Testimony and Autonomy in Mathematics

Kenny Easwaran

November 28, 2011

DRAFT - PLEASE CONTACT ME BEFORE CITING

In "Testimony and Epistemic Autonomy", Elizabeth Fricker makes the important point that most of our knowledge is extremely dependent on the testimony of others - not only would a completely autonomous individual have very little knowledge (if such an individual could even survive as a rational being, given the social requirements of normal human development), but even trying to replace the testimonial foundations of one's knowledge one would be able to get very little. Consider something as basic as my belief that there are kangaroos in Australia. You might say, "but haven't you been to Australian and seen the kangaroos?" Yes, I have, but unless you are a skilled aeronautical navigator, or returning to a place that you are already non-testimonially familiar with, your knowledge of where you are when you get off a plane depends essentially on testimony as well.

This non-autonomy is especially clearly true in the sciences - no matter how specialized a scientist is, she relies on the testimony of others for the data in the experiments that support many of the theories she uses in interpreting her own experiments.

Mathematics appears to be special - although my knowledge of evolution by natural selection, or of the size of the galaxy, or of the age of the earth, could conceivably turn out wrong if I assume a massive conspiracy between all the scientists and textbook authors, this doesn't seem to be true of mathematics. Although I used various textbooks and took classes to initially learn Gödel's theorems, at this point I have worked through every step well enough that it would take skepticism about my own mind to overturn my beliefs. Although my students this semester may be less autonomous for reconstructing all the proofs, they too don't need to take my word for it.

Of course, there is still a lot of dependence on testimony. Many algebraists use the classification of finite simple groups in their own work, but the proof of this result is so massive that arguably no one understands the methods used in every part, much less has followed the whole thing. This happens in more prosaic forms when number theorists use the Davis-Putnam-Robinson-Matiyasevich non-solution to Hilbert's 10th problem, or when string theorists use the completeness of the Reidemeister moves in knot theory, or when anyone uses Fermat's Last Theorem, or the transcendentality of pi, and so on. But for specialists, this dependence on testimony is much smaller - although one gains most knowledge (apart from one's own research) initially from the testimony of others, mathematical testimony seems to work in a special way, getting people to see for themselves, rather than relying on testimony.

In this talk I want to examine the nature of this difference in the role of testimony between mathematics and the other sciences, and I will argue that it can help us understand the role of certain controversial types of mathematical argument, such as computer-aided proof, long collaborative proofs, incomplete proof sketches, and probabilistic proofs.

In an earlier paper ("Probabilistic Proofs and Transferability") I claimed that the role of testimony in mathematics can explain why mathematicians reject probabilistic proofs while accepting these other problematic proofs. Don Fallis ("Probabilistic Proofs and the Collective Epistemic Goals of Mathematicians") raises the worry that there is no individual or collective epistemic goal of mathematicians that can justify this particular restriction on the role of testimony. In this talk I will try to argue that there is in fact a collective goal that can.

1 Testimony

One can gain knowledge from testimony. As I mentioned above, a large part of our knowledge both of the world generally and of mathematics specifically comes from testimony.

How is knowledge generated from testimony? Reductionists, tracing back to David Hume and to the present with people like Elizabeth Fricker, argue that we use our other faculties to figure out who is trustworthy, and then defer to them. Non-reductionists, tracing back to Thomas Reid and to the present with people like JL Austin and Tyler Burge, argue that just as with our other faculties, there is a kind of default entitlement to believe the subject of testimony, in the absence of defeating evidence. Others like Jennifer Lackey argue for some sort of combination of or alternative to these two views. For most of my points, this issue won't matter, although I may consider an issue arising from Burge's view about a priori entitlements if there is time.

These views differ on when it is appropriate to form beliefs on the basis of testimony - does one need positive evidence of the expertise and sincerity of the speaker, or does one merely need to lack defeaters? Fortunately, in the cases I will be concerned with, the positive evidence does seem to be available - one knows from past experience that authors and readers of a particular mathematical journal tend to be well-positioned with respect to the results the journal covers, and they are generally not deceptive. The issues I am dealing with won't get directly to questions of academic fraud, but those questions will raise important challenges for all views.

For now I will follow Fricker's discussion.

Fricker considers a norm to explain knowledge from testimony. If S is a speaker and H is a hearer and S tells H that P, then she says that H properly accepts that P iff the following conditions are all met: (p. 231 of Testimony and Autonomy)

1. S is sincere; 2. S is an expert; 3. S is better epistemically placed than H with respect to P; 4. H recognizes (1,2,3); 5. H is not aware of significant contrary testimony regarding P.

Fricker says that to be an expert means that one would almost certainly have knowledge about P were one to form a belief about whether P. This is not as hard a condition as the term "expert" makes it sound. One is an expert on anything that one knows, and in most cases of ordinary testimony, it is easy for your interlocutor to come to know that you are an expert just by knowing that you wouldn't be saying things that you didn't know.

The fourth condition is a point of contention between reductionists and non reductionists, and like the fifth it won't be relevant for the cases at hand. But in "Second-Hand Knowledge", Fricker argues that this fourth condition shows the "second-hand" nature of knowledge by testimony - the way that H comes to know that P is that by (1) she knows that S believes that P, by (2) she knows that S most likely knows that P, and by factivity of knowledge H thus knows that P. H's knowledge of P is derived from second-order knowledge that S knows that P. Conditions (3) and (5) are just there to make sure that various defeaters are absent.

I don't want to be committed to precisely this story - a non-reductionist will say that 4 is unnecessary, and will explain the knowledge in a different way. But for my purposes, (2) and (3) are the important points, and I think there is some sense in which non-reductionists will also admit that something like them is needed for testimony to give knowledge. Condition (2) means that S has good information regarding P, and condition (3) means that S's information regarding P is better than that of H. Thus, Fricker suggests in "Testimony and Epistemic Autonomy", that we have an important reason to accept testimony in these conditions - it is the only way to take advantage of the best evidence regarding some phenomenon.

If I were to ignore the experts about global warming, and conduct my own experiments with thermometers outside my house for a few years, I would almost certainly come to a less justified conclusion than I would by paying attention to the experts. Their evidence is far better than mine, and an insistence on a kind of autonomy would leave me epistemically much worse off. This is one of her major points in "Testimony and Epistemic Autonomy" - although autonomy has often been seen as a kind of ideal, especially in the early modern period, it would leave us much worse off.

However, she does point out a kind of worry for knowledge by testimony:

> Epistemic dependence on others, while it extends one's knowledge base so enormously, also lessens

> one's ability rationally to police one's belief system for falsity. There are many things a layperson believes for which she would not know how to assess the scientific evidence which supports them, even if presented with it. This being so, these beliefs of one will lack the characteristic sensitivity to defeating evidence, should it come along, which is usually taken to be a hallmark of belief which amounts to knowledge. (p. 242)

Thus, there is a kind of tradeoff - when I defer to experts, I get access to much better evidence regarding P, but at the cost of lacking sensitivity to defeating evidence, were I to come across it later. The experts believe something for a reason. If I run into someone who happens to have new research that undercuts this reason, then I won't necessarily recognize the relevance of the new fact, and won't give up my belief. When I read that chemists have discovered that old ideas about the temperature dependence of isotope ratios in mineral deposits are mistaken, I may realize that this should revise some estimates of past climate information, which should revise our thoughts about global warming. But since my knowledge comes from testimony by experts, rather than direct contact with the evidence and reasoning involved, I don't know whether this should be a serious revision or a minor worry. If I can bring this new information to other experts and engage them in dialogue, I might be able to maintain an indirect sensitivity, but there is some loss of sensitivity when one learns from experts with whom one is not in constant communication.

At any rate, this cost of deference explains why Fricker insists on condition (3) - only if the speaker has better evidence than the hearer is the access to the evidence going to be able to trump the lack of sensitivity to its defeaters.

Thus, if we ignore the sincerity and lack of competing testimony, and assume that the hearer is in a good position to know all the relevant facts about the speaker, then the important conditions are that S should be an expert with respect to P, and S should be in a better position with respect to P than H. If either is lacking, then there may be good motivation for H not to defer to S.

If expertise is missing, then coming to believe that P runs a serious risk of forming a belief that is not knowledge. If S is no better off than H with respect to P, then deference is not likely to improve H's epistemic state. But if S is better off and an expert, then deference is likely to improve H's standpoint and produce knowledge. Thus, bracketing issues of sincerity, competing testimony, and H's knowledge about S, it does seem plausible that S's expertise and being better off than H are necessary and sufficient for H to be justified in deferring to S.

Fricker points to two ways in which S might have better status than H - one is transient and contingent, while the other is deeper. Her examples of the two kinds both involve knowledge about what happened on stage at a concert. The first kind of expertise occurs when your friend was at the concert but you weren't - then she is an expert for what went on there, compared to you. The second kind of expertise occurs when you are both at the concert, but your friend has good vision and you left your glasses at home - in the latter case Fricker argues that you should allow deference to your friend to reverse your judgment about various propositions, while in the former case you only allow it to form judgments where you didn't have any. I won't have much to say about the difference in your deference, but the two different kinds of expertise will be relevant.

At any rate, the fact that testimony requires the speaker to be an expert (even if just a temporary one, as in the case where the speaker also learned the fact via testimony) means that someone earlier in the chain must have knowledge. Jennifer Lackey has apparent counterexamples involving a creationist science teacher who disbelieves evolution but teaches it to his students on the advice of the school board - the students seem to gain knowledge even though the teacher doesn't believe and therefore doesn't know. However, this is not a counterexample to the broader claim - the students only gain knowledge because the school board has knowledge of evolution; the teacher's role is more complex in this case.

This gives a general principle: if something can be known by testimony, then it can be (or could have been) known by some other means as well. (p. 605 of Fricker, "Second Hand Knowledge") (The same is true of memory. These two sources of knowledge thus contrast with perception and a priori reasoning, and other basic sources of knowledge.) The "could have been" part is there for historical facts - there may be events in the past for which the only possible evidence is now missing, apart from the testimonial reports of eyewitnesses. Some may argue that even in this case, a suitable knowledge of physics may allow one to reconstruct the event from the present state of the universe - but if this is wrong, and historical testimony and memory may be the only present sources of information, then the “could have been” clause is important. (This feature of historical events will be important later.)

2 Testimony in math

An interesting thing about published mathematics is that much of this general theorizing about testimony doesn’t really apply. Formal mathematical writing doesn’t generally consist of “tellings” in this sense.

Reductionists like Fricker emphasize that testimony gives a second-hand kind of knowledge - H only comes to know that P via coming to know that S knows that P. Non-reductionists like Edward Hinchman say that H’s a priori entitlement to believe that P comes in part from S’s invitation to trust.

But in mathematics, we can avoid this. The author can actually give the evidence to the reader. (Note that I am talking about “evidence” in a sense that includes formal and informal proofs, not just empirical evidence.) When the author asserts the sequence of claims that form a proof of the conclusion, the reader (if she is a mathematician working in the relevant field) comes to know each of these claims is true, not by means of trusting the author, but rather by seeing for herself that each one is true. And mathematics appears to have a norm of this sort - when writing a paper, one must back up each claim with a proof, so that the reader doesn’t have to trust the author.

The sort of communication that is possible in mathematics is a kind of counterexample to Grice’s ideas on meaning. Grice thought that speaker meaning involves not only an intention to get one’s hearer to believe something, but that they should come to have this belief specifically through recognition of the speaker’s intention. As Grice pointed out, this means that showing someone a photograph of an event is importantly different from telling them about it (and perhaps even from drawing a picture of the event). (Perhaps the clause about photographs needs a modification for Photoshop - last night I saw a photograph of Mahmoud Abbas kissing Benjamin Netanyahu, but I think it was a Benetton ad rather than a real photo.) Mathematics can work like the photograph, except that it is done entirely with utterances, rather than demonstrations.

Tyler Burge makes this point in footnote 19, p. 480, in his paper “Content Preservation”. (He makes a similar point in footnote 10 of “Computer Proof, Apriori Knowledge, and Other Minds”.) He wants to argue that testimony gives the hearer a kind of default entitlement to believe the content of what is said, independent of any positive justification. He draws an analogy between considering a logical truth and hearing a meaningful sentence - in both cases, the content forces itself on one. He notes that there is a disanalogy in that the entitlement to the utterance depends on their being knowledge earlier in the chain, while the justification for the logical truth doesn’t, but he suggests there is a substantial analogy. My purpose for now is to focus on the disanalogy. Just by thinking about a simple logical truth, you come to have knowledge of it. If you are enough of an expert in the relevant area of mathematics, then just by thinking about the claim that a quantifier elimination argument shows that the theory of algebraically closed fields is decidable is enough for you to come to have knowledge of it.

Note that this is very different from science - if a biologist claims that 60% of the control group got better within a week and 80% of the experimental group did, then the reader must trust her for these facts. Once the reader knows these facts about the two different groups, and knows that the groups were selected randomly, then she can conclude on her own that the treatment is significantly better. But she is dependent on the testimony for the data, and the randomization.

You can’t get around this by showing someone a graph of your data - unlike a photograph (in which a purely non-speaker, natural causal system produces the image - or at least did when Grice was around), a graph of your data is, like a drawn picture, or an utterance, an attempt to get one’s audience to believe one’s claim at least in part through recognizing your intention, your expertise, and your sincerity.

My claim is that this difference in testimonial practice is a consequence of a distinction in Fricker’s notion of expertise. Recall that someone is said to be an expert on P if, were she to form a belief with respect to P, it would very likely be knowledge. In mathematics, in most cases, the reader of the paper is likely to be just as much an expert as the author - were they both to consider the next step in the proof, they are both just as likely to come to know it. In science however, this is not the case - were the reader to come to form a belief about what percentage of subjects in the control study got better within a week, it would likely be a random guess.

Of course, in this case, the difference in expertise is entirely of Fricker's first, historical, sort - the author, but not the reader, was present when the experiment was conducted, and thus knows what went on, how the subjects were randomized, and how they looked as they got better or worse, and so on. Had the reader been present for all this, she may have been just as much of an expert of the author. The author is not an expert relative to the reader of the second, more intrinsic, sort.

It is an interesting fact about mathematics that the evidence for a claim can generally be shared, so that expertise of the first sort can generally be eliminated. Thus, I claim there is a norm in mathematics that published claims should be such that all evidence is shared. This is what I called a norm of "transferability" in my earlier paper, borrowing a term from Michael Rabin.

In published mathematics, an author should not *tell* her claims to the reader. Instead, she should provide a proof that will allow an expert reader to come to know the conclusion for herself. I call such a proof 'transferable'.

Using facts about the presumed reader and author, one can give an individual-based justification of this norm. (I will provide a justification that I am more willing to endorse later on.)

The author should not assume she is more of an expert of the second kind than the reader. (There will of course be many individual readers that don't share the author's expertise in her subject matter, but in a professional journal, there will be many readers that do share her expertise.) The author should not seek to maintain expertise of the first sort - since she can provide a transferable proof, this sort of expertise can be eliminated. (Of course, there will be many individual readers that will just accept the author's claim as testimony and ignore the proof, but since some will follow it if it is provided, the author should provide it.)

Thus, the author shouldn't expect the reader to defer to her. Many readers will happen to defer to the author, either because they don't feel like putting in the work to understand the provided proof, or because they are non-experts that can't put in the work. But many readers will share all the expertise of the author.

Fricker's principle is that H should defer to S about P when and only when S is an expert with respect to P and better placed than H is. It seems plausible to further say that S should tell H that P only when S reasonably assumes that H will defer to S on P. (This contrasts telling with showing - telling assumes that H will come to believe that P by trusting S, while showing assumes that H will come to believe that P by considering the evidence that is shown.) Thus, since in mathematics, we see that H is generally just as well-placed as S, we see that S should always show and never tell.

This assumption of expertise on the part of the readers has an interesting consequence - you can *tell* the readers who aren't experts and only *show* the readers who are experts. Thus, one doesn't need to go into deep detail in a proof - a sketch that will allow an expert to come to know the conclusion is good enough. This, I claim, explains why mathematicians allow, in addition to complete proofs, also proof sketches, as pointed out repeatedly by Don Fallis in a series of earlier papers (especially "Intentional Gaps in Mathematicals Proofs").[1]

Another feature of this assumed expertise is that it doesn't depend on the assumed reader being an expert in every part of the proof. In collaborative proofs where neither author can follow the whole thing, and they don't expect the reader to, they still want to present each part in a way that a reader who happens to be an expert in that part will follow it. Thus, there may be mathematical proofs that do in practice require every reader to trust the author(s) for part of the work - but each step is such that an appropriate reader won't have to. This can explain the existence of major proofs like the classification of finite simple groups - although no mathematician autonomously knows the validity every step of the proof, every step is autonomously known by some readers, and can be autonomously known by more, because of the norm not to tell readers but rather to show them, if they are experts.

This sort of proof is transferable - by presenting the proof in this way an expert can come to know the conclusion of the argument, but not in the second-hand way described by Fricker. By assuming a reader

[1]His example is Gödel's second incompleteness theorem - the important step in this proof consists of the assertion that the entire proof of the first incompleteness theorem could be written in a formal system and arithmetized, and the existence of such an arithmetical proof then shows that the consistency of arithmetic can't be proved formally within arithmetic. Of course, no one ever actually give the formal arithmetized proof of the first incompleteness theorem until decades later.

who is an expert of the second kind, and making her an expert of the first kind, one can get the reader to be an expert just like the author, who has her own first-hand knowledge, and doesn't know the claim merely by knowing that the author knows it.

Consider what this means for Fricker's initial motivation for accepting testimony. When an expert tells you something, you can get better reason to believe it than you could otherwise, if you are not an expert. Of course, one loses sensitivity to defeating evidence, but one gains indirect access to this good evidence.

If mathematics is organized the way that I claim, then we don't have this problem - the reader gains *direct* access to the best evidence, which thus *preserves* the sensitivity to defeaters. This gives another reason for authors to publish in this way - it keeps the best aspects of reliance on testimony (access to the best evidence) and also the best aspects of non-reliance on testimony (sensitivity to defeaters for all of one's evidence). This is only possible because the readers are just as likely to be experts as the authors. I will eventually put more weight on this justification for the practice than on the earlier one, and flesh it out more in terms of the community, but the earlier discussion gives some independent motivation for the norm.

3 Probabilistic Proofs and Testimony

These two arguments give different results for a particular class of cases, namely probabilistic proofs. This was the subject of my earlier paper, "Probabilistic Proofs and Transferability". Don Fallis had argued in a series of papers that probabilistic proofs are just as good as the sorts of proofs that mathematicians do accept, and therefore he claims mathematicians are unjustified in rejecting them. In my earlier paper I gave transferability as a reason why they reject them, and my point today is to reply to his response.

A probabilistic proof works by means of something like the Miller-Rabin or Solovay-Strassen method. In both of these methods, one uses the fact that if p is prime, then every number less than it will bear a certain relation R to it; if p is not prime, then at most a quarter of the numbers less than it will bear R to p. Thus, to test if p is prime, one tests a random sample of the numbers less than p. If all of them have R, then we can conclude with very high confidence that p is prime. Don Fallis has argued (and I agree) that if this sample is large enough, then one can come to have knowledge that p is prime.

I will assume a quasi-Bayesian model of belief to explain how this works. (This is discussed more fully in my earlier paper.) It will not be fully Bayesian, because one can have middling credences in mathematical propositions, despite their being consequences of the axioms, which one believes very strongly. One starts with a particular low credence in the claim that p is prime. One then generates a sequence of 100 random numbers less than p, and tests each for the relation R to p. One's credence that all these numbers will bear R to p given that p is prime is 1, and one's credence that all these numbers will bear R to p given that p is not prime is at most $(1/4)^{100}$, because of the randomness of the process that generated them. If all the 100 numbers generated bear R to p, then after updating by conditionalization, one's credence that p is not prime goes down by a factor proportional to $(1/4)^{100}$. If this was enough lower than one's initial credence that p was prime, then one can end up with a justified very high credence that p is prime, and Fallis claims (and I accept) that this credence can amount to knowledge. It may not be a priori knowledge (though see the first footnote in Burge's paper on computer proof for an argument that it might be), but I think Fallis correctly argues that this can't be the relevant distinction between probabilistic proofs and acceptable proofs.

Instead, consider the situation that a reader is in when she reads of this proof. She can follow every step of the calculation and see that all of the numbers generated bear R to p. However, what she can't do is *see* that the numbers tested were generated in a "random" way. If she thinks that for some reason these numbers were generated in a biased way, she may not update with the same probabilities, and may not come to have very high credence that p is prime. Without *some* level of trust in the author, she could reasonably think that the author just randomly generated 1000 numbers, found 100 of them that had the relation, and then showed only those. Even if there is no reason to distrust the author, the fact that trust is needed shows that this step depends essentially on testimony, just as in scientific cases.

Thus, this sort of probabilistic argument has a step that essentially relies on deference by the reader. This violates the requirement of transferability. Thus, although short deductive proofs, long collaborative proofs that no individual can follow, and proof sketches with missing steps can all count as transferable (as long as experts can follow them and find them convincing), probabilistic proofs cannot, and I claim that this distinction gives a reason for the mathematical community to reject probabilistic proofs.

Don Fallis raises an important worry in his new paper, "Probabilistic Proofs and the Collective Epistemic Goals of Mathematicians". He argues that transferability serves no epistemic goal either of individuals or the community, and thus it is no good as a requirement.

I agree with him that it serves no general goal for the individual mathematician. Mathematicians are clearly willing to accept many mathematical claims merely on the basis of testimony, and rejecting probabilistic proofs on the basis of non-transferability would mean banning this.

In fact, the possibility of probabilistic proofs calls into question my earlier, individualistic justification for the norm of transferability. In the case of a probabilistic proof, the fact that is not transferable is the fact that the numbers were generated by a random process. This is a historical fact that the author really is an expert on compared to the reader - the author was present when she generated the numbers, and the reader has no evidence apart from the testimony of the author. (There also are no tests for randomness that can really be applied to show that a set of numbers really was generated randomly, rather than by means of some cleverly designed mathematical algorithm.) Thus, if the justification for the insistence on transferable proofs is merely that the type 1 expertise can be removed in most cases, then this is one of those cases where it can't be, so the reason would go away.

However, I have a second justification for insisting on transferability, which results from taking the notion of the mathematical community in a more robust way.

> Perhaps we should take seriously the idea that [scientific] knowledge is the property not of individuals but of communities. Taking this idea seriously would mean, I suppose, thinking of epistemic norms as applying not to the individuals operating within a scientific community but to the community itself. (Fricker, "Trusting Others in the Sciences", p. 374)

She goes on to say that she understands knowledge by the community in a reductionist way, in terms of the knowledge by the individuals. Thus, she talks about how testimony preserves (and creates, when someone puts multiple testimonies together) knowledge.

I will consider the idea that perhaps we should rather take the community to be a real knower (or at least, a fiction that is part of the social practice of academic knowledge). "The community" is an entity that has access to all and only the reasons that are in published work (at least, as read by relevant experts). If the proofs are transferable, then the community can be said to know the result for itself. If the proofs are not transferable, then the community knows the result only by testimony. The community is taken to be a fallible individual, just like mathematicians - it can be taken in by errors if journal referees are. But because the argument is in print, the community has access to defeaters.

When Kempe gave his faulty "proof" of the four-color theorem in the late 19th century, the mathematical community believed it. However, because his argument was public, other mathematicians were aware of the lemmas it depended on, and thus when a counterexample to one of the lemmas was found, the community was able to return to proper skepticism. By contrast, if a probabilistic proof relied on a faulty random number generator, and the author dies, then the community can't fix its beliefs except by someone else generating a *new* randomized test. Merely discovering a particular flaw in certain random number generators would do nothing specific to shake the community's confidence if the community relied on the testimony of the original mathematician.

One must only give transferable proofs, because these are the ones that preserve sensitivity to defeating conditions for the community. And that does seem to capture what's going on with the probabilistic proof - the only evidence that anyone could ever have that these numbers were generated in a suitably random way is by having been there at the time of generation and seen how it was done, or by trusting the author (or someone else who was present). This is one of those historical questions for which the best evidence is spatiotemporally bound - assuming that the author was alone, then he is the only one who can ever have access to the actual evidence for the claim, and the community must rely on his word. Even if there is no specific reason to doubt his word, it does mean that the community will lack sensitivity to defeaters - if it turns out that some unremarked feature of the way the author generated the numbers tends to produce a biased sample, then the author is the only person who is in a position to respond to this bias. And if the author dies, then the community can never catch the error.

The same is true if the author uses a canonical published list of random numbers, like the RAND corporation's book *A Million Random Digits with 100,000 Normal Deviates*. In this case, the author is deferring her knowledge of the randomness just as much as the reader is, and only the authors at RAND are (or perhaps, were) sensitive to potential defeaters.

Fallis worries that this requirement of transferability is not justified for the community. The sensitivity for particular claims can be regained by individuals who run their own tests verifying the claims of probabilistic proofs.

If we are considering this as an epistemic value for individual mathematicians, I think this is right. But for the community as a fictional entity above the individual mathematicians, this may not make sense - the "community" can't replicate a randomized test, only individuals can. If other individuals repeat the test and publish their result, the community still only has access to testimony.

Should this hypothetical community-agent matter? That is a difficult question. But if it does, then this will answer Fallis' worries.

4 Alternate norms for a community

Mathematicians know lots of mathematics only by testimony. However, this is not the primary means of mathematical knowledge. Similarly, mathematicians know lots of mathematics by other routes. Probabilistic methods can give knowledge - however, no mathematician other than the originator can know this directly; they must know it only through testimony. More specifically, the fact about the random generation of the test numbers can only be known through testimony - the fact about the number being prime can be inferred from the randomization fact and the facts of the calculations that were carried out.

A community of experimenters who each independently verify each other's work can be said to have some sort of collective private knowledge - each knows on her own, not by testimony from the others (though the testimony of the others was important in getting her to do the experiment on her own), and each knows in a way that could only be shared by testimony. This is the situation of a very thorough community of scientists that verify every experiment independently.

Mathematics as currently practiced is not like that - the community of mathematics actually has *shared* knowledge, not just through testimony, but through transferred proofs. The reasons are public, rather than every individual having her own private reasons for the same belief. But the fact that there is a kind of collective knowledge to be had by a community suggests a variety of ways that an epistemic community can be organized.

- Tell the reader P is true. (newspaper)
- Tell the reader the theoretical framework in which P was shown. (textbook)
- Tell the reader the structure of the experiment that was performed, so she can repeat it herself. (traditional scientific journal)
- Do this plus include all data that were generated along the way. (newer scientific journals)
- Give an argument that an expert can use to convince herself, without having the repeat the work. (mathematical journal)
- Give an argument that a non-expert can use to convince herself. (computer verifiable proof)

Why settle for one rather than another?

> Scientists' basis for trusting each other lies in their knowledge of each other's commitment to, and embedding within, the norms and institutions of their profession. Unreliability is likely to

> be subsequently discovered and highly penalized in such a setting, and this gives one strong empirical reason, amongst others, to expect informants to be trustworthy. Scientists perforce rely on each other's reports of their findings. But they need and surely do not do so blindly. (Fricker, "Trusting Others in the Sciences", p. 383)

I claim that the norm of mathematics introduces some important differences - we really can be blind to the social role of the author and accept papers by outsiders. It doesn't matter if you're an undergrad or a Fields medalist - if the proof you publish is transferable then experts will accept it once they read it and work through it. If it is not, then they will reject it. Social standing may be required in order to get your foot in the door at a journal, but once your paper is in, the social standing of the author has no further epistemic importance. This is quite different from the situation in the less stringent norms - one trusts results from labs one knows to be very scrupulous far more than one trusts results from labs that are known to be questionable. If this value of social equality is of epistemic importance to the community, then there is an epistemic importance to the requirement of transferability.

Of course, the existence of scientific communities and journalistic communities with different standards for publication shows that this norm is not essential for a well-functioning epistemic community. Fallis is right to say that transferability is not epistemically necessary. But we can consider the epistemic costs and benefits to organizing a community in different ways.

The scientific community has been going through such debates lately as well. Although traditionally, experiments are expected to be repeatable, there are clearly cases in which no one ever can or will repeat them. Studies of the wild behavior of passenger pigeons or dodos can never be repeated, and neither will studies resulting from the Large Hadron Collider. Limitations of space have also made it permissible for authors to publish their conclusion with only brief summaries of the data. However, many journals now insist on publication of the full raw data in a web archive (with suitable anonymization of medical data and the like). (However, a recent study by Ioannidis and colleagues has shown that a majority of papers published in those journals don't actually follow the requirements and even those that do don't provide as much data as one might like. "Public Availability of Published Research Data in High-Impact Journas", Alawi A. Alsheikh-Ali, Waqas Qureshi, Mouaz H. Al-Mallah, John P. A. Ioannidis, PLoS ONE, September 7, 2011)

And in mathematics there have been calls to change the standards as well. The mathematician Doron Zeilberger, at Rutgers, has made calls for an "experimental mathematics" that seems to agree with Fallis' view that probabilistic proofs should be considered just as good for the community at generating knowledge. Since far more results can be tested far more quickly in this way, he argues that expanding the list of known truths will outweigh the epistemic virtues of sensitivity to the evidence.

Georges Gonthier and others working on computer-aided proofs have argued for a change in the opposite direction. Rather than relying on fallible human experts to verify proofs, they argue that *all* mathematical proofs should be published in a way that a computer could check as thoroughly as the proofs of the four color theorem and the like have been checked. A human could verify any part of the proof due to the triviality of the inference rules involved, but the proofs are far too long for individuals to survey. *Principia Mathematica* was an early attempt at something like this, though of course Russell and Whitehead didn't actually give a proper formalization of very much (they didn't even explain their notation fully!), and they only reconstructed a few of the major results of mathematics. If Gonthier gets his way, then we may some day have a library of all the proofs of all mathematical results, in fully formal detail.

I claim that there are epistemic values for the community that are preserved by keeping a requirement of transferability to human experts, even though there are other epistemic values that would be better served by stricter or looser standards.