

## Week Three: Eliminating Analysis

### 1. Stich on Cognitive Diversity and Analytic Epistemology

Stich (1988) is interested in tracing through the philosophical consequences of the existence of cognitive diversity. As he notes, it is obvious that different people reason in many varied and different ways, and that this kind of difference is accentuated across different cultures. Now many things vary between different people and different cultures; Stich gives language as an obvious example. But there is a salient difference between linguistic differences and epistemological differences. We will, in general, say that each language is as good as another. Actually, we do discriminate slightly in noting how successful different languages are at performing different tasks, but the differences are quite slight. Epistemic differences are another story. We quite definitely regard some epistemic practices as being better or worse than other practices. It would be good to know what grounds these judgements, and indeed if they can be justified at all.

One popular answer is that epistemic practices are justified when they are in reflective equilibrium. The motivation here is Goodman on deduction and induction, and Rawls on ethics. The process of normative discovery is one of constantly checking our beliefs about specific cases against our general rules. Our norms, be they epistemic or moral, are in reflective equilibrium when all of our views about specific cases coincide with what our general theory says we should think about such cases. For a practice to be justified just means that we would endorse it were we in reflective equilibrium. (There are some tricky phrasing matters here if we wish to avoid what Shope calls the conditional fallacy, but these aren't too hard, and are a little off our main interest.)

As Stich points out, the position sketched here is ambiguous in several ways. Stich's interest, and mine, is in a particular disambiguation of it, which it is worthwhile spelling out. First, we are interested primarily not in the justification of particular abstract arguments, but of relatively concrete inferential processes. Secondly, we take the fact that a practice would be endorsed under reflective equilibrium not to be *evidence* that it is justified, but *constitutive* of its justification. Finally, we take this view to claim that it is a conceptual truth that justification is constituted by endorsement under reflective equilibrium. As Stich rightly notes, saying that it is a conceptual truth is consistent with saying that adopting it may involve a slight change in our usage of conceptual terms. (And it is vague just how much change our usage must undergo to conform to the new theory before the theory ceases to be an analysis and starts to become a stipulative reform.)

Stich's first objection to this position is that it is vulnerable to counterexamples. As he notes, at this stage of the paper Stich is quite happily playing the old-fashioned 'analysis and counterexample' game. This is relevant of course because he will criticise the assumptions of this game later in the paper. Does this mean he is inconsistent? No; the argument here is best read as saying, "Granting your dubious methodological assumptions, your theory is still wrong, but the theory needs these dubious assumptions, so it must be wrong." The counterexamples are of people who not only make silly errors, usually in probabilistic reasoning, but they endorse these errors when they are presented to them in abstract

theoretical form. That is, in this part of their mind, they are in equilibrium, because their specific judgements conform to their general theory, but they are nevertheless mistaken. Hence Goodman *et al* are mistaken to think that being in reflective equilibrium constitutes being justified.

Well, maybe they aren't really in equilibrium, appearances to the contrary. Or maybe we should look for a wider concept of equilibrium, where to be in equilibrium a subject shouldn't just have her specific and general beliefs in equilibrium, but a whole host of epistemic, metaphysical, semantic, moral, *etc* in equilibrium, and it seems clear enough that isn't true of our subjects. Or perhaps what matters isn't what her beliefs would say in equilibrium, but what the beliefs of the people she takes to be experts would be were those beliefs in equilibrium. There seems to be indefinitely many ways to patch the reflective equilibrium story if we just want to avoid problem gamblers.

But Stich has a general purpose objection to this move. Even if we can't find actual people who are in equilibrium in this strong sense but nevertheless have unjustified beliefs (probably because we can't find any people at all who are in equilibrium in this strong sense) it still seems possible that there should be such people. Any belief at all, no matter how crazy, can be worked into a maximally coherent system. And that is enough to refute the theory that reflective equilibrium could provide an *analysis* of justification. This shouldn't be too surprising; it looks much too subjectivist to possibly be right, and Stich's counterexamples play on that.

Maybe, though, the problem is just with the details. (Well, assuming the hefty strain of subjectivism is a 'detail'.) There's got to be some property we're picking out when we say some beliefs or inferential processes are justified, and Goodman showed us how to start the project of finding what it is. We haven't finished this project yet, but it is a worthwhile project to carry on with. This is what Stich calls the neo-Goodmanian project, and it is what he most strongly wants to argue against. And it is what we are most interested in.

Stich thinks the neo-Goodmanian project makes a number of false assumptions, and these undermine the worth of the project. We'll look at these in some detail. (The underlined headings are the dubious assumptions.)

#### *There is a unique interpersonal notion of justification*

Maybe when you and I use the term 'justification', we simply mean different things. "I would not be at all surprised to learn that what I mean by terms like 'morally right' and 'freedom' is very different from what the followers of what the Rev. Falwell or admirers of Col. Khadafi mean." (103) I would be surprised if this were true, for two reasons.

First, just because we use words differently does not mean that we mean different things by them. Sometimes the words have a (hidden or overt) indexical element. So when you assert, and I deny, "Parachuting is fun", this is no indication that we mean different things by any of the words. Rather the right thing to say is that the word 'fun' picks out different properties when used by different speakers. This does not mean it means something different in the mouths of different speakers, any more than 'I' means something different in the mouths of different speakers because it picks out a different

person.

Secondly, we want to say that when followers of the Rev say things like, “Homosexuality is morally wrong,” that they are saying something *false*. If they are speaking some other idiolect, where ‘morally wrong’ means, say, ‘disapproved of by some salient book in the Bible’, then they are speaking truly when they make these utterances. This is wholly incredible. I know that sometimes it is indeterminate whether a dispute is a factual dispute or a verbal dispute, a dispute about how the world is or a dispute about how words are, but this doesn’t look like one of the indeterminate cases.

*Digression.* I think it is an unfortunate (and unintended) legacy of Quine and Davidson’s work on charity that there is a great reluctance on the part of modern theorists to attribute falsity to beliefs or utterances. This is most prominent in discussions of reference; most of the examples of putative reference failure in the recent literature seem to me straightforward cases of speakers successfully referring to some object, and saying something false about it. *End of Digression.*

If someone used our normative terms in a way that made them completely disconnected from motivation, they acknowledged that some beliefs were unreasonable but held them regardless, then they may mean something different by their words to us. This example isn’t meant to be fanciful; some people use ‘reasonable’ and like terms to mean ‘endorsed by some authority’, without any concession that they should submit to that authority. In such cases we have a difference in use of our normative terms, but not a difference in normative concepts, because such people do not regard terms like ‘reasonable’ as picking out normative concepts. If, on the other hand, speakers just think that these normative terms apply to different acts and beliefs, in all likelihood they are just wrong. We’ll look at this in more detail when we get to the network analysis of moral terms in Jackson’s book.

In sum, I don’t think Stich has given us much reason for doubting this assumption.<sup>1</sup>

#### *There is a unique intertemporal notion of justification*

We use different methods for forming beliefs in different circumstances. Can anyone really say I use the same epistemic methods when I am trying to work out who will win the World Series as when I am trying to tell whether ethical consequentialism is tenable? But the whole neo-Goodmanian project relies on their being a single concept which is applied to all cases. This is false, so the project collapses. Stich links this back to the cluster theory of meaning endorsed by the later Wittgenstein. He thinks that our notion of justification might be a cluster of closely related concepts, not a single clear concept.

If justification works like this, there had better be some rules for working out which element of the cluster should

---

<sup>1</sup> Though see Ted Sider’s “The Problem of Personal Identity Deflated” for an argument that (for reasons peculiar to that debate) there might be no shared concept of personal identity.

be applied in each different case. Presumably it is irrational to apply the methods of normative ethics to predicting baseball results. If not then there is a single concept of justification after all. But if this is right there had better be some normative concept overlying all these elements of the cluster, saying which should be employed on which occasion. Now why not identify justification with this overarching concept?

Well, one reason why not, at least if we are sticking to the classical paradigm, is that it may well be too gruesome to use. Maybe there is one property which is picked out by 'justification' in this way, but it is one property in the sense that all of the games in the world share one (highly disjunctive) property. So this might be a problem for a classical view of epistemic concepts, but I doubt it's a problem for the neo-Goodmanian project. We know why the classical view of concepts has to restrict itself to not too gruesome attributes, the view becomes trivial otherwise. We don't have any reason for thinking the neo-Goodmanian project should be so restricted, for thinking that the project could not find these gruesome attributes of justification.

### *Intuitions about Justification are Contaminated*

As we saw in the McNamara and Sternberg experiments, subjects have a great deal of trouble discriminating between necessary and contingent attributes of objects. We would like to think that we philosophers are better than the rabble at this job, but maybe this ain't so. If this is right, then our intuitions about what beliefs and processes are justified will depend not only on *a priori* facts about justification, but on the specific nature of this world. As Stich points out, there is every reason to think that our views about justification will be relatively finely tuned to the nature of this world.

It does seem to be something of a leap of faith to think that we can disentangle the conceptual and contingent elements of our beliefs. But remember that the neo-Goodmanian project is a project, and we might be able to make progress on this front. Remember also that the project can survive even if it should turn out that there is no unique disentanglement to be had. Concepts can be vague lots of ways, and perhaps it will be vague whether a certain belief will turn out to be conceptual or contingent. (That is, maybe the content of a particular concept will be vague, and on some precisifications the belief will turn out to be conceptual, and on others contingent.)

I think there is at least something to be said in response to each of the assumptions behind the neo-Goodmanian project as it is currently practiced. These aren't the important issues, however, for as Stich soon makes clear, these are merely preliminary skirmishes. The real interest is in whether anything like this project could succeed.

## *2. Against Analytic Epistemology*

I'll try and set out Stich's argument in the final section as clearly as I can, then make some comments on how well it works.

- (1) Some cognitive processes are justified and some are not.
- (2) Hence there is a (perhaps infinitely long) set of rules separating the justified from the unjustified. These are what

Goldman calls J-rules

- (3) Philosophers disagree on which rules should be on this list (for example, some favour foundationalist rules, some coherentist, some reliabilist, and so on.)
- (4) Some of these philosophers are mistaken.
- (5) Hence there are (meta-)rules for separating the good rules from the bad (again, subject to the proviso that these rules may be infinitely long.) These are what Goldman calls criteria of rightness.
- (6) If these criteria of rightness turn out to be based on conceptual or linguistic analysis, then the evaluative epistemic concepts we should use will be those embedded in our practices, i.e. our thought and language.
- (7) The evaluative concepts embedded in our practices may differ from those embedded in other practices.
- (8) It is possible that these alternative concepts are superior to ours; just as we should use different cognitive processes to those we actually use, we should use different standards for evaluating cognitive processes to those we actually use.
- (9) If justification is ultimately grounded in the evaluative epistemic concepts embedded in our thought and language, it is impossible that those evaluative epistemic concepts can be mistaken, or that alternative concepts can be superior to ours.
- (10) So by *modus tollens* from (8) and (9), justification is not ultimately grounded in the evaluative epistemic concepts embedded in our thought and language.
- (11) So by *modus tollens* from (6) and (10), criteria of rightness cannot be based on conceptual or linguistic analysis.

The qualifications at steps (2) and (5) are so the ‘rules’ may simply be a list of the justified and unjustified cognitive processes or good and bad epistemic rules. It is convenient to have a catch-all term for whatever tracks the separates the wheat from the chaff, so I posit (perhaps degenerate) rules.

The importance of (6) is that the antecedent is meant to cover most work currently done in English-language epistemology. Here’s a common way to do epistemology. We look at someone’s theory, note that it would allow that the gambler’s fallacy is a legitimate cognitive process, intuit that the gambler’s fallacy is illegitimate, and hence give up the theory. Actually, that’s a bit quick. Probably we intuit some principles from which we derive the conclusion that the gambler’s fallacy is illegitimate, but it amounts to the same thing. Or again, if someone’s epistemology allowed me to reasonably believe the moon is made of green cheese, we can deploy our intuition that this would be massively unreasonable to conclude this epistemic theory is mistaken.

The point to note is the reliance on intuition in each case. Without the intuition that the gambler’s fallacy is illegitimate, or that it would be unreasonable for me to believe the moon is made of green cheese, the objection could not get off the ground. As Stich reasonably claims, “If a philosophical project proceeds by offering definitins or “truth

conditions,” and testing them against our intuitions about real and imaginary cases, then the project should be viewed as an attempt at conceptual analysis or explication.” (111, fn. 12) Hence anyone who uses this popular method of testing epistemological claims is subject to the antecedent of (6). And since (11) is the denial of this antecedent, any such person is subject to Stich’s argument.

It isn’t entirely obvious that (7) is true; it might depend on how we interpret the earlier premises. What is being claimed in (7) is that not only the actual belief-acquisition processes we use might differ from the belief-acquisition processes of other cultures, nor even that the standards we use for judging such processes may differ. It isn’t enough to ground (7), for example, should it turn out that the embedded standards happen to be foundationalist in New York and coherentist in Massachusetts. What we need is that the criteria we use to judge such standards are different. Let’s set this out pedantically.

At level one, we may believe  $p$  in New York, while they believe  $\neg p$  in Massachusetts; of course this doesn’t support (7). At level two, we in New York may believe that the question of whether we should believe  $p$  or not is settled by foundationalist considerations, while in Massachusetts they believe it is settled by coherentist considerations. Again, this isn’t enough to support (7). What is needed is that at level three, while we in New York believe in one kind of test for deciding whether foundationalism or coherentism is correct, in Massachusetts they use a different standard. And we need these to differ despite the fact that very few people in either state will have explicitly considered the question. That is, we need the answer to this level three question implicit in our practices to differ.

One might think that for Davidsonian reasons, this couldn’t be possible. It is untenably uncharitable to assign anything other than the correct “criterion of rightness” to each of the cultures in question. Well, I’m not sure about that, but I do think there are interesting questions here which Stich has glided over by simply asserting that “the theoretical disputes emerge at a higher level.”

The interesting step, where I think all the action is, is at (8). The intuitive appeal of (8) should not be understated, but I don’t need to say a lot to convince you of that. It is very appealing, especially to those of us who are not reactionary xenophobes on normative matters generally, to think that our criteria of rightness could be improved. But it isn’t, I think, a position which can coherently be upheld.

The fact to note is that (8) we are making a normative assessment of normative standards. (Or better, of meta-normative standards, since we are using these standards to test more practical normative standards, if foundationalism can be viewed as a ‘practical’ doctrine.) Note especially the appearance of words like ‘superior’ and ‘should’ in my statement of (8). I think that studying Stich’s text closely reveals similarly normative language. Now we can just apply the main move of the argument again; what are the meanings of these normative claims?; what are the grounds of this claim?; in virtue of what is one criteria of rightness superior to another?; and so on.

If the ultimate answer here is that these norms are ultimately grounded in our practices, if they are things which

can be revealed by close inspection of our language, then (8) will be false. As long as our system of norms is in something like equilibrium, any norm will be judged favourably if it is just judged by other standards within the system. (In fact, this qualification may not be necessary given the fact that a criterion of rightness is a meta-norm, but that's a different digression.)

Perhaps, however, Stich is referring to other grounds for norms here. Given the reference to Wittgenstein's view about norms, I doubt this is so, but we should at least look at the possibility. I'm not sure what this something else could be. Without hoping to go through all the possible candidates, there are some fairly general objections we can make to all of them. An alternative ground must be (a) plausible, (b) accessible to us and (c) really an alternative. Should there be norms grounded in facts about Platonic heaven, that would satisfy (a) and I suppose (c), but fail (b) dramatically. I think some of the more naturalistic alternatives Stich would favour will in this case fail (c). Say that norms are grounded in what would best assist the propagation of our species, as some crude versions of evolutionary ethics claim. This seems to make ethics much more *a posteriori* than we supposed. But looked at closely, it claims there is just one criterion of rightness for 'first-level' norms, be they epistemic or ethical, that adoption of those norms would promote survival. Hence those theorists cannot say that there could be a superior criterion of rightness to their own; and to the extent that they say that we share their criterion, they cannot say there could be a superior criterion of rightness to ours. (Interesting questions: if someone is trying to convince us to accept some form of evolutionary ethics, hopefully more subtle than this one, must they accept a common criterion of rightness? If not, on what do our discussions rest?)

So I think there are grounds for denying that (8) is true. Given the 'depth' of criteria of rightness in our system of norms, there is just nowhere to stand from which we can acknowledge the superiority of alternative criteria. (Someone who held a 'view from nowhere'-ism about ultimate norms, but not epistemic norms, could accept Stich's argument here; but that view is manifestly implausible, and there is no reason to shoulder Stich with it.) Stich has a reply to this objection. While agreeing that normative inquiry into our own practices cannot continue indefinitely, he thinks this would be a 'disasterous' stopping place. This is because he thinks there are so many more things that go into 'ultimate' norms than just truth-tracking, or even worse, agreement with intuition. This is the argument against epistemic xenophobia that ends the paper. As Stich acknowledges, he has nothing to say against the committed xenophobe except some fairly unsubtle name-calling. But I think he has under-estimated what can be said on the xenophobe's behalf.

For one thing, the kinds of disagreements Stich is presupposing to exist near the top of page 108 might be more apparent than real. It seems there is a dispute between the person who thinks beliefs should be such as to foster happiness, and the person who thinks that they should be such as to satisfy our intuitions about what is justified. (The latter can be restated, close enough to equivalently, as saying our beliefs should fall within the extension of the predicate 'justified'.) But this appearance may be chimerical. Say that X tells me I should admire Bill Clinton, and Y tells me I should admire the President. These bits of advice appear to be in conflict, but we can resolve them when we find that Bill Clinton is the

President. (The conflict disappears entirely if we find that X gave his advice *because* Bill Clinton is President.) So to if intuitions about what beliefs are justified co-incide with beliefs that foster happiness, the conflict here will disappear.

More problematically, if there is a real dispute here, it's hard to know what could be said on Stich's behalf. Surely at some stage when he is arguing that we should have beliefs which foster happiness, or whatever, he will have to appeal to our notions of what we should do. It's no use just pointing out that empirical facts are thus-and-so. Unless we commit a particularly egregious version of the naturalistic fallacy, some serious argument will be needed to show that given the facts are thus-and-so we should do/believe this-and-that. And since we have assumed in the 'antecedent' that the empirical facts are in, this argument cannot be based on empirical discovery. So it had better be based on intuition. (Exercise: can this last argument be rephrased 'constructively', as a positive argument for the use of intuition, rather than an argument by elimination?)

But Stich still has a challenge, and this might seem to be the primary challenge the last two pages raise. Why think the intuitions here will be distinctively epistemic? Why not think that other normative intuitions play a role? This is what Stich might be suggesting by saying that we shouldn't stop at epistemology. Well, this strikes me as a reasonable objection to some varieties of epistemology, but not particularly to modern analytic epistemology. It seems to me that when conducting the meta-normative inquiry that Stich is discussing, when using criteria of rightness to adjudicate between rival cognitive processes, everything is on the table. We don't, in practice, isolate epistemic norms from other norms, and perhaps Stich's argument gives us a good reason why we shouldn't. To give just one example, Bayesianism (a system of epistemic norms) is based on prudential considerations. Even those theorists who think that in practice prudential and moral considerations have little role to play in a particular cognitive process, nevertheless acknowledge that this is something that must be argued for. And implicitly at least, I think that even these theorists accept that the criteria of rightness do refer to other normative standards. So there is little reason to think that analytic epistemology, as it is actually practiced, stops inquiry anywhere earlier than Stich acknowledges to be a viable stopping point.