MICHAEL J. SHAFFER

BAYESIANISM, CONVERGENCE AND SOCIAL EPISTEMOLOGY

ABSTRACT

Following the standard practice in sociology, cultural anthropology and history, sociologists, historians of science and some philosophers of science define scientific communities as groups with shared beliefs, values and practices. In this paper it is argued that in real cases the beliefs of the members of such communities often vary significantly in important ways. This has rather dire implications for the convergence defense against the charge of the excessive subjectivity of subjective Bayesianism because that defense requires that communities of Bayesian inquirers share a significant set of modal beliefs. The important implication is then that given the actual variation in modal beliefs across individuals, either Bayesians cannot claim that actual theories have been objectively confirmed or they must accept that such theories have been confirmed relative only to epistemically insignificant communities.

1. INTRODUCTION

Following the standard practice in sociology, cultural anthropology and history, sociologists, historians of science, various epistemologists and some philosophers of science define epistemic communities as groups with shared beliefs, values and practices. Moreover, many of these philosophers argue that the concepts of knowledge and/or rational belief are themselves dependent in some way on such social entities. So the crucial unique feature of this general sort of approach is that such epistemologists define or characterize certain aspects of knowledge – or of rational belief – in terms of social groups. So, the existence of knowledge or rational belief of the sort in question is essentially dependent on the existence of the requisite kinds of social groups. This is in fact one central thesis of social epistemologies of the more radical stripe.[1] Here discussion will be limited to the concept of an epistemic community understood to being defined by shared belief. However, it is dubious that such communities exist in any epistemically interesting sense. This is because beliefs surely vary across individuals in actual communities of inquiry. Once this point is acknowledged, a problem arises for Bayesians who subscribe to the convergence defense of their theory against the

DOI: 10.3366/E1742360008000324

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

charge of excessive subjectivity. This is because that defense requires that the members of the community of inquiry share exactly in common an important set of modal beliefs. Since it is highly unlikely that this condition has ever actually been met with respect to wide groups of individuals, Bayesians must accept that no actual theories have likely been objectively confirmed or they must accept that actual convergence occurs only relative to rather narrowly defined and rather insignificant groups.

2. THE ORIGIN OF THE MYTH: "KUHNIAN" RELATIVISM

At least since the publication of Kuhn's seminal work on the nature of scientific revolutions (1962/1970), it has become popular to endorse a version of social epistemology in the philosophy of science (and elsewhere) whereby the conduct of inquiry is taken to be a matter of the shared beliefs, values and practices of a social group defined in terms of those shared beliefs, values and practices.[2] Moreover, following the early Kuhn's apparent rejection of *any* trans-temporal methodological standards, this take on the matter is often extended to the even more radical view that all there is to knowledge is the shared beliefs, values and practices of groups. In other words knowledge is nothing more than a social construction maintained by social agreement and negotiation.[3] Kuhn tells us, for example, that

> Successive paradigms tell us different things about the population of the universe and about that population's behavior. They differ, that is, about such questions as the existence of subatomic particles, the materiality of light, and the conservation of heat or of energy. These are the substantive differences between successive paradigms, and they require no further illustration. But paradigms differ in more than substance, for they are directed not only to nature but also back upon the science that produced them. They are the source of the methods, problem-field, and standards of solution accepted by any mature scientific community at any given time. As a result, the reception of a new paradigm often necessitates a redefinition of the corresponding science … as the problems change, so, often, does the standard that distinguishes a real scientific solution from a mere metaphysical speculation … (103)

Straightforwardly, Kuhn also tells us at the conclusion of the second edition of *The Structure of Scientific Revolutions* that

> Scientific knowledge, like language, is intrinsically the common property of a group or else nothing at all. To understand it we shall need to know the special characteristics of the groups that create it and use it. (210)

Compare this with Martin Kusch's recent claims that "the existence of knowledge is dependent on the existence of communities" (2002, 1) and, in discussing individual knowers, that "for the communitarian usually there is no such knower." (3)

There are, however, serious problems with such views that become glaringly apparent on close inspection.[4] Most obviously, the concept on which such social

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

epistemologies are based is thoroughly relativistic. However, there is a further and equally troubling point. Specifically, the concept of an epistemic community where there is substantial agreement is an almost pure fiction. Nevertheless, this myth survives in some form not only in the relativistic views of Kuhnian-inspired constructivists such as David Bloor and in the views of communitarians like Martin Kusch, but also in the methodological views of some Bayesians.

Of course, Kuhn's most notable and most controversial contributions to the philosophy of science are the introduction of the concept of the research paradigm, and his claims to the effect that there are no timeless methodological norms that define and govern the conduct of inquiry. In the development of Kuhn's views, these theses are closely related in the following sense. When we properly understand the concept of a paradigm as a historically situated shared set of fundamental beliefs and practices it becomes apparent that this *includes* the methodological standards that traditional philosophers of science had taken to be constitutive of both science itself and of good scientific practice.

Insofar as the adoption of all such beliefs, values and practices are then the result of purely socio-historical factors parochial to a given time, their origin is primarily causal and not purely evidential. As such, methodological standards are more or less just socially adopted and maintained products of the agreement of the community in question. So they have far less to do with objective truth and evidence than more traditional views suggested. Given this radical reading of the early Kuhn, and those who follow it, we are to reject essentialism with respect to science – the view that there is an essence of science – and simply adopt the view that science is a social construct from top to bottom. This inference is of course made on the basis of the historical observation that the features of what we refer to as 'science' have undergone radical change over time.[5] The methods, values and practices that are characteristic of any particular theory or era are therefore all subject to revision in the long run, and all that is supposed to be required for their rejection is prevailing community agreement on the matter. Borrowing Imre Lakatos's (1977) characterization of this radical interpretation of Kuhnian science, science is then merely a matter of "mob psychology".

However, for obvious reasons this view is inadequate as a characterization of the *epistemic* rationality of scientific methodology. This is because it is both a form of relativism and a form of skepticism.[6] Even Kuhn came to recognize this by the time of his (1977). In response to the perceived inadequacies of both Kuhnian and Popperian views of inquiry, a group of particularly influential philosophers of science and epistemologists have turned to developing a purely evidential theory of scientific methodology in terms of partial (subjective) belief and the probability calculus.[7] Such subjective Bayesian epistemologists have made great strides in characterizing scientific inquiry in these terms. The most successful version is that of the subjective Bayesians. However, it is rather ironic that in defense of the objectivity of their view they have subsequently been forced to accept one of the core theses central to Kuhnian-style social epistemology.

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

2.1 What are Epistemic Communities?

In order to dispel some of the ambiguity concerning the general characterization of the concept of an epistemic community as a group with shared beliefs we can more clearly interpret that claim either weakly (W) or strongly (S) as follows:

> (W-community) A scientific community is defined as a group sharing *similar* beliefs in common.
>
> (S-community) A scientific community is defined as a group sharing *exactly the same* beliefs in common.

Moreover, we can further interpret these concepts as being restricted – applying to a specific set of beliefs – or as unrestricted – as applying to all beliefs. Thus there can be restricted W-communities that have similar beliefs about some specific subset of beliefs and unrestricted W-communities that have similar beliefs *tout court*. Similarly, there can be restricted S-communities that exactly share some subset of beliefs and unrestricted S-communities that share all beliefs in common.

As should be apparent, the weak concept of a community is problematic if it is interpreted in the manner just noted. How much similarity is required for membership in the community in question?[8] Need all the members share any particular beliefs, values and/or practices? Absent precise answers to these questions, the concept of a W-community cannot reasonably serve as a clear explication of a scientific community. As such this concept of a scientific community cannot serve any real purpose – whether it is descriptive or normative – without a more serious treatment. In short, the concept is hopelessly vague as it stands. However, as we shall see, although it is plausible to believe that there may actually be communities of this sort, in responding to criticisms, Bayesians have been pressed into accepting the more dubious claim that scientific communities are kinds of S-communities.

The problem then is that the concept of an S-community is realistically empty or that such communities are epistemically insignificant. There are in all likelihood no such things in virtue of the fact that beliefs vary to a significant degree across individuals at a given time and across individuals at different times. Surely there are no *de facto* unrestricted S-communities. But are there any restricted S-communities? Are there the requisite kinds of communities defined by some shared set of specific preferred beliefs that is *epistemically important*? One might be skeptical even of this claim. For example, one scientist's particular beliefs about quantum mechanics (QM), or about the neo-Darwinian theory of evolution by natural selection (DAR) and its competitors, may be appreciably different from those of another even though they are plausibly both members of the some community of scientific inquirers. Beliefs about QM can vary radically and the various disputes about DAR surely point to the fact that there is significant variation in beliefs within even reasonably well-defined scientific communities. For example, genetic selectionists hold that selection operates on genes and not on individual organisms. Other

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

avowed Darwinians argue that selection operates only on individuals. Pluralists argue that selection operates on genes, individual organisms and even perhaps on groups. More radically, structuralists even oppose DAR with respect to how species are to be defined. Darwinians define species by appeal to phylogeny, while structuralists appeal to morphological similarity in delineating species. With respect to QM, while there is significant agreement about the acceptability of the basic equations of that theory (e.g. the Schrödinger equation), there is considerable disagreement about whether standard QM needs to be supplemented by an additional postulate governing the collapse of the wave function. Of course, the interpretation of QM is also a source of considerable disagreement. At most, what may then be shared by various proponents and practitioners of some theory is no more than a small set of semantically uninterpreted equations or some small set of very general core beliefs about the subject matter in question.[9]

Consider, however, the larger community of physicists or biologists as wholes. Is it plausible to believe that this broader community of inquirers constitutes a restricted S-community defined in terms of sharing exactly the same relevant beliefs about all of physics or about all of biology? What seems to be obvious from even these cursory accounts of two cases is that (1) in virtue of the general variability of beliefs it is rather implausible to believe that there are any *large scale* S-communities at all that share a significant set of beliefs in common and that (2) while it is more plausible to believe that there are *bona fide* restricted S-communities than it is to believe that there are unrestricted S-communities, the plausibility of such a belief generally diminishes greatly as we increase the size of the shared belief set. These points should be obvious because as the size of a belief set grows it will then be harder for members to satisfy the requirements for community membership and this is supported by the general principle of parsimony (i.e. that simpler hypotheses are more probably true). In any case, given this brief attempt to make the concept of an epistemic community a bit more lucid, let us now turn our attention to the role that such communities play in Bayesian epistemology.

3. THE MYTH WITH METHODOLOGY: BAYESIANISM

Recognizing how poorly the myth of the scientific community fares from the epistemic perspective in the context of the general relativistic rejection of methodology, one might suspect that the concept of an epistemic community might well make more sense in the context of epistemological views that do not suffer from the twin intellectual scourges of relativism and skepticism. In other words, we might suspect that there is an epistemologically interesting sense in which there are communities in (something like) the radical social epistemologist's sense, but which are governed by more coherent normative constraints of some methodological views of the nature and function of inquiry. Here this question will be addressed in the limited context of the role that the concept of an epistemic community plays in standard subjective Bayesian methodologies in epistemology

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

and the philosophy of science. In standard Bayesian methodologies, epistemic communities – if they are to have any significance at all – are identified with groups sharing certain crucial beliefs. As such, the Bayesian version of the myth employs the notion of community in the sense of a restricted S-community with precise identity conditions. So is this non-relativistic methodologically inspired version of the myth at all tenable? Let us begin to examine its particular manifestation in the context of subjective Bayesianism in more detail.

3.1 Subjective Bayesianism and the Myth of Shared Modal Beliefs

Subjective Bayesian epistemology is a thoroughly fallibilist and subjectivist view that holds that the goal of inquiry is to converge in the long run on probabilistic agreement concerning epistemic matters. This is supposed to be achieved via the repeated application of Bayes' Theorem to sets of sentences representing the beliefs of individuals (i.e. by Bayesian conditionalization). So, Bayesians accept a methodological view that incorporates a sophisticated theory of theory confirmation based on a specific algorithm. This algorithm allows us to assess the degree of the impact that a given item of evidence has with respect to a given theory. Given this cursory characterization of Bayesian methodology, it might appear to be the case that the concept of the epistemic community plays only a minor role in the Bayesian doctrine. However, this appearance is deeply deceptive, especially as it applies to the Bayesian characterization of convergence. As it is understood in the various Bayesian convergence theorems, *rational* convergence requires community agreement on certain substantive issues, specifically on the set of statements that are possible. This is a deeply important aspect of the Bayesian methodology. This is the case because the primary criticism of subjective Bayesianism is that in endorsing the subjectivity of prior probabilities Bayesian methodology is disastrously subjective and so leads to arbitrarily large disagreement concerning the probabilities of propositions across Bayesian agents. However, before demonstrating that the myth of the scientific community is alive and well in subjective Bayesianism, let us establish a characterization of that view adequate for our purposes.

Standard Bayesian confirmation theory holds that degrees of belief ought to conform to the axioms of the probability calculus. This requirement is referred to as the requirement of coherence, and is typically supported by appeal to various forms of so-called Dutch book arguments. These arguments are designed to show that it would be irrational to have a probability distribution over one's beliefs that did not obey the probability calculus. In any case, following the presentation in Howson and Urbach (1993), the axioms of the probability calculus are as follows:

(A.1) $P(a) \geq 0$ for all a in the domain of $P(\cdot)$.
(A.2) $P(t) = 1$ if t is a tautology.
(A.3) $P(a \vee b) = P(a) + P(b)$ if a and b and $a \vee b$ are all in the domain of $P(\cdot)$, and a and b are mutually exclusive.

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

In A.1–A.3, P(·) is a function whose domain is a complete set of statements closed under Boolean operations. Furthermore, from these axioms Bayes' Theorem can be derived. In its most well known form, Bayes' Theorem holds that

(B.1) $P(h|e) = P(e|h)P(h)/P(e)$, provided that $P(e) > 0$.

In other words, B.1 tells us how well e supports h. The quantity $C(h,e) = P(h|e) - P(h)$ is then a measure of the confirmational impact of e with respect to h. $P(h|e)$ is referred to as the posterior probability, $P(e)$ the probability of the evidence, $P(e|h)$ the likelihood, and $P(h)$ the prior probability of h. The sort of Bayesian confirmation with which we are here concerned interprets the function P(·) as the assignment of credal – or subjective – probabilities over an agent's set of beliefs. This function is governed by the axioms of the probability calculus and the prior probability is the only term in B.1 that is supposed to be wholly subjective.[10] However, the subjectivity of priors is not supposed to be a serious problem. This is alleged to be the case because B.1 is supposed to imply a formal rule for dynamic updating of degrees of belief based on experiential learning in the form of Bayesian conditionalization. Again, following Howson and Urbach (1993), the Principle of Bayesian Conditionalization states that:

> (BC) When your degree of belief in e goes to 1, but no stronger proposition acquires probability 1, set $P'(a) = P(a|e)$ for all a in the domain of P(·), where P(·) is your probability function immediately prior to the change.

Given this view, it can be shown mathematically that convergence on a consensus concerning the probabilities of multiple Bayesian agents will be achieved in the long run provided that the following conditions are met:

(C.1) The agents' degrees of belief satisfy A.1–A.3.
(C.2) All experiential learning is just change of probability governed by strict (Bayesian) conditionalization; i.e. by BC.
(C.3) All the agents initially assign probability 0 to the same beliefs; i.e. they essentially agree on what is possible.

What should then be obvious is that in order to secure some sense of objectivity in subjective Bayesianism, Bayesians must accept a version of the myth of the scientific community in accepting C.3.[11] Abner Shimony (1970) and Leonard J. Savage (1954) are most famous for introducing and defending this view about convergence and arguing that convergence can objectively be secured in the long run provided Bayesian agents agree in the sense of C.3. In short, they must agree on what are seriously proposed statements or hypotheses. What they must agree on, according to C.3, is the content of a set with the following form:

M_n: $\{\Diamond h_1, \Diamond h_2, \Diamond h_3, \ldots, \Diamond h_i\}$.

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

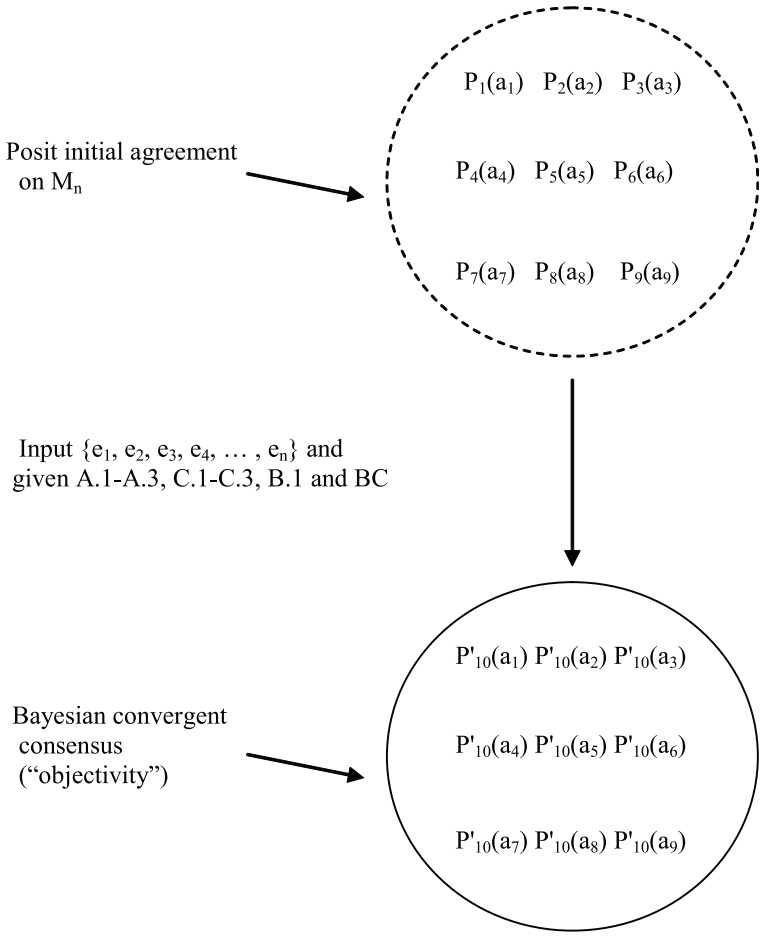


Figure 1. Idealized Bayesian Model of Convergent Inquiry with C.3. Each agent a_i begins with a probability function defined over his/her set of beliefs, $P_i(a_i)$. $P_i(a_i)$ then "evolves" to $P'_i(a_i)$. Recall also that $C(h,e) = P(h|e) - P(h)$.

In cases where this occurs, Figure 1 then illustrates how such agreement is supposed to lead to objective consensus.

However after introducing C.3, Shimony does acknowledge that,

> The phrase "seriously proposed hypothesis" in the definition of "tempered personalist probability function" requires clarification. Although there is often general agreement in the context of a specific investigation as to which hypotheses are to be considered seriously proposed, it is very difficult to state a reasonable set of conditions upon the intrinsic characteristic of hypotheses and upon the circumstances of their proposal which would permit one to distinguish unambiguously those which are seriously proposed from those which are not. (1970, 215)

Difficult indeed! In any case, subjective Bayesians accept the view that a community of inquirers will objectively converge on consensus only if they share all modal

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

beliefs about the relevant hypotheses in common. How specifically does this happen? Let us look at one of the most important and relevant results to this effect.

3.2 A Classic Example of a Convergence Proof

There are a number of different proofs alleged to secure the results noted above, but they all share C.3 in common. Some of the most well-known versions of such convergence proofs are those of Savage (1954), Doob (1971) and Gaifman and Snir (1982). These results are all usefully presented in Earman (1992, ch. 6). Presentation of Savage's particular result here will follow Earman (141–4) very closely. Only minor symbolic alterations are introduced for the sake of consistency. This exercise is nevertheless useful because Earman's reconstruction of Savage's proof is carefully reconstructed so as to be particularly applicable to confirmation theory.

Savage (1954) is duly famous for having introduced the germ of the idea of Bayesian convergence and his basic proof of such convergence can be illustrated as follows. We are to suppose that a set of Bayesian agents accept that there exist K trials of a coin flip that are independently and identically distributed. Moreover, and most importantly for our purposes here, we are to also assume that C.3 holds. Given C.3, all of the Bayesian agents assign non-zero prior probabilities to every member of $\{H_i\}$. Each member of this set, each h_i, is a hypothesis that the objective probability that the result of a coin flip is heads is p_i such that $p_i \neq p_j$ if $i \neq j$. Then, given that Lewis's principal principle holds,[12] that e_j is the report of the result of the jth flip, and that m is the total number of heads results in n flips, all the agents will evaluate the likelihoods as follows:

$$P(\&_{j\leq n}e_j|h_i \,\&\, K) = p_i^m(1 - p_i)^{n-m}$$

We can then select any particular Bayesian agent in the community we are considering and apply Bayes' Theorem to generate the ratio of posterior probabilities of two of the competing hypotheses h_i and h_k as follows:

$$P(h_i|\&_{j\leq n}e_j \,\&\, K)/P(h_k|\&_{j\leq n}e_j \,\&\, K) =$$
$$p_i^m(1 - p_i)^{n-m}P(h_i|K)/p_k^m(1 - p_k)^{n-m}P(h_k|K)$$

Given the strong law of large numbers, the relative frequency of heads approaches the true value of the chance of heads in the limit as the number of flips goes to infinity. Suppose, for example, the real value of that chance is p_R (i.e. the real probability that a coin flip will come up heads), then the likelihood ratio for $i = R$ and $k \neq R$ rapidly grows large for obvious reasons. More importantly this also implies that $P(h_R|\&_{j\leq n}\, e_j \,\&\, K) \to 1$ as $n \to \infty$. Given C.2, it then follows that the probability function $P_n(\cdot)$ for an agent at n who begins with probability function $P_o(\cdot) = P(\cdot|K)$ is $P(h_k|\&_{j\leq n}\, e_j \,\&\, K)$. As $n \to \infty$ the opinions of all Bayesian agents will then converge because they all independently converge on h_R.

So it should be amply clear that C.3 is presented as a sufficient condition for convergence in the context of Bayesian methodology. This, of course, does

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

not imply that convergence of opinion *could not possibly occur* some other way. It could happen, say, by accident, or by stubborn refusal of a group to change its opinions from the outset or by the application of some as yet unknown principles of rationality that are compatible with the basic tenets of Bayesian methodology. But this is only to deny that C.3 is a strictly necessary condition for convergence. Irrespective of this point C.3 is an assumption required for (standard) Bayesians to reach rational consensus *within the purview of the standard accepted Bayesian concept of rationality and absent any known alternative.*[13] Nevertheless, if C.3 fails to be acceptable, then this sort of standard Bayesianism is not even itself a way in which convergence to consensus can be achieved.

3.3 What to Make of the Convergence Proofs

What should we make of this solution to the charge of excessive subjectivity against subjective Bayesianism that crucially involves C.3? First, we should note that Bayesianism says nothing – and can say nothing – methodologically about such agreement. This is because the initial assignment of prior probabilities in Bayesian methodology is utterly subjective. So the solution to the charge of debilitating subjectivity is purchased at the expense of positing such modal agreement without any argument whatsoever. It is simply then a matter of stipulating the *sui generis* existence of a community of inquirers who share a significant set of beliefs. Moreover, *pace* Shimony et al., it is not clear at all that there is ever such agreement if that is meant to indicate a totality of shared modal beliefs with respect to a given domain. The exact content of one scientist's modal beliefs about a given domain of inquiry may, even in relatively ideal circumstances, bear only a weak similarity relation to those of another. As noted in the immediately previous section, as the size of the belief set in question grows, it becomes increasingly unlikely that it is satisfied by any group of significant size.

Given the deeply unrealistic nature of this assumption, we can see that there is an implicit subscription to the myth of the epistemic community that has been incorporated into this methodological project and which is importantly similar to the views of other social epistemologists in this respect. Subjective Bayesians must accept total agreement on the relevant set of modal beliefs (i.e. the set of possible hypotheses) if they are to avoid the charge of utter subjectivism, and so subjective Bayesians must accept a restricted form of the theses that epistemic communities are S-communities with respect to these modal beliefs.[14] However like most myths this is utterly unrealistic. But this myth is also problematic from the perspective of Bayesian methodology itself on two crucial counts because that methodology is construed as aiming at the truth and avoiding error via the concept of probability.

Why is the myth problematic here? First, simply consider what adopting this myth by adopting C.3 without further justification would amount to. What then would occur is Bayesian convergence on some stable probability assignments defined over the space of modal beliefs assumed to be adopted by all Bayesian

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

agents in the initial state, the state where prior probabilities are to be assigned. This may well then involve convergence to consensus with respect to a set of modal beliefs that does not contain the truth with respect to one or more, or even all, scientific problems *as the true hypothesis may have been ruled out as impossible at the outset*. In essence this result is the case because in Bayesian methodology the probabilities assigned to hypotheses can change only when they are non-extreme (i.e. when they are neither 0 nor 1). In mandating such initial modal agreement without any method of justifying modal beliefs, resolving modal disputes or revising modal beliefs (none of which, incidentally, can be done in a manner consistent with Bayesian principles) subjective Bayesianism reveals itself to be deeply conservative, unrealistic and essentially dogmatic. Bayesian progress is constrained by initial and utterly arbitrary assignments of possibility and so Bayesian convergence to consensus may have nothing at all to do with truth.

Second, and perhaps more crucially, if we recognize that the existence of restricted S-communities that share such modal beliefs in common is highly unlikely, then Bayesians are likely faced with the following unhappy dilemma. In virtue of the claims that objectivity is secured by convergence, that *de facto* convergence seems to require substantial *de facto* agreement with respect to modal belief (in the absence of any other alternative account of rational convergence), and since it is highly unlikely that, even in the restricted case of QM and in the case of DAR and its non-Darwinian competitors, all quantum physicists or biologists share exactly the same relevant modal beliefs *tout court*, such Bayesians must accept that at least as far as we know it is likely that either these groups do not constitute relevantly interesting communities or QM and DAR have not been objectively confirmed in the sense of Bayesian convergence.

Neither horn of this dilemma seems like it will be very appealing for Bayesians. On the one hand, accepting that quantum physicists and biologists constitute communities seems to lead to the conclusion that there are in all likelihood no communities of inquirers of sufficient size to be epistemically interesting as the numbers of individuals sharing the relevant modal beliefs exactly in common will be, at best, negligible. This point is, of course, easily generalized beyond these particular cases and the problem is that in opting for this horn of the dilemma, Bayesians would render the appeal to community agreement as fixing *sufficient* objectivity moot. So this tactic would constitute an especially Pyrrhic victory. Why should we think it epistemically relevant to the issue of *objective* confirmation that there is convergent inter-subjective agreement between – potentially numerous – very small groups of inquirers that happen to share the same relevant modal beliefs? The gist of this problem can be seen more easily in Figure 2.

On the other hand, if Bayesians accept the similarly dubious claim that quantum physicists and biologists do not constitute restricted S-communities, then they must, by their own methodological standards, accept that QM and DAR are not objectively confirmed in the standard sense of Bayesian convergence. This follows from the requirement that the only known way convergence can be achieved is

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

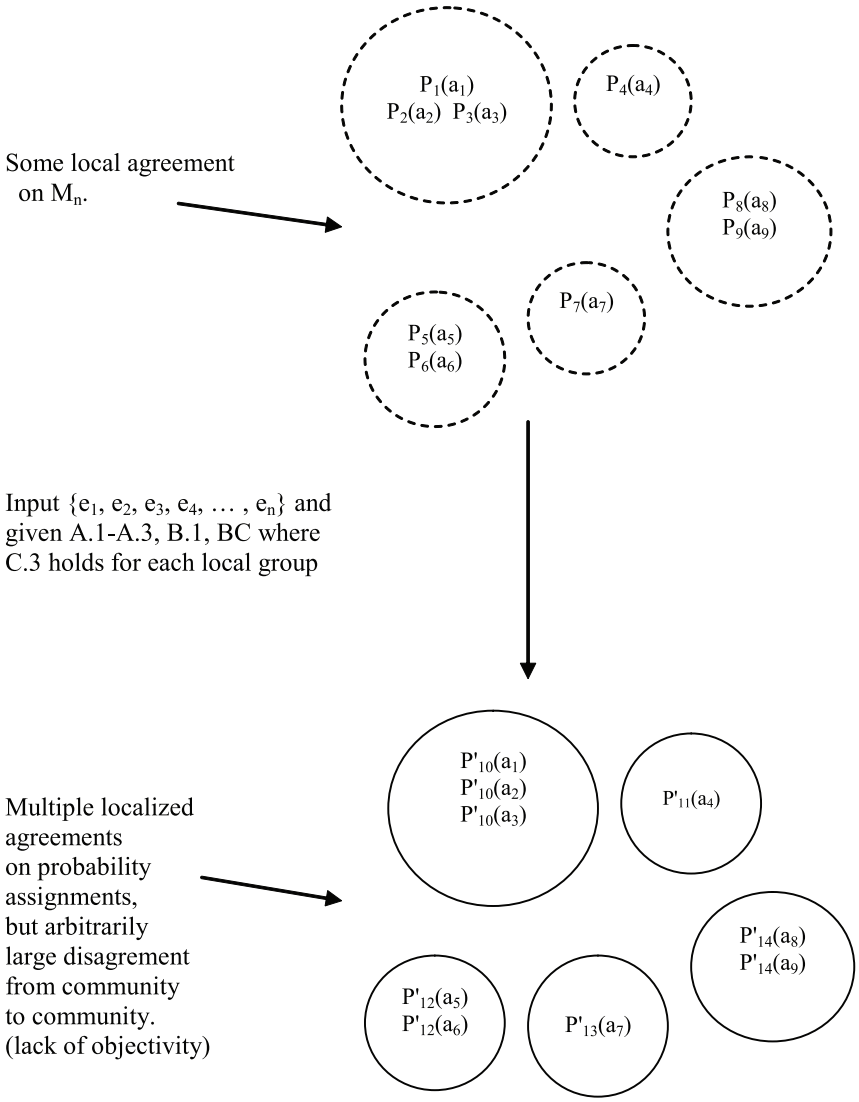


Figure 2. Less Idealized Bayesian Model of Inquiry. Each agent a_i begins with a probability function defined over his/her set of beliefs, $P_i(a_i)$. $P_i(a_i)$ then "evolves" to $P'_i(a_i)$. Recall also that $C(h,e) = P(h|e) - P(h)$.

from an initial state in which there is total modal agreement with respect to the relevant hypotheses being considered. This seems to be a painful pill to swallow as well, for obvious reasons.

3.4 One Possible (But Ultimately Unsuccessful) Solution

One way that Bayesians might suppose that they can avoid the charges leveled above would be to claim that for any epistemic community and some epistemic

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

problem, the members of the set M_n should include *all relevant contingent hypotheses*. In effect, this would be to suppose that every logically possible and non-tautological theory in that domain must be assigned a non-extreme probability (i.e. neither 0 nor 1) by an epistemic community. If this is the case, subjective Bayesians might suppose that epistemic communities will have no difficulties meeting C.3, except where there are disputes over whether a given theory is a logical possibility. Of course, there may well be disputes among the members of a given community about just which hypotheses are logically possible, and this alone is enough to question the adequacy of this solution. If there are such disputes and there is no Bayesian way to settle them, then reaching objective consensus by convergence will be threatened in any case where there is any significant disagreement about such matters. It is also not unreasonable to suppose that there are actually disputes of just this sort in the case of actual epistemic communities, particularly as there is no real consensus about the epistemology of modality. As Gendler and Hawthorne (2002) demonstrates, there is considerable disagreement about how to determine what is and what isn't possible and so there is actual controversy about how to demarcate the possible from the impossible.

More problematically, however, this proposed solution simply involves trading one questionable epistemological myth for another. If we ignore the worries about disputes over which hypotheses are logical possibilities and posit that the members of epistemic communities simply assign non-extremal probabilities to all contingent hypotheses relevant to a given domain of inquiry, then it turns out that subjective Bayesians are committed in a different respect to a deeply unrealistic conception of inquiry. This is simply the case because in order to secure objective convergence, epistemic communities would be forced to consider *every* logically possible hypothesis in the course of their epistemic conduct. But surely this is not the case in real scientific communities. Most logically possible hypotheses have never – and will never – be formulated and assigned probabilities. Moreover, on purely logical grounds, it seems plausible to believe that each complete set of logically possible hypotheses relevant to a domain of inquiry will have an infinite number of members. This would then result in a terribly unrealistic methodological view of the conduct of actual science and other forms of inquiry wherein hypotheses with very marginal probabilities must be treated as serious objects of inquiry. For obvious pragmatic reasons this is never actually the case. Resources are far too valuable and limited to bother with *most hypotheses*. More crucially, this proposed solution should be rejected purely on the grounds that it would imply that most theories that we take to be highly confirmed currently have not been *actually* confirmed at all. This would be the case because most of the alternative hypotheses with which those hypotheses compete have not been considered or, indeed, even formulated by most members of any epistemic community. Again, it is reasonable to believe that this is a rather painful pill to swallow to avoid the charge of excessive subjectivity leveled at subjective Bayesians.

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

4. CONCLUSION

So what are we to make of this? First, the general project of radical social epistemology needs to be more carefully scrutinized. Second, we should be careful not to fall prey to the temptation to conflate objective rationality and consensus. Thirdly, we ought to more closely explore how actual rationality and objectivity come from *diversity* of belief.[15] Finally, and far more specifically, subjective Bayesians should be aware that the convergence defense against the charge of excessive subjectivity is something of a failure. Of course, this does not imply that there may not be some other, more realistic principles of rationality compatible with Bayesian methodology that might yield convergence, but, as we all know, "the proof of the pudding is in the eating". To be sure, the actual force of the case made here does to an extent depend on the empirical question – a question in the sociology of science – of how frequent disputes about which hypotheses are logically possible are when different scientific hypotheses are under investigation. But it is not unreasonable to believe that they are quite frequent and frequent enough to trouble subjective Bayesianism for the reasons given here.

ACKNOWLEDGEMENTS

I would like to thank Peter Lewis, Hal Brown, A. J. Krieder and an anonymous referee for comments on this paper.

REFERENCES

Barnes, Barry. 2003. "Thomas Kuhn and the Problem of Social Order in Science." In T. Nickles (ed.), *Thomas Kuhn*, pp. 122–41. Cambridge: Cambridge University Press.

Barnes, Barry and David Bloor. 1982. "Relativism, Rationalism and the Sociology of Knowledge." In M. Hollis and S. Lukes (eds.), *Rationality and Relativism*, pp. 21–47. Cambridge, MA: MIT Press.

Bird, Alexander. 2000. *Thomas Kuhn*. Princeton: Princeton University Press.

Bloor, David. 1991. *Knowledge and Social Imagery*. 2nd ed. Chicago: University of Chicago Press.

Brown, James. 1989. *The Rational and the Social*. London: Routledge.

Doob, J. L. 1971. "What is a Martingale?" *American Mathematical Monthly* 78: 451–62.

Earman, John. 1992. *Bayes or Bust?* Cambridge, MA: MIT Press.

Feyerabend, Paul. 1975. *Against Method*. London: Verso.

Friedman, Michael. 2001. *The Dynamics of Reason*. Stanford: CSLI.

Fuller, Steven. 2000. *Thomas Kuhn: A Philosophical History of our Times*. Chicago: University of Chicago Press.

Gaifman, H. and M. Snir. 1982. "Probabilities Over Rich Languages," *Journal of Symbolic Logic* 47: 495–548.

Gendler, T. and J. Hawthorne. 2002. *Conceivability and Possibility*. Oxford: Oxford University Press.

Goldman, Alvin. 1999. *Knowledge in a Social World*. Oxford: Oxford University Press.

Haack, Susan. 1998. *Manifesto of a Passionate Moderate*. Chicago: University of Chicago Press.

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

Hacking, Ian. 1999. *The Social Construction of What?* Cambridge, MA: Harvard University Press.

Hankinson Nelson, Lynn. 1993. "Epistemological Communities." In L. Alcoff and E. Potter (eds.), *Feminist Epistemologies*, pp. 121–59. London: Routledge.

Hoyningen-Huene, Paul. 1993. *Reconstructing Scientific Revolutions*. Chicago: University of Chicago Press.

Howson, Colin and Peter Urbach. 1993. *Scientific Reasoning: The Bayesian Approach*. 2nd ed. Chicago: Open Court.

Hull, David. 1988. *Science as a Process*. Chicago: University of Chicago Press.

Kitcher, Philip. 1993. *The Advancement of Science*. Oxford: Oxford University Press.

Kuhn, Thomas. 1962/1970. *The Structure of Scientific Revolutions*. 2nd ed. Chicago: University of Chicago Press.

Kuhn, Thomas. 1977. "Objectivity, Value Judgment and Theory Choice." In his *The Essential Tension*, pp. 320–39. Chicago: The University of Chicago Press.

Kukla, Andre. 2000. *Social Construction and the Philosophy of Science*. London: Routledge.

Kusch, Martin. 2002. *Knowledge by Agreement: The Programme of Communitarian Epistemology*. Oxford: Oxford University Press.

Lakatos, Imre. 1977. *The Methodology of Scientific Research Programmes. Vol. 1: Philosophical Papers*. J. Worral and G. Currie (eds.). Cambridge: Cambridge University Press.

Rescher, Nicholas. 1993. *Pluralism*. Oxford: Clarendon Press.

Rorty, Richard. 1979. *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press.

Rorty, Richard. 1991. *Objectivity, Relativism and Truth*. Cambridge: Cambridge University Press.

Savage, Leonard J. 1954. *The Foundations of Statistics*. New York: Wiley.

Schmitt, Frederick. 1994. *Socializing Epistemology*. New York: Roman and Littlefield.

Schmitt, Frederick. 1999. "Social Epistemology." In J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology*, pp. 354–82. Oxford: Blackwell.

Shimony, Abner. 1970. "Scientific Inference." Reprinted in his *Search for a Naturalistic World View*, pp. 183–273. Cambridge: Cambridge University Press.

Siegel, Harvey. 1987. *Relativism Refuted*. Dordrecht: D. Reidel.

Sperber, Daniel. 1996. *Explaining Culture*. Oxford: Blackwell.

NOTES

1 See, for example, Kitcher (1993), Rorty (1991), Barnes and Bloor (1982), Bloor (1991), Barnes (2003), Hankinson Nelson (1993), Schmitt (1994), Schmitt (1999) and Goldman (1999).

2 Less radical social epistemologies merely focus on issues like the central role of testimony in knowledge.

3 There has been, of course, considerable debate about this interpretation of Kuhn as Bird (2000) and Hoynigen-Huene (1993) both demonstrate, and their point is supposedly supported by reference to Kuhn (1977). Essentially, they argue that Kuhn never believed that there were no atemporal methodological standards. Moreover, there has been considerable debate also about the exact nature of social constructions as addressed, for example, Kukla (2000), Hacking (1999), Sperber (1996) and Brown

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

(1989). However, several points are in order concerning these various but related issues. First, the issues addressed here do not directly concern Kuhn scholarship as such. What is important is what the relevant followers of Kuhn *take* his view to be, and many of his followers take Kuhn's point to be that there are no atemporal methodological standards, regardless of what his actual view was. Second, given Kuhn's proclivity to change his views and retract statements, it is not clear that the Kuhn apologists are even correct in claiming that Kuhn never held the views that science is a social construction and that there are no atemporal methodological standards (see Fuller 2000). There is considerable textual evidence that he did believe this in *Structure* and it should be taken at face value unless there is some overwhelming reason to do otherwise. Third, following Friedman (2001, ch. 3), it is not clear that Kuhn's post-*Structure* attempts to distance himself from such views were successful. Concerning the issue of the nature of social constructions, all that needs to be pointed out here is that the views in question hold that epistemic standards and knowledge itself depend essentially on the existence of social groups and this is not especially controversial among social constructivists.

4 Such criticisms can be found, for example, in Siegel (1987).

5 Notably, Feyerabend (1975) also supported this view in perhaps its most explicit manner. See also Barnes (2003), Barnes and Bloor (1982), Bloor (1991), Rorty (1979) and Rorty (1991) for comparison.

6 See Siegel (1987).

7 The canonical source is Howson and Urbach (1993).

8 Hoynigen-Huene (1993, sect. 3.4) interprets Kuhn as holding that community membership involves learning certain similarity relationships, but fails to address the issue of whether this means that members of such communities must learn exactly the same similarity relationships or only ones that are sufficiently similar to some degree. If the latter view is correct, then we are owed an account of how similar they must be, and if the former is true then it is likely that there are no such communities.

9 It is important to add that it is dubious that theories can be identified with semantics free equations. Theories understood as sets of beliefs must be understood as contentful entities and so variation in interpretations of those equations then also becomes a source of significant disagreement among members of scientific communities. The issue of the interpretation of QM illustrates this well, as there are numerous interpretations of the equations of QM, including the Copenhagen interpretation, the many worlds interpretation, the many minds interpretation, the decoherence interpretation, etc.

10 See Howson and Urbach (1993).

11 The reader is referred to Earman (1992, ch. 6) for both the relevant proofs and for the details of various inadequate attempts to steer around technical criticisms of the convergence results. The details are not especially important to the argument of this paper. What is important here in this respect is that these attempts to secure objectivity by appeal to convergence by conditionalizing share C.3 in common and so are suspect on the grounds that in real cases there is little or no reason to believe that such agreement ever actually obtains. Earman himself alludes to this problem in passing and he suggests that perhaps convergence can be relativized to communities (142), but he does not appear to see the dire implications that accompany such a maneuver. To wit, this approach would then cede its very purpose, i.e. the securing of *objectivity* via conditionalization that yields consensus.

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press

12 See Earman (1992, ch. 2 sect. 7) for a discussion of Lewis's principal principle

13 It would be interesting to investigate how C.3 might be weakened in order to avoid some of the problems raised here. However, my suspicion is that what is needed here is not weaker Bayesian principles that show how convergence can be achieved, but rather what is needed is an adequate "extra-Bayesian" theory of *modal rationality* sufficient to show how modal beliefs posited at the outset of Bayesian inquiry converge prior to the operation of the Bayesian scheme.

14 One might be tempted to argue that Bayesianism incorporates an implicit commitment to fallibilism and so all Bayesian agents assign non-zero, non-unity probabilities to all propositions that are not logical truths. This, of course, would allow one to side-step many of the worries raised here about modal agreement, but this assumption is itself very unrealistic and so would open Bayesian methodology to yet another charge of excessive idealization.

15 Such an approach might begin with historical work like that done in Hull (1988) and might benefit from the more philosophical insights of Rescher (1993).

Michael J. Shaffer is an assistant professor of philosophy at St. Cloud State University. His principle areas of interest are in the philosophy of science, epistemology and logic.

https://doi.org/10.3366/E1742360008000324 Published online by Cambridge University Press