
CA☆ FORUM ON THEORY IN ANTHROPOLOGY

CULTURE AS CONSENSUS

I Against Idealism/ Contra Consensus

by Robert Aunger¹

From a philosophical point of view, the most significant feature of current definitions of culture is the fact that they presuppose either a realistic or an idealistic approach. . . . [Idealists] regard culture as a heritage of ideas that have a transcendent reality independent of the individuals or societies which happen to bear them. . . . [Realists] conceive of culture as an attribute of human social behavior and usually define culture in terms of acquired habits, customs, and institutions.

DAVID BIDNEY, *Theoretical Anthropology*, 1967

At least until recently, idealism has held sway in ethnography (Vayda 1994). As Bidney suggests, it is both historically and naturally linked to the idea of culture as shared knowledge, while realism has been affiliated with a conception of culture as socially learned information. Idealism and its associated definition of culture are, however, fundamentally flawed. This is not a purely philosophical or academic matter, since both the means and the goals of ethnography are influenced by the culture concept adopted in practice.

In the following, I first delineate a number of conceptual and methodological problems associated with idealism as an approach to ethnographic description. Then, since purely theoretical arguments are often unconvincing—especially to those with different metatheoretical commitments—I go on to pursue an empirical line of investigation. In particular, I use a recent approach with idealistic underpinnings, cultural consensus analysis (Romney, Weller, and Batchelder 1986), to examine

whether the central conceptual goal of idealism—finding the one true ethnographic depiction of a cultural group from among the conflicting reports and perspectives that constitute ethnographic data—can in fact be achieved.

Cultural consensus analysis has a number of virtues which make it a “best case” for accomplishing this end: as a formal model, it permits testing of its claims, its assumptions are made explicit, and it has been validated using simulated data against “modest violations” of its underlying assumptions (Romney, Batchelder, and Weller 1987:164). However, to my knowledge, the central, idealistic claim of cultural consensus analysis—to provide a *robust* prediction of cultural knowledge—has not been tested using ethnographic data. I therefore report quasi-experimental tests using the method on a suitable ethnographic database. Significant problems of reliability and interpretation are uncovered. Further, the cultural consensus literature is examined in light of this analysis, and evidence is found which corroborates the conclusion that the method does not provide sufficient internal clues to identify consensual beliefs. Since this represents a failure of the best available method for depicting life in cultural groups using idealized constructs, I believe that the problems highlighted in these investigations are significant. In fact, I posit that the road to idealized ethnographic representation is *necessarily* filled with both practical and theoretical potholes. But there is another, less traveled road which provides a smoother ride to a meaningful understanding of lifeways in particular groups. I conclude, therefore, by suggesting that the alternative to idealism, realism, avoids the pitfalls illuminated in the course of this paper and serves as a better foundation for ethnography.

Conceptual Problems with Idealism

Many of the conceptual problems with an idealistic approach to culture are reflected in, and perhaps spawned by, its definition of culture. Definitions of culture as shared knowledge have been around since the beginning of anthropology (see the reviews by Kroeber and Kluckhohn 1952 and Keesing 1974). As Borofsky (1994:331) notes, most current introductory textbooks continue to emphasize the shared nature of culture. Further, since cognitivism gradually conquered behaviorism in the 1960s, culture is now generally viewed as being “in the head” (D’Andrade 1995). This tends to exclude artifacts and other material manifestations of cultural knowledge from the culture concept. Thus cognitivism, together with the shared-knowledge constraint, makes culture equivalent to widely shared beliefs and values. The notion of culture becomes further refined in ethnographic practice when only the coherent and consistent aspects of cultural knowledge are emphasized (in line with idealistic objectives). In effect, culture becomes consensual knowledge.

One contemporary stream of culture-as-consensus theory, processual in nature, is interested in why some

1. King’s College, University of Cambridge, Cambridge CB2 1ST, England (rvazo@cam.ac.uk).

beliefs and attitudes come to be widespread. A representative is Sperber (1991, 1994), who argues that culture theory should explain why only some of the beliefs invented by human imagination become widely distributed. He postulates that widely shared beliefs are those which are easily communicated and minimally transformed in the process. A more restrictive version is that only widespread *normative* beliefs are cultural. D'Andrade (1989a:824), for example, believes that what makes knowledge cultural is not just whether it is shared but whether ignorance or inappropriate use of that knowledge is sanctioned by others in the group: "What makes something a cultural model is the use of it, not how it is learned."

However, a definition of culture as shared, whether normative or not, has a number of conceptual problems. First is the problem of origins. Any newly minted belief necessarily begins at low frequency in a population—a frequency of one. Some beliefs become popular, and therefore today's marginal belief is tomorrow's cultural convention. Do beliefs become more "cultural" as they increase in prevalence from necessarily humble beginnings? This reliance on degree of prevalence to define culture also creates operational difficulties. For example, it suggests that the nature of beliefs changes once they reach a critical degree of prevalence. Thus, it is wrong to argue that "at what point in the continuum of sharedness we decide to call a given schema [or model] 'cultural' is simply a matter of taste" (Strauss and Quinn 1994:293). What is this point—51%, 75%, or 90% of the population? And what is the nature of this change?

For those who adopt the normative view of cultural consensus, this crucial point is when belief becomes sanctioned. But problems may arise if the criteria of normativeness and popularity conflict. Some beliefs are acquired because they are rare (a bias toward beliefs of relatively *low* frequency in the group). Is such an "anti-norm" also a norm, even if it never leads to high sharedness? What if, even after a belief has become popular, most individuals acquire the belief simply on its merits rather than because it is sanctioned? And if norms are defined by the strength of commitment, why can't the source of commitment be nonsocial? Eccentric personal revelations such as belief in having been abducted by extraterrestrials can nevertheless change an individual's life in myriad profound ways (Mack 1994). Requiring cultural beliefs to serve as norms involves complicated distinctions.

Second, the predominant justification for defining culture as shared belief is that it facilitates communication and hence serves as the backbone for social behavior. The underlying assumption is that "people must share some degree of understanding if they are to communicate effectively with one another, if they are to participate in the same tasks" (Borofsky 1994:331). However, what social interaction requires is not commonality of beliefs but general role-playing abilities that lead, iteratively and mutually, to the creation and coordination of complementary expectations in particular social contexts (Wallace 1970, Swartz 1991).

Third, it can also be inappropriate to characterize

groups by some single value when intracultural variation is significant. "For a long time there has been a minor scandal at the heart of the study of culture. . . . Culture is shared knowledge and belief; but when we study human groups, we find that there is considerable disagreement concerning most items" (D'Andrade 1987: 194). In such cases, it is questionable whether any summary measure based on agreement can serve to "represent" a group. In statistical parlance, deviations from the mean dominate any central tendency: too much of the distribution lies "in the tails," so the mean or mode represents only a few of the group's members. It then seems perverse to characterize an entire group by such a minority position.

Fourth, it is not a belief's normative status but its ability to replicate itself in the minds of others that matters to cultural dynamics.² Beliefs which remain only in the heads of those who originate them die with their originators; they do not enter the social circle and are not transmitted to subsequent generations (Cavalli-Sforza and Feldman 1981, Boyd and Richerson 1985). If one admits that cultural dynamics are of interest, then idealistic approaches fall short.

Fifth, since variation is a necessary precondition for endogenous change (homogeneous objects require exogenous forces to be perturbed), idealistic investigations of cultural change must invoke a *deus ex machina*. Postulating such transcendental entities (including ideal informants or groups) is unparsimonious both in the sense of involving an extra step in analysis after describing variability in data and ontologically, since these entities exist in addition to the individuals that make up groups (unless one is willing to relegate individuals to the unreal). Further, to reify the group is to commit the "culturalistic fallacy" (Bidney 1967; Vayda 1994: 320).

Sixth, the importance of consensus could also be defended if a belief's degree of sharedness significantly influenced how people behaved. However, a robust finding from social psychology is that individuals tend to overestimate the degree of sharedness of opinion (the false-consensus effect—see Marks and Miller 1987 for review); since an individual's social network is usually composed of like-minded people, the error consists of generalizing correct knowledge about the beliefs of associates to larger social groups. The ubiquity of this effect suggests that social order can persist even when in-

2. However, normative beliefs may also be cognitively distinct from non-normative ones in their motivational foundations. D'Andrade (1989b:114–18), for example, has argued that widespread beliefs gain new emotional support when they are socially enforced. In effect, individuals begin to think, "I will believe X because I see that most others around me believe X" rather than believing X because of some intrinsic quality of that belief. In such cases, the normative (i.e., shared) status of beliefs can also affect cultural dynamics, if the motivational change associated with normative status results in a new bias toward adopting those beliefs. Once a belief becomes normative, its rate of further reproduction may increase because of the higher probability that it will be both transmitted and adopted by individuals. However, it is not the psychological difference (i.e., salience) itself but the adoption bias that matters to cultural dynamics.

tracultural variability is relatively high—indeed, this bias may increase one's sense of community with a large group despite necessarily limited social contact because people assume that unmet individuals are also like themselves.

Seventh, as Boyer (1994) points out, the ethnographic tradition of describing cultural groups using ideal knowledge systems is problematic because such portrayals describe not the thoughts that people actually have but those that the ethnographer infers are necessary to make sense of what they say and do. The problem is that the ethnographer then treats these constructs as "direct, literal descriptions of people's mental representations, which of course leads to rather extravagant interpretations" (p. 51). In fact, ritualized behavior and some culturally informed beliefs are purposely counterintuitive to demand attention, to set them apart from the backdrop of ordinary life. Thus, many "collective representations," "worldviews," and "[primitive] modes of thought" have to be considered scholastic concoctions rather than descriptions of psychological realities. Indeed, considerable effort has been expended in symbolic anthropology developing sophisticated representations of culture which exist only in the anthropologist's mind. This undertaking has been sustained by the belief that these ethnographers were describing ways of life characterizing whole cultural groups. However, what people actually know does not necessarily exhibit the qualities of idealized systems, such as logical consistency and completeness. Cultures do not smoothly hinge at the joints between belief systems (Barth 1993). Indeed, we probably should not call topical domains of cultural knowledge (e.g., kinship terminology) systematic, because, if we admit that culture is "in the head," then cognitive science tells us that there are limitations on the ways in which people can represent knowledge. Human reasoning relies upon a variety of domain-specific heuristics and biases which deviate from the standards of logic (Hirschfeld and Gelman 1994). Further, knowledge acquired from others will have additional idiosyncracies resulting from problems of access and interpretation of information outside personal experience. The typical ethnographic picture of an integrated, coherent, and stable body of knowledge from which appropriate deductions can always be made is therefore not cognitively realistic (Boyer 1994). Constructing idealistic representations in ethnography involves "filling in" what a number of informants report with implications that may not depict what any collection of individuals actually knows about its culture. Idealism justifies such inferences by arguing that what individuals actually believe is only an imperfect reflection or manifestation of the abstraction, which is more highly valued because of its logical qualities.

Methodological Problems with Idealism: Cultural Consensus Analysis

Some may be unimpressed with the apparent absurdity of describing cultural groups using traits that no mem-

ber of the group need possess and so continue to pursue the idealistic objective of a single best ethnographic representation. However, methodological and interpretive problems also beset such efforts. These problems begin with the fact that the amount of knowledge required to be fully competent in all social tasks exceeds what single individuals can remember (Roberts 1964). As a result, there is necessarily (mostly role-based) specialization in cultural knowledge (Swartz 1991). The methodological consequence is that no individual can serve as an ethnographer's sole informant, with every answer trusted to be correct. Rather, true cultural knowledge—even within restricted domains—must be cobbled together using some principle from the error-ridden information provided by a number of informants. But whom can one trust about which aspect of culture? Each informant's knowledge is imperfect in some unknown way. A crucial assumption is required to escape from this potential hall of mirrors.

The standard recourse in ethnography has simply been for ethnographers to rely on their own corpus of knowledge and understanding, acquired through participant observation, to discriminate between alternative characterizations of cultural knowledge. The ethnographer's judgment can be given a stamp of authority through effective use of a confident writing style (without cavil or hedging), which induces in the reader a sense of consistent and coherent lifeways in the group portrayed (Geertz 1988). Imbuing an ethnographic report with such features, however, requires idealized abstraction from the varied mass of observations and information that ethnographers usually collect. Weber (1971 [1949]), who is particularly clear about the process of idealization, argues that such representations (which he calls "ideal types") are constructed through the analytic accentuation of certain elements synthesized from many exemplars to make the type's characteristic features clear and understandable. He emphasizes that this is a purely logical operation, to be distinguished from aesthetic or moral appreciation of the idealized fabrication as a model of what *ought* to exist. "It is a matter here of constructing relationships which our imagination accepts as plausibly motivated and hence as 'objectively possible'" (Weber 1971 [1949]:498). However, this method is subject to a fatal flaw: the argument that the ethnographer cannot rely on any single individual to tell the whole truth pertains to ethnographers themselves; ethnographers are fallible people too (Clifford and Marcus 1986).

A second recourse is to hand the job of idealization to a formal algorithm, which is presumably not subject to human foibles. This is the approach adopted by cultural consensus analysis (Romney, Weller, and Batchelder 1986). In particular, cultural consensus analysis assumes "that the correspondence between the answers of any two informants is a function of the extent to which each is correlated with the truth" (p. 316). Competent informants, each with good knowledge of a cultural domain, should therefore exhibit a relatively high degree of concordance with each other. Cultural competence is then an individual's degree of agreement with

other informants, and true knowledge becomes consensual knowledge. Further, since culture is defined as the "information pool" which is shared (p. 316), consensual knowledge again becomes cultural knowledge.

Cultural consensus analysis is still an idealistic approach: the two kinds of output it produces—the consensual knowledge set and informant competence measures—constitute examples of the two kinds of representation that ethnographers have traditionally used to create idealized depictions of cultural groups. Culture can either (1) be made into an abstract structure external to the relevant population of individuals (e.g., Durkheim's "social facts") or (2) be personified in the form of a single hypothetical, idealized individual (e.g., the culture-and-personality school, ethnoscience). The first representation of culture tends to serve as an isolated causal model of the cultural system (i.e., how culture "works" in that society) and corresponds to the consensual belief-set made up by cultural consensus analysis from bits and pieces of informant testimony. The second representation, embodied in the "ideal informant" concept, provides a normative description of life in the cultural group (i.e., how people should behave). In cultural consensus analysis, this personified ideal is defined by the scale of cultural competence; a highly competent informant is expected to have a number of personal virtues (such as intelligence and experience) associated with this social ability (D'Andrade 1995:212–13). The question which remains is whether by avoiding the primary methodological problem associated with idealism (the fallibility of human judgment) cultural consensus analysis also avoids the central conceptual problem (providing a robust determination of the single best ethnographic representation of a cultural group).

CULTURAL CONSENSUS IN THE ITURI

Cultural consensus analysis gives ethnographers the ability to ascertain simultaneously (1) whether an unknown belief system constitutes a "high concordance code" (Romney et al. 1986:316), (2) what the consensual beliefs in a topical domain are, and (3) estimates of individual informants' competence in this cultural belief system. Further, it draws these inferences from standard ethnographic data: responses by a small sample of informants to a suite of questions designed to tap knowledge in some particular topical domain. Formally, the method factors a matrix which measures the cultural similarity of all possible pairings of sampled individuals in the group. If the factor-analytic output exhibits several characteristics, the domain can properly be considered consensual.³ Under these conditions, the

3. Three criteria are used together to determine whether a particular belief system is consensual: (1) The first eigenvalue from the factor analysis should be at least three times the size of the second. Two nearly equal eigenvalues would suggest two meaningful vectors of variation (i.e., two underlying belief systems). If this condition holds, then individuals' loadings on the first eigenvector estimate their competence in the relevant belief system. (2) Mean

response to a question most likely to be culturally correct becomes the one that receives the highest score when the frequency of each alternative response is weighted by the average competence of those who made that response (individual competence also being estimated by the method). These inferred beliefs can be reasonably applied to the group as a whole, with observed variation being considered methodological or idiosyncratic in origin (Romney et al. 1996:4704).

To test the method's ability to identify consensual beliefs, I used a database concerning the edibility of foods from the Ituri Forest in the northeast corner of the Democratic Republic of the Congo (formerly Zaire).⁴ In particular, these data concern reported food avoidances from formal interviews with a representative sample of 450 individuals from the various ethnic groups living in the study area. In the course of this investigation (reported more fully in Aunger n.d.a), I adopted a tactic that does not conform with standard cultural consensus analytic practice but that shows the fragility of the conclusions that can be derived from the method. In essence, I used a variety of criteria to select subsamples from the available data. Such manipulations represent reasonably "natural" variation in the database that might be obtained—given the vagaries of ethnographic fieldwork—and hence are legitimate tests of the method's ability to produce the same consensus consistently for a cultural group.

To select data subsamples, I first made assumptions about the nature of the data available for testing (e.g., I changed sample composition by removing all the youngest informants). Second, I simulated other decisions required to formulate an empirical test (e.g., by selecting particular sets of questions which might be more representative of domain knowledge or by changing the coding of responses). After each such manipulation I performed the standard consensus analysis on the resulting data to see if substantial disagreements in the characterization of the belief system resulted.

In fact, the method produced significantly different sets of beliefs considered consensual in each case. For example, randomly selecting informant samples of different sizes from one ethnic group produces a degree of variation in the consensual response set which approaches in significance the variation produced by running the analysis on samples from different Ituri Forest ethnic groups. Indeed, it is fairly easy to construct a subsample which does not exhibit the characteristics necessary for food avoidances in the Ituri to be called consensual, even though all of the other analyses conducted indicate that this is a consensual belief system.

competence should be reasonably high; otherwise convergence in responses could be due to personal biases rather than cultural similarity (Weller and Romney 1988). (3) Competence scores should not be negative, since this would also indicate that individuals are competent in a secondary belief system at odds with the more dominant one.

4. The empirical case study can only be summarized here. Full documentation of the claims made in this paper is available in Aunger (n.d. a).

Cultural consensus analysis can be used to construct both types of representation of a cultural ideal. In the light of the above-mentioned problems with the measurement of consensus, the question remains whether the method is also troubled by difficulties surrounding its implementation of the notion of competence and the ideal informant. I asked two questions about its estimates of individual competence in the consensual belief system. First, how stable are these measures? The answer is similar to that for consensual responses: individuals' cultural competence can vary significantly depending on with whom they are grouped: the degree to which their response sets approach that of the group is a function of group composition. Of course, this is sensible. The problem is that it is not clear which group provides the relevant comparison, unless some criterion external to the analysis can provide a theoretical rationale for selection of a particular group (ethnic groups do not necessarily exhibit the highest consensus measures).

Second, I asked, Do the individuals estimated to have high competence also have the characteristics which make it appropriate to call them "ideal informants" or "experts"?⁵ Employing a variety of methods, I found that culturally competent individuals in the Ituri are *not* those with higher intelligence, greater schooling, culturally designated roles as "keepers of traditional knowledge," normality of experience, or central roles in social networks—each of these characteristics having been found to characterize high-competence informants in at least one previous cultural consensus study.

The procedures I used to design my field study and to prepare the resulting data for cultural consensus analysis are not outside the bounds of common practice in such studies.⁶ I therefore conclude that the problems uncovered in the course of this investigation are likely to characterize any such study that undertakes similar testing of its data.

CULTURAL SYSTEMS ≠ CONSENSUAL SYSTEMS

Indeed, if the scope of argument is broadened to include comparisons between other studies in the literature, similar, confirming problems with cultural consensus analysis come to light. In particular, there are cases involving systems that by almost any definition would be called cultural except that they do not exhibit the necessary degree of unanimity. Alternatively, there can also be a relatively high degree of measured consensus in a cultural group despite the lack of a cultural system for the determination of those beliefs.

Two cultural consensus studies (Weller 1984a, Boster and Weller 1990) concerning humoral beliefs (i.e., the application to foods of values such as "hot" or "cold")

qualify as examples of the first possibility. It is very unlikely that individuals constituting a large fraction of a cultural group would independently invent a system of valences for foods from their own personal experience. Cultural transmission between individuals is almost certainly the primary means by which humoral beliefs are replicated, since they derive from Hippocrates's notion of bodily "humors," which then diffused to particular cultural groups (Mathews 1983:827). However, case studies consistently show that there is an insufficient degree of consensus in humoral beliefs to call them cultural. In such cases, cultural consensus theorists are left in a quandary as to what such beliefs should be called (e.g., Weller 1984a).

It is also possible, however, for individuals in a cultural group to show a significant degree of uniformity in response to interview questions despite the lack of a socially defined and culturally transmitted system of belief about that domain of experience. For example, Boster and Weller (1990) found a degree of consistency between American informants in the assignment of valences to foods equal to the consistency in a group with a culturally explicit humoral system. Without the knowledge—derived from being members of that cultural group themselves—that no cultural system for valence assignment exists in America, the researchers would have had to argue that a humoral system exists in American culture. As Boster and Weller (1990:178) suggest, the American group appears to have consistently applied shared cultural knowledge concerning food spiciness and the temperature of prepared foods to the unfamiliar valence-assignment task.

Similarly, Romney et al. (1986:331) suggest that mean competence should also serve as a criterion of consensus, since some studies show a low competence despite a high eigenvalue ratio. For example, they found a relatively high eigenvalue ratio in their study of a "general information task" among American undergraduate students but a low informant-by-informant average correlation and were therefore hesitant to call it a "cultural pattern," presumably because of a feeling that "general information" is not a natural category or domain of knowledge (p. 332).⁷

It therefore appears that significant interinformant agreement is a poor indicator of the existence of cul-

5. The correlation between measured competence and these reasonable measures of expertise has been claimed by Romney (1994: 270) to validate cultural consensus analysis.

6. On the basis of referees' comments, however, I believe this to be a major area of contention.

7. Their "low" average competence is .54, higher than a significant number of consensus values from other cultural consensus studies (Boster 1991, Brewer 1995, Chavez et al. 1995, Johnson, Mervis, and Boster 1992). Nevertheless, these researchers—plus Iannucci and Romney (1994), with average competence less than .54 and eigenvalue ratios of around 2.5:1—claim consensus for their domains, whereas Romney et al. (1986) are loath to do so for the "general information task," probably because that domain is not a "named" category of knowledge but consists of randomly selected questions about history, sports, current events, etc. Obviously, Romney et al. doubt that the domain is "real." However, their "general information" could be seen as belonging to a "pop culture" domain, while Chavez et al.'s (1995) breast cancer etiology might not be a domain for those outside the Western biomedical model. This only points out again that definition of a cultural domain can be fraught with operational difficulties and requires some theoretical rationale.

tural systems of belief. However, a primary use of cultural consensus modeling in an anthropological context is to infer that a belief system is in fact cultural when no a priori expectations are available (Batchelder and Romney 1989:229). Given the unreliability of consensus in a sample of informants as an indicator of socially transmitted beliefs in a domain, cultural consensus analysis seems unsuited to revealing the existence of a previously unknown cultural belief system or confirming the existence of a known one. Thus consensus in beliefs cannot be used to describe the strength of a cultural system, much less determine whether such a system exists at all.

The other primary measure derived from cultural consensus analysis—competence—also cannot be used to determine whether a belief system is consensual or not. D'Andrade (1995:214) points out that cultural conformity effects could be responsible for the correlation between competence and expertise: by definition, those who unquestioningly adhere to cultural standards are more likely to give the normative response and to give it consistently and without hesitation. Nevertheless, according to D'Andrade (1989b: 122), that "one can be an expert in a cultural system implies that such cultural systems are indeed real." However, individuals can show "expertise" (i.e., exhibit a higher percentage of consensual responses than other individuals, while also being more normal, intelligent, etc.) in noncultural systems as well, such as word association tests: "even when there are no right or wrong answers, and no pressure of conformity, and not even an obvious kind of knowledge involved, the same pattern [of association between competence and other desirable individual qualities] is found" (D'Andrade 1995:215).

Because of this inability to distinguish between the significantly different consensus sets and competence values that cultural consensus analysis produces, neither competence values nor especially the "answer keys" to cultural knowledge produced by the method can be trusted. Thus, even if the results from the Ituri study are ignored, I believe that published studies conducted by researchers trained in the method which have passed peer review make the same points: there can be significant intracultural variation underlying an apparent consensus (Johnson and Griffith 1996); a belief system can exhibit high consensus (Boster and Weller 1990) and people can be "expert" in such knowledge (D'Andrade 1995) even though it is not a cultural system and vice versa (Weller 1984a).

This inability to discriminate between consensual and nonconsensual systems is significant because researchers argue that the primary objective of the method is to determine when it is legitimate to aggregate to the "majority view" (Borgatti 1994:275; cf. Weller et al. 1993:115). They consider this move legitimate when there is consensus because cultural systems are defined as commonly shared beliefs (Romney et al. 1986:316) or shared cognitive representations (Romney et al. 1996). Borgatti (1994:276) believes that consensus

theory is an important theoretical development because it provides clear criteria for when and how to aggregate data, "which is the fundamental operation of analysis."

However, even when no consensus is found, it is argued, cultural consensus analysis remains useful. For example, Garro (n.d.), Johnson and Griffith (1996), McMullin (1996), and Weller et al. (1993) believe that it can be used to investigate intracultural variation. But when it is used for this purpose individuals (via their competences) are compared with an idealized representation of the group rather than with each other (unless the agreement matrix itself is analyzed—but this is not typical practice). The argument is that individuals deviate from this construct in ways which can be analyzed to provide new insights into the sources of such deviation. Romney (1994:269, 273) goes so far as to claim that "detailed studies of intracultural variability depend upon estimates of cultural competence" and that "the application of the model makes possible a far deeper understanding of individual differences in cultural knowledge than heretofore." But I have just argued that neither consensus nor competence can reliably distinguish between what these researchers themselves would call cultural and noncultural belief systems, and therefore the point at which aggregation to specific group-level values becomes legitimate is underdetermined. The set of values with which individuals' values should be compared is therefore also in question and the project of characterizing deviance jeopardized.

Further, the method itself makes no direct reference to any theory of consensual belief; this pure instrumentality leaves it without the kinds of a priori expectations that would permit a researcher to choose the correct characterization of variation or consensus from among its outputs. Of course, theoretical concerns and background knowledge acquired through field experience play important roles in constraining the analytical strategies employed by any ethnographer. Certainly, the studies of humoral belief cited above have shown that knowing whether people in a particular group attach social value to choices in a domain of knowledge is still required when using cultural consensus analysis. In fact, researchers must have independent knowledge of whether the system they study is consensual (that is, cultural) *before* they conduct their analysis. Even given the perception that a consensus should be expected, the choice among possible representations of that consensus is uncertain. I therefore submit that cultural consensus analysis is not a valuable addition to the panoply of techniques for the description of cultural groups.⁸

8. I want to emphasize that my point has not been to critique this particular method laboriously. Since it is one of the few formal approaches to cultural analysis and has gained a substantial following, cultural consensus analysis simply provided the most effective foil for my more general argument against idealism. My complaint is not against such methods per se. Especially since the present climate of opinion seems contrary to scientific approaches to the analysis of culture, I do not intend to give comfort to those who would argue against quantitative methods or to be seen as an advo-

The Alternative: Realism

Cultural consensus analysis has nevertheless proven extremely attractive to researchers, probably because its goal has been to rejuvenate and legitimate the traditional ethnographic project: to paint a true representation of cultural life using relatively few informants, independent of the foibles of both anthropologist and informant. It could be seen in this light as a potential panacea for the postmodern malaise in cultural anthropology, returning us to an earlier age of certainty and objectivity.

But my investigations have belied this promise. Instead, I have uncovered a variety of empirical and conceptual problems which issue primarily from the method's implicit idealism. Even if a more reliable algorithm were to be found, it probably would not solve the basic problem: that there is no means besides individual predilection for choosing among the various possible representations that the method produces—even assuming that it is justifiable to argue that the domain in question is cultural on external grounds.⁹ This is because it is the decisions *prior* to activation of the algorithm (associated with study design) which significantly influence the character of the consensus, not the algorithm itself. Alternatively, other justifications for idealism besides expert consensus might be defended using independently recognized theoretical principles. For example, a "democratic" goal might represent a group by its modal beliefs and values. However, I would argue that a fundamental difficulty remains. The need to choose without a clear rationale for choice is likely to characterize any idealistic approach because it necessarily follows from the combination of a particular desideratum (the one true characterization of a cultural group) with the fact that aspects of research design (such as the choice of an informant sample or set of questions to probe knowledge in a cultural domain) do not have a single best solution. This condition implies that arguments against idealism based on cultural consensus analysis can be couched in more general terms.

I conclude from this investigation that the operational and conceptual difficulties associated with idealism are sufficiently dire to require turning to the alter-

native approach, namely, realism. I believe that the realist perspective can avoid both the conceptual and the empirical problems associated with idealism. This ability derives from a definition of culture which is operationally and conceptually clearer than its status as shared knowledge or a norm. I suggest that if a belief is learned from others, then it is cultural; if it is invented or inferred from individual experience, it is not—at least until it is imparted to others (Boyd and Richerson 1985:33; see also Swartz 1982:316; 1991:7).¹⁰ Realism is thus concerned with how culture is learned and, once learned, transmitted between individuals (Aunger n.d. b, Hewlett and Cavalli-Sforza 1986) and groups (Barth 1987, Soltis, Boyd, and Richerson 1995, Guglielmino et al. 1995). By virtue of being learned socially, such knowledge is cognitively distinct and hence relatively easy to recognize: only socially acquired beliefs can be what Sperber (1985:51) calls "semi-propositional." Because of incomplete information about causal antecedents and consequences, semi-propositional beliefs are consistent with a variety of interpretations rather than a single, precise referent. Thus, Strauss and Quinn (1994:293) are wrong to argue that "cultural schemas differ not at all from other schemas except in being shared"—they can be inferentially limited in ways which individually learned schemas are not. I thus argue for a different cognitive distinction for culture from D'Andrade's: not motivational force but inferential robustness. Further, from a social point of view, what makes human culture unique is its ability to accumulate knowledge through generations (Tomasello, Kruger, and Ratner 1993). That function is crucially dependent on the mode of information acquisition, not on its normative status. I therefore want to argue that D'Andrade, as quoted above, has it exactly wrong: it is the transmission of a belief, not its frequency within a population (or its status as a norm), that makes it cultural.¹¹

How, then, would a realistic ethnographic analysis be

cate of purely interpretive work in this area (see Aunger 1995). However, it is important that the tools developed in anthropology be given a theoretical underpinning and include internal controls for methodological biases (e.g., due to the data elicitation technique and observer-based biases).

9. Admittedly, as with any research tradition, there is an art of practice into which users of cultural consensus analysis become enculturated, and in my empirical study (Aunger n.d. a) I violate a number of the unwritten rules of that practice. For example, the construction of samples, recoding of data as missing, and use of multiple coding schemes from the same set of data are irregular. As a result, many more results are presented than are considered legitimate by those who advocate the approach. However, the art of cultural consensus analysis does not cover all of the necessary decisions in research design, so the sensitivity of consensual representations to such decisions remains a problem.

10. The definition of culture as similarity can cause researchers to make peculiar statements. For example, Boster (1991:223), surveying ethnofaunal classification studies, argues that "culturally diverse groups of informants can converge on a single consensus; they can agree (share *culture*) without the benefit of social information transmission. It is ironic that the cultural consensus model may work best when . . . individuals agree by virtue of their independent insights into the task." In fact, in this case, the cross-cultural similarity is due to universal cognitive mechanisms for the perception of animate objects in nature, as Boster himself has argued (Boster and D'Andrade 1989). Thus, in effect, Boster here suggests that cultural consensus analysis is most appropriate for domains which are significantly genetic in origin. Defining culture as socially learned information would guarantee that culture remained a social scientific concept rather than a biological one.

11. To be fair, there have been recent shifts among anthropologists away from an idealist conception of culture (see Borofsky 1994). Indeed, D'Andrade (1995:216) has come to the conclusion that there are two kinds of cultural domains: consensual, sanctioned domains created by "the need to communicate effectively and share expertise" and less shared knowledge systems marking subgroups as somehow distinct. Now "the issue is not 'how shared is culture,' but rather how to understand both distributed and high consensus aspects of cultural knowledge" (D'Andrade 1995:216).

conducted? A general realist approach to the study of cultural belief systems is "cultural epidemiology," outlined by Sperber (1996). This constitutes a fairly fundamental reorientation of the cultural anthropological project. Whereas the traditional ethnographic enterprise has been interpretive (e.g., to depict cultures in a static, idealized fashion), the quest of cultural epidemiology is explanatory: to uncover how social and cognitive mechanisms work over time to produce the distribution of cultural beliefs both within and between cultural groups. Its fundamental premise is that culture consists of meaningful units of information which are replicated during transmission between individual minds. Transmission occurs either through imitation of the behaviors inspired by beliefs or by the communication of signals related to belief content. Beliefs (mental representations) are then transformed into observable phenomena such as rituals or expressed opinions (public representations) just as a genotype determines its respective phenotype. The explanatory goals of cultural epidemiology are to track the life history of these representations as they metamorphose from one form to the other, to understand the psychology of choice among competing beliefs, and to uncover the forces which influence social access to these representations. Together, these factors determine the distribution of beliefs across individuals. Since both mental and public representations have material manifestations (e.g., as sound waves in the case of speech-as-belief-expression and as brain states in the case of mental representations), cultural epidemiology is fully materialistic (Sperber 1996:26).

The similarity of cultural epidemiology to biological evolution is not coincidental. Both cognition and cultural transmission can be studied from a Darwinian viewpoint. Decision making among alternative beliefs is the cultural analog to natural selection in that it leads to the differential replication of beliefs within groups (Boyd and Richerson 1985), while the origins of these cognitive biases can be explained as the result of the evolution of the brain as an information-processing device (Tooby and Cosmides 1992). Cultural epidemiology, although stretching across both psychological and social scientific levels of explanation, is united under a single theoretical umbrella (for an empirical example, see Aunger n.d. c).

"A central part of a theory of natural selection—functional adaptation—is millennia old, universal, and easily grasped by young pre-school children, whereas natural selection [as the differential reproduction of genetic variants in a population] seems to have emerged only when Darwin and Wallace abandoned strongly held ideas of species' having essences" (Keil 1995:266). Cultural anthropology remains stuck in a pre-Darwinian state, with a functionalist/structuralist underpinning, but the same liberation must now occur for it as has already transpired for other sciences. Just as biology as a science separated from folk biology (Atran 1990), so must the study of culture abandon the folk idea that cultural groups have essences.

But does adoption of such an approach mean that we

can never do principled aggregation or cross-cultural comparison? Certainly this would be an unappealing conclusion and a major strike against realism if true. But fortunately we are no longer constrained to the kind of cross-cultural work until recently characteristic of cultural anthropology: the bivariate cross-tabulation of societies (e.g., Ford 1967). Instead we can do quantitative comparisons of structures with techniques such as log-linear modeling for cross-tabulated data (Agresti 1990), quadratic assignment for similarity matrices (Hubert 1987), and various scaling methods for categorical data (Weller and Romney 1990). This means that each group can be represented not by a single value but by a uni- or multivariate distribution of values, thus preserving much of its unique character. Of course, there will always be situations in which one wishes to characterize groups using single values—for example, when data on distributions are not available or when norms with respect to some domain of behavior are to be discussed. In such cases it can simply be made clear what claims are being made about sharedness or the applicability of such characterizations to particular individuals.

Thus, Weber's (1971 [1949]) claim that idealization is forced on the researcher by the myriad manifestations of cultural traits in individual minds no longer has power: this variation can now be treated analytically. The great benefit of idealization claimed by Weber—that it permits a succinct and persuasive representation—also no longer holds in the face of the postmodern critique of such representations (e.g., Clifford 1988, Clifford and Marcus 1986). Traditional practice privileges the outsider ethnographer's perspective to the exclusion of native voices, which presumably are more legitimate or truthful.

A further point made by recent ethnographic critiques is that some individuals tend to be excluded from active participation in culture formation (Carspecken and Apple 1992). Concern with consensus can thus obscure inequalities in social power. Johnson and Griffith (1996), in discussing their finding of underlying variation within an American consensus, quote Keesing's (1987:166) argument that consensus values "may be shared (at least in surface observance) even though they sustain the interests of some and work against the interests of others. We must . . . dig beneath surface consensuality to seek counterideologies and cultural expression of subaltern struggle. The overlay of consensuality, viewed uncritically, can make an anthropology of meaning insidious as well as politically naive." This perspective suggests that the interesting question is not whether consensus exists but who *makes* consensus and how social elites influence mass opinion.

The most important benefit of realism is that it forces the ethnographer to focus on intracultural variation. A powerful objection to idealism is that it can too easily lead to a denial of the metaphysical and ethical preeminence of the individual. As Weber (1971 [1949]:507) notes, the great temptation is to treat ideals as real—indeed, to "do violence to reality in order to prove the

real validity of the construct."¹² This is because idealism is always normative (at least implicitly), as is suggested by the use of such language as "correct" or "true" belief (e.g., Romney, Weller, and Batchelder 1986; Romney et al. 1996:4704). The notion of competence also implies a hierarchy of value—in particular, a gradually increasing degree of approximation to the ideal informant, who represents what everyone in the group should be striving for, because competence is correlated with intelligence, social status, reliability, and so on. However, the devaluation of individuals through

comparison with a fabricated ideal is contrary to a democratic humanism. Since nationalism and xenophobia rely upon the objectification and normative idealization of one's own group in comparison with others, it is a moral imperative that individuals be considered real, their minds unique. They should be treated with respect and valued for their diversity of experience and opinion. Indeed, the survival of humankind as a species probably depends upon the maintenance of and development of tolerance for both intra- and intercultural diversity.

12. For example, Romney et al. (1996:4704), discussing high-consensus results concerning kinship terminology, claim that "the cultural [i.e., consensus] definition [of a domain by cultural consensus analysis and scaling techniques] is a better estimate of what is in the mind of the subject than an estimate of a cognitive representation based on the subject's own responses. This is because of the vastly increased reliability of aggregate measures compared with single measures." Thus, a group-level representation of the individual should be substituted for the individual's own on the basis of reduced error in measuring the group-level construct. The infor-

mant's differences from that construct are implicitly considered to be due strictly to methodological causes. This is because the informant's responses are just a "test," whereas the group-level statistical construction represents the "truth" (p. 4701). As Romney et al. (p. 4704) note, this use of "single pictures based on aggregate data" is common practice in psychological research. But as Boyer (1994) has argued, it is a fundamental fallacy to assume that beliefs or values characteristic of a group actually exist in the minds of individual members of that group.

