining whether or not a hypothesis of interest passes or fails the test by reference to critical values of that test statistic, also specified by the model. The model will also specify the space of hypotheses (parameter values) between which the data will be used to discriminate. The model will involve assumptions about the *distribution* of values for the random variable, about the *dependence* of the outcome of a trial on the outcome of previous trials, and about *heterogeneity*, i.e., whether or not the distribution of outcomes changes from one trial to the next.

It is this model from which the probabilities used to characterize the testing procedure are derived. The model should encompass all of the hypotheses among which one is attempting to discriminate. It will be these specific alternatives to *H* that must be considered when evaluating probabilities in the context of requirement *ES2* above, and an examination that distinguishes between alternatives against which *H* has and has not been severely tested is crucial to an error-statistical analysis of the exact nature of the inference to be drawn with regard to *H* (cf. Mayo and Spanos

2006).

For our purposes, it is important to emphasize that for the investigator to *correctly* judge which hypotheses have and have not severely passed test *T* with data $\boldsymbol{x}_0$, the statistical model used in such an error-statistical analysis must be *statistically adequate*, in the sense that it "captures the statistical systematic information contained in the data" (Spanos 1999, 16). By ensuring that the actual error probabilities of the test are at least approximately equal to those assumed, statistical adequacy amounts to adequacy for the purpose of reliably drawing primary inferences (Spanos 1999; Mayo and Spanos 2004). How statistical adequacy is evaluated is a point to which we shall return, but it is important to note that statistical adequacy is an objective characteristic of statistical models. Even if one is unaware of any reason to doubt the model one employs in drawing a statistical inference—and even if any such reasons were inaccessible to the investigator—it may fail to be statistically adequate, in which case one's judgments about which hypotheses are evidentially supported by the data will be mistaken.

Thus we note an apparent anti-internalist* methodological consequence of the error-statistical view of evidence:

> *AI:* Error-statistical evidence claims can be rendered false by facts[9] to which the investigator has no access.

We will consider later in this section a challenge to this claim that purports to show that such defeaters are, after all, accessible. We will show that accessibility can be vouchsafed only for some error-statistical evidential inferences, but not in general. For now, we wish to pursue the consequences of accepting *AI*. It is anti-internalist

insofar as the error-statistical account depicts scientific justification as drawing both upon those reasons that are accessible to the investigator *and* the reasons implicated in the objectively obtaining evidential relations, even if those are not accessible to the investigator. Thus one might, in the internalist* sense, appear to be justified in asserting a hypothesis, while in the sense of justification that requires an objective evidential relationship between $\boldsymbol{x}_0$ and H, one is not justified in asserting H.

Hence, just as externalist views of justification hold that an individual's belief can be justified by reasons to which the believer has no access, error-statistical evidence relations can be satisfied or fail to be satisfied in virtue of facts to which the investigator has no access. Moreover, there seems to be a resemblance between *ES* and a paradigmatically externalist account of justification in epistemology. Just as Alvin Goldman's reliabilist theory makes justification rest on the tendency of a belief-forming process to produce true rather than false beliefs (Goldman 1986; 1999), *ES* links the justification of an inference to its having resulted from a testing procedure with low error probabilities (Woodward 2000). Contrary to what might be suggested by this similarity, however, there is good reason to think that the error-statistician will not hold a strictly externalist* view of justification. Seeing why requires, however, looking beyond the schematization of the error-statistical view of evidence discussed above, to the methodological framework that error-statistics draws upon for making what Mayo calls "arguments from error."

That methodology, developed so as to enable the investigator to pursue evidence that meets the requirements schematized above as *ES1* and *ES2*, emphasizes that scientific investigations must "severely probe" for error in the drawing of

inferences from data. In the laboratory, this amounts to the need to engage in a wide variety of activities aimed at checking for errors in assumptions about instrumentation, about the control of confounding variables, about the nature of the data-generating process under investigation, about auxiliary theoretical assumptions, ceteris paribus factors, etc.

This line of thought in Mayo's work has intersected with the work of the econometrician Aris Spanos in the development of a methodology aimed at testing the assumptions employed in statistical inference and modeling. Because any statistical inference will rely on an assumed statistical model, such inferences must always answer to the worry that flaws in the model assumptions defeat the inference that has been drawn. Mayo and Spanos develop a "methodology of mis-specification (M-S) testing" to help researchers address this need. Such a methodology, they observe, should provide

> methods for uncovering and probing model assumptions, isolating sources of any anomalous results, and iterative procedures for accommodating any flawed assumptions in respecified models until arriving at a statistically adequate model—a model that is adequate for subsequent (primary) statistical inferences. (Mayo and Spanos 2004, 1008)

Laid out in some detail by Spanos in his (1999), the methodology of mis-specification testing and respecification can be characterized for our purposes as the attempt to assess a candidate statistical model considered for use in drawing a primary inference. Such an evaluation involves embedding the candidate model within a larger encompassing model that includes alternatives based on different assumptions, and then using the data already in hand to test for departures from the assumptions of the candidate model in the directions reflected in the alternatives included in the encompassing model. A series of such tests will be directed at testing for departures

with regard to the different types of assumptions (distribution, dependence, heterogeneity) that define a statistical model. Should the candidate model fail such a test, this will not be taken as evidence for a specific alternative model, but will rather serve as the occasion for *respecifying* the model, and then reiterating the testing process with the new, respecified candidate model until one has a model that can be vindicated as statistically adequate.

Here one can see the methodologically internalist flip-side of the externalist possibility of defeat by facts that are not accessible. For the justification of an evidence claim the mere *de re* satisfaction of the *ES* requirements is insufficient; the statistical model through which the satisfaction of those requirements is evaluated must in turn be validated as statistically adequate. By carrying out the mis-specification testing required for such validation, the investigator simultaneously acquires the ability to articulate the grounds for accepting that model.

Indeed, one might think that the possibility of mis-specification testing advances the internalist* aspect of scientific justification *so* far as to provide grounds for rejecting our "anti-internalist*" claim *AI.* Because mis-specification testing requires only the data already in hand and the use of readily-available statistical techniques, one might argue, any facts that would result in statistical inadequacy of a model used in making an error-statistical inference *are* accessible, whether or not the investigator avails herself of these techniques.

Granting this point, we note that it shows only that mis-specification testing constitutes grounds for rejecting *AI* as applied to *formal statistical inferences*. This objection, however, does not undermine *AI* as a general thesis about error-statistical

evidence claims, insofar as some such claims pertain not only to formal statistical inferences but to inferences to substantive scientific claims going beyond what can be secured via mis-specification testing alone (see Mayo and Spanos 2006, 341–42). Often a statistical hypothesis "stands in for" or is a model of a substantive claim about causal or other relationships (e.g., "these data include decays of the Higgs boson" or "these structures are fossilized remains of pigment-producing organelles"), such that an inference to that claim involves non-statistical assumptions regarding the adequacy of the experimental set-up, or the adequacy of the statistical hypothesis as a model of the substantive claim. Mis-specification testing underwrites the statistical reliability of the procedures used to assess substantive hypotheses or claims of interest, but does not in general subvent substantive reliability.[10] Mis-specification testing does not underwrite a general argument against *AI.*[11]

We regard mis-specification testing as addressing two related problems. First, it helps to prevent the investigator from being misled as to which hypotheses have and have not been severely tested (the problem of *misleading evidence*). Second, it helps the investigator to articulate the reasons that support the use of the statistical model employed (the problem of *justification*).

Our view is that justification in science is externalist* in character insofar as the evidential relations that are of concern in addressing the problem of misleading evidence are objective (as they are on the *ES* view), and internalist* in character insofar as addressing the problem of justification requires the capacity to access and provide reasons that support one's inferences from the data.

A thoroughgoing externalist, of course, would not accept our identification

between the problem of justification and the question of one's ability to articulate supporting reasons, for on an externalist* account one can be justified in drawing conclusions even if one cannot access any reasons that support such a conclusion. Have we not simply assumed an internalist* point of view in the way we frame the question?

In reply to this concern, we should first restate that our concern is with justification in the *socially situated* contexts of scientific inquiry and communication; it is the nature of these contexts, and not a prior commitment to internalism*, that grounds our understanding of the problem of justification. As we noted in section two, investigators drawing conclusions from data are responsible for vindicating their assertions and inferences in response to critical questioning from the community of investigators. In the absence of such a capacity for vindicating a conclusion, an investigator may be able to make statements that are objectively supported by evidence, but does not, thereby, contribute to the scientific pursuit of knowledge.

Moreover, it is not merely mis-specification testing or other methods of model criticism that serve this dual function. Rather, statistical methods in general can be thought of as directed at both the avoidance of being misled and at providing resources for the articulation of justifying reasons. These dual purposes are hinted at as well in some of Mayo's own work:

> [e]rror statisticians appeal to statistical tools as protection from the many ways they know they can be misled by data as well as by their own beliefs and desires. The value of statistical tools is that they allow one to develop strategies that capitalize on knowledge of mistakes: strategies for collecting data, for efficiently checking an assortment of errors, and for communicating results in a form that promotes their extension by others (Mayo 1996, 337).

The error-statistical emphasis on methodology seeks to provide strategies that enable investigators to vindicate their evidence claims by appealing to methods employed that either eliminate errors or take them into account in the final inference. Such vindication employs lines of reasoning based on the characteristics of those very same strategies. The *ES* conditions explain what characteristics such strategies should have (they should be reliable in the sense articulated in *ES*), and thus guide methodological development. But one cannot justifiably infer *H* on the basis of the bare satisfaction of the *ES* conditions alone. Rather, the claimant has to give reasons to show how appropriate methodological precautions against error have been taken (and hence must have access to such reasons). This aspect of internalism* finds a natural place within error-statistics.

The conceptual landscape of our account appears incomplete, however. Evidential relations on the *ES* account are objective in a rather strong sense that they are not relativized to any epistemic situation. Whether data that enable a hypothesis to pass a particular test really do provide evidence for that hypothesis is independent of the epistemic situation of anyone seeking to draw inferences from those data. Yet the justification of any such inferences making use of the evidence is in some sense dependent on such epistemic situations. In the next section, we explicate an epistemic notion that shares this dependence on epistemic situations and illuminates the nature of the additional epistemic work that takes us beyond *de facto* evidence for *H* to the justification of an inference to *H*.

4. Securing Experimental Conclusions

A researcher presents a conclusion from data gathered during research. The

decision to present a conclusion indicates that the researcher and her collaborators are convinced that they are prepared to justify their inference in response to whatever challenges they might plausibly 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 of the conclusions: Are there biases in the sampling procedure? Have confounding variables been ruled out? Is the correct model being employed? To what extent have alternative explanations been considered? Are estimates of background reliable? Can the conclusion be reconciled with the results of other experiments? Have instruments been adequately shielded, calibrated, and maintained?

To a large extent, such challenges can be thought of as presenting possible scenarios in which the experimenters have gone wrong in drawing the conclusions that they do. But such challenges are not posed arbitrarily. Being logically possible does not suffice, for example, to constitute a challenge that the experimenter is responsible for addressing. Rather, both experimenters in anticipating challenges and their audience in posing them draw upon a body of knowledge in determining the kinds of challenges that are significant (Staley 2008). Indeed, as Mayo has argued (1996, 200–203), there is good reason on error-statistical grounds for not regarding the mere logical possibility of error as grounds for rejecting an inference. Such a strategy is highly unreliable in that it always prevents one from accepting a true hypothesis, and in that sense has a maximum error rate.

It would be valuable to articulate some general principles for determining

those challenges to which an experimenter must be able to respond in order to justify an inference from data. Here we merely propose a modest first step toward this aim. We propose, specifically, to articulate a general conceptualization of the problem that such justifying responses address. Our aim is to provide a heuristic that might serve to systematize the strategies that experimenters use in responding to such challenges and allow for a clearer understanding of the epistemic function of such strategies.

Our discussion above highlights certain features that can guide us in formulating the concept at which we aim. Responses to the kinds of challenges we have in mind are concerned with scenarios in which the inference drawn would be invalid; they are posed as more than mere logical possibilities, but as scenarios judged significant by those in a certain kind of epistemic situation, incorporating relevant disciplinary knowledge; and an appropriate response needs to provide a basis for concluding that the scenario in question is not actual.

We conceive of the practices of justifying an inference as the *securing* of that inference against scenarios under which it would be invalid. Here we explicate the concept of security as follows:

> *SEC:* Let Ω_0 be the set of all scenarios that are epistemically possible relative to an epistemic situation K. Suppose that $\Omega_1 \subseteq \Omega_0$. Proposition P is *secure throughout* Ω_1 relative to K iff for every scenario $\omega \in \Omega_1$, P is true in ω. If P is secure throughout Ω_0, then P is *fully secure* relative to K.

Before proceeding, some explanation of terminology is in order. This definition employs the notion of *epistemic possibility*, which can be thought of as the modality employed in such expressions as "For all I know, there might be a third-generation

leptoquark with a rest of mass of 250 GeV/c^2" and "For all I know, I might have left my sunglasses on the train." Hintikka, whose (1962) provides the origins for contemporary discussions, there 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*."[12] Borrowing Chalmers' notion of a *scenario* for heuristic purposes, we use that term to refer to what might be intuitively thought of as a "maximally specific way things might be" (Chalmers forthcoming). In practice, no one ever considers scenarios as such, of course, but rather focuses on salient differences between one scenario and another.

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. Applied to inferences from data, we will say that an inference from data *x* to a hypothesis *h*, based on results of test *T*, is secure relative to *K* insofar as the proposition "data $\boldsymbol{x}_0$ from test *T* are good evidence for *h*" is secure relative to *K*. In the context of the error statistical account, this amounts to making the security of such an inference depend on the security of the principles *ES1* and *ES2* as applied to the relevant evidence claim.

In order to address the pressing concern that we are constructing a useless bit of conceptual apparatus without methodological applicability, let us emphasize two points. First, the notion of a fully secure inference is something we regard as an ideal to be employed only in articulating an account of justification. Second, we do not propose that investigators can or should attempt to determine some *degree of security* of any of their inferences. (Doing so would require, for example, that one determine

just what scenarios *are* epistemically possible for a given epistemic situation, thus drawing us into debates over the semantics of epistemic possibility that we are eager to avoid.)

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. Thus, although we have defined a concept that we call *security*, the methodologically significant notion is not security *per se*, but the *securing* of inferences, which we understand in terms of the use of methods that serve to increase the *relative* security of an inference, either by expanding the range of validity of an inference across a fixed space of possible scenarios, or by decreasing the range of possible scenarios in which the inference would be invalid. One can thus *secure* an inference without ever needing to determine its *degree of security*.

Returning, then, to justification, we wish to relate justification to security in the following way:

> *JUS:* An assertion of H as a conclusion inferred from data $\boldsymbol{x}_0$ on the basis of test T is fully justified relative to epistemic situation K only if:
>
> (1) on the basis of test T, data $\boldsymbol{x}_0$ are good evidence for H (in error-statistical terms, ($\boldsymbol{x}_0$, T, H) satisfy *ES1* and *ES2*); and
>
> (2) the proposition "on the basis of T, data $\boldsymbol{x}_0$ are good evidence for H" is secure throughout all scenarios that are epistemically possible relative to K.[13]

This account articulates a notion of full justification as an epistemic ideal. The point is that methods of justification serve two epistemically distinguishable 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 than they were before.

Such methods therefore can be seen as underwritten by a general methodological dictum for investigators considering a potential inference: Consider those scenarios which, for all you know, might obtain that would invalidate the evidence supporting your inference and take the measures necessary to secure your inference against those scenarios.

5. Is Security a Strictly Internalist* Concept?

In the past three sections we have argued that although scientists ultimately aim to produce evidence that objectively connects data and a hypothesis under investigation, they also seek to vindicate their experimental conclusions by reference to information and evidence available to those within their epistemic situation. The overarching picture this might suggest is one in which evidence is treated as an externalist* notion and security is treated as internalist* notion. We believe that this is false because security itself is not a strictly internalist* notion. In this section, we discuss the internalist*-externalist* duality of the notion of security and then conclude with some brief reflections on the problematic notion of accessibility for epistemic justification in the sciences.

There are clearly internalist* aspects to the notion of security. In particular, it is important to distinguish between satisfying the conditions for evidence and the security

of an experimental conclusion. The former is an externalist* notion concerning the objective connection between experimental test procedures and the truth of the inferences formed on the basis of these results. But even if one's test procedures are truth-conducive, this does not entail that the conclusions one derives from these experimental results are secure since one may not have access to the information essential to making the case that one's test procedures are reliably connected to the truth of the assertion. Even if an investigator's tests are, as a matter of fact, reliable, since the investigator lacks access to this information and is aware of the fallibility of the various assumptions of his tests, there is good reason for the investigator to attempt to make his evidence claim more secure from defeat.

Vindicating an experimental conclusion to the relevant scientific community involves showing, on the basis of available evidence produced through various tests and statistical analyses, that one's inferences are not likely to be defeated due to a false fundamental assumption or a mis-specified model. Given the social structure of scientific inquiry and the fact that the demand for security emerges as a response to the critical scrutiny an assertion must undergo to be added to the corpus of scientific knowledge, the argumentative practice of vindicating one's primary evidence claim requires an appeal to information that is available within the relevant epistemic situation. To secure an experimental conclusion requires that the reasons ruling these scenarios out are made accessible to the scientific community and this is clearly an internalist requirement. As such, accessibility is (at least) a necessary condition for one's inference to be secure.

But security, as we have defined it, is not a strictly internalist* notion because

the attempt to secure an experimental conclusion is not sufficient for the conclusion to be secure. Access to the information and evidence available within a particular epistemic situation is not sufficient for an experimental conclusion to be secure. This follows from the fact that security is an objective notion—that is, whether some inference is secure relative to an epistemic situation *K* depends objectively on the scenarios that are epistemically possible relative to *K*. Hence, the claim that a scientific assertion is relatively secure can be *defeated* by facts inaccessible to those making the assertion. As such, an investigator or collaboration, having taken steps to secure an inference, can be mistaken in thinking that the evidence claim or inference is relatively more secure than it was prior to employing these methodologies. Thus, security depends not merely upon information accessible to those within the epistemic situation but also upon factors that are beyond the scope of what is accessible within the relevant epistemic situation. Although accessibility is necessary for security it is not sufficient and, as such, security is not a strictly internalist* notion.

6. Conclusion

The preceding discussion of security provides a fruitful way of connecting our earlier discussion of the intuitive grounds for both the internalist* and the externalist* theses. Recall, that the primary appeal of internalism* was the ability to vindicate an empirical claim against skeptical questions by appealing to accessible evidence grounding the claim. Since the notion of security is partly an internalist* notion it satisfies this intuition. But one of the primary intuitions supporting externalism* was the idea that justification needs to be properly connected to truth; in fact, the disconnect between what is available to one within an epistemic situation and what is

relevant to the justification of an assertion is one of the primary deficiencies of the internalist* thesis. Since security is not a strictly internalist* notion, it satisfies the intuition that justification ought to be connected to truth in a strong way.

What, then, is the epistemic significance of the internalist* aspects of security? As we noted above, contemporary epistemologists tend to construe accessibility in terms of what is available to the epistemic subject on the basis of reflection alone. Our emphasis upon collaborations and the socially-situated nature of scientific inquiry (section two) precludes thinking of access in this way. Instead, we proposed that accessibility is a relation holding between that relevant scientific community and the evidence available within its epistemic situation. Given the arguments of this paper, it should be clear that accessibility concerns the availability *to a relevant scientific community* of reasons sufficient to vindicate an experimental conclusion in the face of legitimate questions about its justificatory status. Access thus become distributed, reflecting the broader distribution among the group members of the relevant epistemic tasks that must be undertaken in order to produce and secure evidence (Giere 2002).

The significance of this social structure reinforces the importance of access for scientists who seek to secure their assertions. When an experimental conclusion is advanced, the audience to which it is directed may raise questions about ways in which the underlying assumptions on which that evidence claim rests might be wrong. The claimant needs to be able to address these questions, and the security of an evidence claim might be thought of as measuring how well a collaboration can do this in principle. The securing of experimental conclusions thus is a manifestation of the capacity of the collaboration to defend their claims. This is, of course, a fallible and

corrigible process and it is clearly possible that the purported defenses or attempts at securing an assertion will fail for reasons that might not be accessible at the point the collaboration is asserting their conclusions. But the vindication of these claims requires reference to the available information if scientists are going to proffer their assertions as contributions to scientific knowledge. Absent a defense of their assertions, scientists within the relevant epistemic situation may continue to raise legitimate concerns about the epistemic status of the proposed assertion.

Hence, epistemic justification, at least within the context of paradigmatic objectivist theories of evidence, requires the ability to defend an assertion from legitimate concerns about its epistemic status. This, in turn, requires appealing to information and evidence that is available within an epistemic situation. Clearly this information might not be connected objectively with the truth and, as such, the attempt to secure an assertion might, in fact, fail to secure the claim. There is no guarantee that the methods employed will be truth-conducive, but security itself increases the epistemic standing of the evidence claim relative to the epistemic situation at issue. So, while the appeal to available information may not be sufficient for epistemic justification in the context of objectivist theories of evidence, it is necessary. Likewise, establishing an objective connection between justification and truth is necessary for justification in an objectivist account but it is not sufficient by itself for an assertion to be justified. Justification requires both internalist* and externalist* elements.

Acknowledgments. We would like to thank two anonymous referees for *Synthese* for

their helpful comments. An earlier version of this paper was presented at the Second Meeting of the Society for the Philosophy of Science in Practice in Minneapolis, Minnesota. We are grateful for the helpful comments from audience members at our talk, especially Deborah Mayo and Aris Spanos. Aaron Cobb's research was supported by a Saint Louis University Research Fellowship. Kent Staley's research was supported by National Science Foundation Award SES-0750691.

References

Achinstein, Peter (2001). *The book of evidence*. (New York: Oxford University Press)

Alston, William (1989). *Epistemic justification: Essays in the theory of knowledge*. (Ithaca: Cornell University Press)

Baird, Davis (2004). *Thing knowledge: A philosophy of scientific instruments. (*Berkeley: University of California Press)

Bonjour, Laurence and Ernest Sosa (2003). *Epistemic justification: Internalism vs. externalism, foundations vs. virtues*. (Malden, MA: Blackwell Publishing)

Chalmers, David (forthcoming), "The Nature of Epistemic Space." In *Epistemic Modality*, ed. Andy Egan and Brian Weatherson. Oxford: Oxford University Press.

Comesana, Juan (forthcoming), Evidentialist reliabilism, URL: http://philosophy.wisc.edu/comesana/ (retrieved 8/19/2009).

Conee, Earl and Richard Feldman (2001). Internalism defended. (In Hilary Kornblith (ed.), *Epistemology: Internalism and externalism* (pp. 231-260). Cambridge: MIT Press.)

Cox, D. R. (2006). *Principles of statistical inference*. (New York: Cambridge University Press)

DeRose, Keith (1991). Epistemic possibilities. *The philosophical review*, 100, 581–605.

Fumerton, Richard (1996). *Metaepistemology and skepticism*. (Boston: Rowman and

Littlefield)

Giere, Ronald (2002). Scientific cognition as distributed cognition. (In Peter Carruthers, Stephen Stitch and Michael Siegal (eds), *Cognitive basis of science* (pp. 285–99). Cambridge: Cambridge University Press.)

Gilbert, Margaret (1994). Remarks on collective belief. (In Frederick Schmitt (ed.), *Socializing epistemology: The social dimensions of knowledge* (pp. 235-256). Lanham, MD: Rowman and Littlefield.)

Goldberg, Sanford (2007), *Internalism and externalism in semantics and epistemology*. Oxford: Oxford University Press.

Goldman, Alvin (1986). *Epistemology and cognition*. (Cambridge: Harvard University Press)

Goldman, Alvin (1999). *Knowledge in a social world*. (Oxford: Oxford University Press)

Goldman, Alvin (forthcoming). Toward a synthesis of reliabilism and evidentialism? Or: Evidentialism's troubles, reliabilism's rescue package. URL http://fas-philosophy.rutgers.edu/goldman/Papers.htm (retrieved 8/19/2009).

Hintikka, Jaakko (1962). *Knowledge and belief: An introduction to the logic of the two notions*. (Ithaca: Cornell University Press)

Jaynes, E. T. (2003). *Probability theory: The logic of science*. (New York: Cambridge University Press)

Kitcher, Philip (1993). *The advancement of science: Science without legend, objectivity without illusions*. (New York: Oxford University Press)

Kornblith, Hilary (2001), *Epistemology: Internalism and externalism*. Cambridge: MIT Press.

Lele, Subhash (2004). Evidence functions and the optimality of the law of likelihood. (In Mark Taper and Subhash Lele (eds.), *The nature of scientific evidence: Statistical, philosophical, and empirical considerations* (pp. 191-216). Chicago:

University of Chicago Press.)

Longino, Helen (2002). *The fate of knowledge*. (Princeton: Princeton 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. *British Journal for the Philosophy of Science*, 57, 323–57.

McGrew, Timothy and Lydia McGrew (2006). *Internalism and epistemology: The architecture of reason*. (New York: Routledge Publishing)

Pitt, Joseph (2005). Hume and Peirce on belief, Or, Why belief should not be considered an epistemic category. *Transactions of the Charles S. Peirce Society,* 41, 343–54.

Popper, Karl (1979). *Objective knowledge: An evolutionary approach.* (Oxford: Oxford University Press)

Poston, Ted (2008), Internalism and externalism in epistemology, *Internet Encyclopedia of Philosophy*. http://www.iep.utm.edu/int-ext/ (accessed 8/19/2009).

Pryor, James (2001). Highlights of recent epistemology. *British Journal for the Philosophy of Science*, 52, 95-124.

Roush, Sherrilyn (2005). *Tracking truth: Knowledge, evidence, and science.* (New York: Oxford University Press)

Royall, Richard (1997). *Statistical evidence: A likelihood paradigm*. (Boca Raton, FL: Chapman and Hall)

Sober, Elliott (2008). *Evidence and evolution: The logic behind the science.* (New York: Cambridge University Press)

Spanos, Aris (1999). *Probability theory and statistical inference.* (Cambridge: Cambridge University Press)

Staley, Kent (2004). Robust evidence and secure evidence claims. *Philosophy of Science*, 71, 467–88.

Staley, Kent (2005). Agency and objectivity in the search for the top quark. (In Peter Achinstein (ed.), *Scientific evidence* (pp. 154-183). Baltimore: Johns Hopkins University Press.)

Staley, Kent (2007). Evidential collaborations: Epistemic and pragmatic considerations in 'group belief'. *Social Epistemology*, 21, 321–35.

Staley, Kent (2008). Error-statistical elimination of alternative hypotheses. (In Kent Staley, Jean Miller, and Deborah Mayo (eds.), "Error and Methodology in Practice," special issue of *Synthese* 163: 397–408.

Steup, Matthias (1999). A defense of internalism. (In Louis P. Pojman (ed.), *The theory of knowledge: Classical and contemporary readings* (pp. 373-384). Belmont, CA: Wadsworth/Thomas Learning.)

Suppe, Frederick (1993). Credentialing scientific claims. *Perspectives on Science*, 1, 153-201.

Thagard, Paul (1997). Collaborative knowledge. *Noûs*, 31, 242–61.

Tollefsen, Deborah (2002). Challenging epistemic individualism. *Protosociology*, 16, 86–117.

Wheeler, Gregory and Luís Moniz Pereira (2008). Methodological naturalism and epistemic internalism. (In Kent Staley, Jean Miller, and Deborah Mayo (eds.) "Error and Methodology in Practice," special issue of *Synthese* 163: 315–28.)

Williamson, Jon (2008). Objective bayesianism with predicate languages. (In Kent

Staley, Jean Miller, and Deborah Mayo (eds.) "Error and Methodology in Practice," special issue of *Synthese* 163: 341–56.)

Woodward, James (2000). Data, phenomena, and reliability. *Philosophy of Science*, 67, S163–79.

Wray, K. Brad (2002). The epistemic significance of collaborative research, *Philosophy of Science*, 69, 150–68.

Wray, K. Brad (2007). Who has scientific knowledge? *Social Epistemology*, 21, 337–47.

---

[1] See, however, Wheeler and Pereira (2008) and Roush (2005) for interesting exceptions.

[2] We are framing the debate in terms of accessibility although there are alternatives. Some think of the debate as one concerning whether the factors relevant to the justification of a person's belief are either internal or external to a person's mental life (Feldman and Conee 2001). Others think of the debate as concerning whether one ought to accept a deontological account of epistemic responsibilities (Steup 1999). Although these ways of demarcating internalism and externalism may be fruitful, they are orthogonal to our purposes for reasons we develop more fully below.

[3] Several authors (Alston 1988, Comesana forthcoming, and Goldman forthcoming) have been developing hybrid accounts that combine aspects of both internalism and externalism. Such approaches emerge from a way of framing the debate that does not treat internalism and externalism as mutually exclusive and exhaustive. We regard these attempts as steps in the right direction and see the work in this paper as providing further reasons for developing hybrid accounts since justification in scientific inquiry requires both internalist and externalist aspects.

[4] For the purposes of our paper, we are focused on internalism as a thesis about propositional justification. In the literature on internalism and externalism in contemporary epistemology, many distinguish between propositional and doxastic justification (Poston 2008). We think it is appropriate to focus on propositional justification in the context of scientific inquiry because we are concerned with the justificatory practices essential to justifying experimental conclusions, rather than the status of the beliefs of individual scientists.

[5] See Baird (2004) and Pitt (2005) for views that are critical of the idea of knowledge as belief. Popper's notion of objective knowledge results from a somewhat different version of such a critical stance (Popper 1979).

[6] Furthermore, the relationship between scientific knowledge and individual belief is complicated by the fact that the great majority of evidential claims are issued by groups rather than individuals, and the relationship between these claims, the beliefs of the individuals in those groups, and knowledge is itself a contested issue (see, e.g., Gilbert 1994; Staley 2007; Tollefsen 2002; Wray 2007).

[7] 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).

[8] Following Spanos (1999), we are using the term "statistical model" in the same sense as Cox's "family of models" – i.e., to refer to a mathematical structure that characterizes certain aspects of the data-generating process without specifying fully the values of all parameters that describe that structure. Of course, the term "statistical model" would also

be appropriate for referring to such a structure in which the values of such parameters have been specified, but context will suffice to make clear which sense is intended.

[9] We use the term "facts" here in a broad sense, to refer to anything, including states of affairs and regularities, that might render a statistical model inadequate.

[10] We are indebted to an anonymous referee for *Synthese* for this point.

[11] We thank an anonymous referee for *Synthese* for bringing this issue to our attention.

[12] Just how to formulate the semantics of such statements is, however, contested (see, e.g., DeRose 1991 and Chalmers forthcoming). The central claims of the present proposal are independent of disputed issues regarding the semantics of epistemic possibility.

[13] An epistemic situation might be that of an individual, a research group, or a scientific community. Wray (2007) denies that communities (as opposed to research teams) can be bearers of knowledge, which may be thought to render justification relative to a community's epistemic situation problematic. We take no stance on this and do not think that our view has any direct consequences for this issue.