Transferability and Proofs

Kenny Easwaran

Draft of October 15, 2007

1 Grice on Meaning

[Grice, 1957] argues for the following account of "non-natural meaning" (i.e., ordinary linguistic meaning):

> "A meant$_{\text{NN}}$ something by x" is roughly equivalent to "A uttered x with the intention of inducing a belief by means of the recognition of this intention." [Grice, 1957, p. 384]

In particular, he adds, "A's intending that the recognition should play this part implies ... that he does not regard it as a foregone conclusion that the belief will be induced in the audience whether or not the intention behind the utterance is recognized."

The motivation for this particular condition is a series of examples in which A utters x with the intention of inducing a belief, and with the intention that this intention be recognized, but where (because the recognition of the intention plays no essential role in the formation of the belief) we don't want to say that A "means" anything by x. These examples include Herod presenting Salome with the head of John the Baptist (to get her to believe that John the Baptist is dead); a child letting its mother see how pale it is (to get her to believe that the child is ill); and a contrast between the case of presenting Mr. X with a photograph of another man being unduly familiar with his wife, and drawing a picture of the same situation.

However, counterexamples to this requirement have been proposed.

> "If Grice's account of what it is for someone to mean something were correct, an unwelcome and somewhat ironic consequence would be that although Grice will have written and published an article of several pages on what it is for someone to mean something, Grice will have meant almost nothing by what he wrote." [Schiffer, 1972, p. 42]

I will suggest that there is another series of examples in which precisely this condition fails to be met, and yet we do want to say there is meaning (in the relevant sense). My main concern will be to argue that in these cases, the lack of this condition is essential to the nature of the communication.

2 Mathematics Journals

The specific type of communication I have in mind is that of a paper in a refereed mathematics journal. I take it that it's clear that such a paper is indeed meaningful, and that it does in fact often result in a transmission of beliefs from the author to the reader. However, I will argue that a recognition of the intention of the author is not necessary for this transfer. That is, contra Grice, it is"a foregone conclusion that the belief will be induced in the audience whether or not the intention behind the utterance is recognized." I think that being non-Gricean in this sense is generally required, and I will suggest that this requirement of non-Gricean-ness will answer a question raised by Don Fallis.

Various sorts of proofs are accepted for publication in mathematical journals. Notably, these proofs are not complete formal proofs of the sort studied in proof theory, but are rather some sort of informal approximation to them. As pointed out in [Fallis, 2003], in addition to being expressed in informal language (rather than formal symbolism), many steps are only gestured at, or even left out completely when they seem obvious enough to relevant specialists. In addition, at least since the 1976 publication of Haken and Appel's proof of the four-color theorem, it has been considered permissible to publish proofs, some of whose steps can only be carried out by computer calculation and won't fit in the published version. The question asked by [Fallis, 1997], which is my primary focus, is whether there is any epistemic reason for mathematics journals not to accept other, "probabilistic", arguments as well. (I will explain the details of probabilistic proofs later.)

I claim that the requirement that a paper be non-Gricean may in fact serve an epistemic purpose, and that it draws the line at just the point Fallis says that mathematicians actually draw the line.

A first relevant question, that should be more fully answered before this claim can be dealt with completely, is just what the proper role of a mathematics journal is. Does it aim to share mathematical discoveries, truths, or knowledge? Or rather than being restricted to one of these categories, should it be seen as a general means of communication to facilitate further discoveries? Or perhaps the role is more purely social, to enforce standards of methodology and productivity among working mathematicians?

Surely, all of these goals are at work at various points in different journals. But the goal I am most interested in is that of sharing information, whether discovery, knowledge, or truth. To the extent that discovery is relevant, conjectures and very programmatic statements may well qualify. If it is truth that must be shared, then some of these will be cut out, although standards for justification won't necessarily be as high as if knowledge is required. Thus I will primarily focus on the goal of transferring knowledge from author to reader. In fact, to make things easier, I will focus on the transfer of justified belief in a (presumably true, in whatever sense mathematical claims are true) proposition. There may well be Gettier cases in mathematics, but I think they will be largely irrelevant to the particular issues at work here.

This is clearly not a complete characterization of the relevant goals of a

journal, though it will suffice for current purposes. But note that the goals of journals in any discipline presumably have some focus on the transmission of justified belief from authors to readers. Thus, considerations that bear on the acceptability of probabilistic proofs in mathematics journals may well bear on the acceptability of various sorts of arguments in other journals.

3 Fallis

Fallis suggests in [Fallis, 1997] that there is no epistemic purpose for which probabilistic proofs are less acceptable than other mathematical methods.

First of all, he is concerned only with proof "as a means of establishing mathematical truths." [Fallis, 1997, p. 166] He leaves open the possibility of other goals for proof, such as providing good *explanations* for conclusions, which very well may favor deductive proofs over probabilistic proofs. But in response to a passage of Wittgenstein advocating deductive proofs for just this reason, he says, "while providing understanding is nice, it is not required."[Fallis, 1997, p. 170] After all, many deductive proofs provide very little understanding or explanation, but they are still published if they are the first proofs of some interesting result.

Fallis argues that for every property conducive to the goal of establishing justified belief, either there are methods acceptable to mathematicians that lack the property, or probabilistic proofs of some sort have them. (He is concerned in this article with "probabilistic DNA proof" used to prove that particular graphs have no paths of a certain sort; in [Fallis, 2000] he considers a computerized method of deciding whether a large number is prime. I will focus more on the latter) However, at the December 2005 meeting of the Association for Symbolic Logic, Michael Rabin suggested that probabilistic proofs are "non-transferable". I will suggest that there is a precise property (which, following Rabin, I will call "transferability") that traditionally acceptable proofs have but probabilistic proofs lack. This property is intimately connected to the fact that a published mathematical paper generally communicates without need of the Gricean condition mentioned above.

The first property Fallis considers is that of providing absolute certainty. As I think he is right to point out, deductive proofs don't in general provide absolute certainty either - when a proof is exceedingly long and complicated, one is in general not absolutely certain that the conclusions in fact follow from the premises. Even when it has passed the referee process at a journal, it is not certain - journals regularly withdraw former publications that turn out to contain serious flaws. Many acceptable proofs also contain steps carried out by computer - in these cases, we can't be absolutely certain that the computer performed as specified unless we are certain the program contained no bugs, and are certain that the physical microchips behave as specified.

He also suggests that the particular degree of certainty can't be relevant - after all, probabilistic methods can grant extremely high certainty, while ac-

ceptable proofs as large as the classification theorem for finite simple groups[1] almost certainly contain some invalid steps.

The next property he considers is that of providing "conditional certainty". Although there is a chance of invalidities in an argument, or failure with computer software, it seems that the conclusion is absolutely certain, *conditional* on the claim that "nothing went wrong". Probabilistic proofs don't provide this guarantee. However, as Fallis points out, proof "sketches" are often considered acceptable - very few published proofs actually cover all the relevant steps, instead relying on the reader to fill some in based on her familiarity with the material. In particular, Fallis discusses the first publication of Gödel's second incompleteness theorem - Gödel relies essentially on the claim that the proof of the first incompleteness theorem can be carried out in a formal system in order to prove the second, though he only ever actually gives an informal proof of the first. No full proof of the second incompleteness theorem was published until decades later. ([Fallis, 2003] shows how standard examples of this sort really are.) However one makes precise the notion of conditional certainty, a proof sketch can't do it. Therefore, since proof sketches are acceptable, Fallis suggests that this property can't be essential.

The final properties Fallis considers are those of giving *a priori* warrant, or a "proof", that the claim that the conclusion is true. However, as Fallis points out, since long deductive proofs require checking, and can be discovered to be invalid, they are not *a priori* under many analyses of this notion. For many calculations, they also rely on the fact that computers (and paper, and blackboards) reliably preserve their data during the calculation, since the entire calculation can't be internalized at once. And since incomplete proof sketches are acceptable, any proper criterion of acceptability must include them. But they provide no better evidence than probabilistic proofs that a deductive, *a priori* proof can be found. The only real difference seems to be a difference in the phenomenology of this evidence, which Fallis points out can't be epistemically relevant. However, I do think that the status of mathematics as something like the *a priori* underlies the fact that mathematicians are able to require transferable, and thus non-Gricean proofs.

Thus, he has surveyed a range of potential properties that could separate probabilistic proofs from acceptable means of establishing mathematical conclusions. Assuming that long deductive proofs, computer calculations, and incomplete proof sketches are all acceptable, Fallis seems to have argued that probabilistic proofs should be accepted as well.

[1] A group is a certain type of abstract algebraic structure, and ones with a particular property are known as "simple groups". In around 1980, it was established that all finite simple groups fell into one of fifteen well-defined infinite classes, except for 26 particular "sporadic" groups. The proof proceeded by a huge enumeration of cases, and was carried out by dozens of mathematicians in hundreds of published papers and books, totaling around 10,000 pages. It has all been refereed, but no one person has been able to follow all of it.

3.1 Probabilistic proofs

The technique Fallis primarily concerns himself with in [Fallis, 1997] is "probabilistic DNA proof". This is a technique for showing that a particular directed graph has no Hamiltonian path.[2] For each vertex in the graph, the mathematician chooses a pair of distinct sequences of DNA bases. Each directed edge in the graph is then represented by a strand of DNA that links up in the appropriate ways to the strands representing its start and end vertices. The mathematician is then able to make many copies of the strands representing the edges and the vertices, and stir them in a test tube to ensure that many long strands are created. With only a little work, she can then select out strands of exactly the right length, and then use physical facts about the particular strands representing each vertex to select out all and only the strands containing that sequence. If any strands remain at the end, then there is a Hamiltonian path. And if there is a Hamiltonian path, then (provided enough copies of the strands were created in the initial stages) it is extremely likely that some strands will remain at the end.

This procedure is not deterministic, but it yields a very high likelihood $P(strands|HP)$ (which can be made arbitrarily close to 1 by making enough copies of each DNA strand at the start) and a likelihood of 0 for $P(strands|\neg HP)$. Thus, it can be used to convince oneself whether a particular graph has a Hamiltonian path. It is clear that $P(HP|strands) = 1$. And by Bayes' Theorem, we can calculate:

$$P(HP|\neg strands) = P(\neg strands|HP)\frac{P(HP)}{P(\neg strands)}$$

We want to make this value as small as possible, so that the lack of strands at the end gives one high posterior confidence that there is no Hamiltonian path. Since $\neg HP$ entails $\neg strands$, we see that $\frac{P(HP)}{P(\neg strands)} \leq \frac{P(HP)}{P(\neg HP)}$, which is one's initial betting odds for the existence of a Hamiltonian path. Since this value is fixed (for a particular graph), we can make one's posterior confidence in the non-existence of such a path as high as we want, just by making $P(strands|HP)$ sufficiently high (that is, by making enough copies of all the strands at the beginning of the process). Thus, we can use this technique as a probabilistic method for proving the non-existence of a Hamiltonian path. (It also generally finds a path if there is one, but for demonstrating a path, we can use completely non-probabilistic methods, even if the path was originally found in this way - non-existence is in general much harder to prove.)

[2] In this technical sense, a graph is a set of points, called "vertices", together with a specification of which pairs of vertices count as adjacent. Such vertices are said to have an "edge" between them. A path is a sequence of vertices, each of which is adjacent to the next one. In a directed graph, the edges may have a direction to them, specifying which direction they are allowed to be traversed in paths - they can go in one direction or the other, or in both directions. A Hamiltonian path is one that contains each vertex of the graph exactly once. Graphs and directed graphs are often used to model computer networks, social networks, highway systems, and many other things.

The technique Fallis discusses in [Fallis, 2000], to which I will pay greater attention, is the Miller-Rabin primality test. Miller and Rabin established that if a number n is composite, then there are at least $3n/4$ integers less than n satisfying a certain relation to n. However, if n is prime, then there are none. Therefore, to test whether a number is prime, one chooses a long sequence of numbers less than n, and checks whether each satisfies this specific relation, which is relatively straightforward. If none of them do, the the number is declared prime; if one does, then it is declared composite.

In this case, particular probability values are easy to come by. If we check k integers less than n, and we choose these numbers to check by some means independent of the process by which we chose n, then it seems clear that we should have $P(yes|prime) = 1$ and $P(yes|\neg prime) \leq 1/4^k$. By a use of Bayes' Theorem similar to the previous case, we can see that $P(prime|no) = 0$ and $P(\neg prime|yes) \leq P(yes|\neg prime)\frac{P(\neg prime)}{P(prime)}$. Thus, if our threshold for belief is $1 - \epsilon$, then to convince ourselves that a number is prime, we just need to make sure that $P(yes|\neg prime) \leq \epsilon\frac{P(prime)}{P(\neg prime)}$. Since $P(yes|\neg prime)$ goes down exponentially based on the number of trials, we see that we just need this number of trials to be proportionate to the logarithm of the prior betting odds against primality. Thus, no matter how unlikely the mathematician originally thought it was that the number was prime, she can use this test to fairly quickly convince herself that it is, or to find a witness to its compositeness otherwise.

An important note for each of these methods is that the relevant conditional probabilities can be arrived at by the mathematician without any substantive constraints on her subjective probability function beyond the probability axioms. She just needs to trust that the DNA strands mix in some suitably thorough way, and that she has some method for generating integers in some range that she regards as equally likely to produce any given number. She doesn't need an actually random process to do this, as long as she has no reason to believe that this process is biased for or against non-witnesses for the particular number in question. No debates about "objective Bayesianism" or "logical probability" are necessary.

There are some worries about applying the notion of probability in mathematics - in particular, traditional formulations of probability theory require that a rational agent assign degrees of belief to the theorems of a system that are at least as great as the degree of belief assigned to the conjunction of the premises. Since the fact of a number's being prime or not is always a consequence of the Peano axioms, and we can assume that most mathematicians are at least very highly confident of the Peano axioms, this would prevent them from being uncertain of the primality of these numbers. Since this is clearly false (rational mathematicians are in fact uncertain of whether or not various large numbers are prime), there is a challenge to the use of probability to measure uncertainty in mathematics.

However, there are ways around this problem. [Garber, 1983] recommends replacing all statements of interest in a particular application of probability by propositional atoms, and then applying probability theory to the resulting

propositional language. In this way, the actual logical relations between statements are opaque to the system, so the requirement of "logical omniscience" doesn't interfere with the possibility of uncertainty. This is an incomplete solution, because there is still logical omniscience in the resulting propositional language, so this picture can't be the right final story about uncertainty in mathematics, but it's a good start that lets us at least represent the relevant situation appropriately. Another approach is suggested by [Gaifman, 2004], where he advocates assigning probabilities only to a subset of the formulas of the language, in accord with a slightly modified set of probability axioms. Any result that can't be proved using only statements of this restricted language will then not be required to have maximal probability. He discusses the example of probabilistic primality proofs extensively in this paper.

Whether or not this particular approach solves this problem, there are reasons to adopt something like a probabilistic framework for thinking of mathematical beliefs. Probability (or something like it) is widely seen as the right framework for discussing partial degrees of belief, and it's clear that some notion of partial belief is required to adequately model mathematical epistemology, or else we will ignore the role of conjecture, hypothesis, and failed proof.

David Corfield suggests one more reason in chapter 5 of [Corfield, 2003]. "To contemplate the reliability of a result in a particular field we should think of someone from outside the field asking a specialist for their advice. If the trustworthy expert says she is very certain that the result may be relied upon, does it matter to the enquirer how the specialist's confidence arises?" [Corfield, 2003, p. 110] For the outsider, confidence in the result will be based entirely on the specialist's confidence. He can't worry how she got the result, because he isn't qualified to decide between methods. So at least for the outsider, some single probabilistic scale of certainty seems to be the right notion of partial belief, just as it is for other non-mathematical areas.

4 Transferability

However, if it's another mathematician asking the insider, then Corfield's argument makes a different suggestion - an insider may be very reliable at recognizing solid proofs, and even outlines of proofs, but very bad at making conjectures, or vice versa. A friend working in model theory once said that the great model theorist Boris Zilber is "everywhere locally wrong but globally right", because he has made a series of conjectures that have each turned out to be false, but have motivated exactly the right sort of thinking to prove interesting, related statements. Whether or not this is the right way to characterize him, it seems plausible that some mathematicians may have this sort of track record, so one wouldn't want to just adopt their credences as one's own in a field.

While an outsider may not be able to judge a mathematician's reliability, an insider may. Thus, this reliability will be an important consideration for a mathematician that wants to get involved in a field. A mathematician reading a paper may want to base her credences on evidence without relying on testi-

mony. If I'm right, this suggests that appeals to authority are to be avoided - a mathematician reading a journal doesn't want her justification to consist just in the fact that she has read something in a prestigious journal, but wants to be confronted directly with evidence of a non-testimonial sort that will raise her credences.

The position that such a mathematician will find herself in is a strange one - she wants to gain true beliefs about mathematics, but wants to do so in a way that doesn't depend on the reliability of other people. This position is obviously untenable in one's ordinary life. If I didn't believe street signs that said "road work ahead", or friends that told me they would meet me for dinner, my life would be very difficult indeed. Surprisingly, in mathematics (unlike most areas of life) this may actually be a tenable position. If someone presents a sequence of propositions for my consideration, and each proposition is such that consideration of it in light of my current beliefs leads me to believe it, then I can learn quite a bit from testimony, even if I don't trust it. For instance, if someone presents a deductive proof of some conclusion, I don't have to believe anything they say, as long as I independently have a high credence in the premises, and see independently that each step follows from previous ones.[3]

Of course, a mathematician can't maintain this position for her entire mathematical life, because she will rely on testimony for relevant results outside her area of work. For instance, a real analyst who is told that a certain tangential claim is equivalent to a large cardinal axiom in set theory will stop working to prove it - she has been told that these axioms are provably independent of ZFC, but doesn't need to work through this whole proof herself. Similarly, a topologist might reduce some claim to an algebraic one, and then just appeal to outside sources to convince herself that this algebraic claim is true. However, directly in the core parts of her own research, she will want to convince herself of everything and avoid trusting testimony.

If this is an important goal for mathematicians, to be justified in their conclusions through non-testimonial means, then it may be relevant for journals to require the relevant sort of exposition. Papers will rely only on premises that the reader can be assumed to antecedently believe, and only make inferences that the reader would be expected to accept on her own consideration. Arguments of this form I will call "transferable", following Rabin's terminology. On receiving a transferable proof, the reader will (if she has the right mathematical expertise) come to believe the conclusion of the argument. She will in fact recognize the intention of the author to get her to believe this conclusion, but this intention is inessential in her coming to believe. Thus, transferable proofs are non-Gricean.

Note that this standard doesn't necessarily require complete deductive proofs - in many cases, mathematicians can be relied upon to be familiar with certain modes of argumentation, so that a presentation of a sequence of propositions in a proof outline can in many cases be sufficient for the reader to convince

[3]Philosophers will be familiar with the strategy of granting premises believed by one's opponent, in order to use them in an argument for one's own claim.

herself of the result (whether by mentally filling in the missing steps, or just being familiar enough with the domain to see that the claim follows). When other papers are cited, if they have all been published at this sort of standard, then a mathematician can in principle go through and convince herself in each case of the relevant result, without relying on testimony. If she decides that certain results are far enough afield that she doesn't care to check, then she can rely on testimony, but by establishing a custom of only publishing in a way that allows mathematicians to convince themselves, these citations will preserve transferability. If the reader is already familiar enough with the cited literature, then the citations will trigger her tacit beliefs in the relevant theorems - she won't rely on the intention of the author at all in coming to form her beliefs.

Computer proofs are slightly harder to transfer - however, I believe that they are in principle transferable as well. If the author provides the code for the relevant computer program, then the reader can presumably see just as well as the author that it does what is claimed. Actually running the program can then be done in the same way. I suggest that it is important that exactly the same calculations can be carried out with exactly the same pattern of reasoning, in both the complete proof case and the deterministic computer proof case. This way, we can see that it really is the same proof that is being transferred from the author to the reader. But at any rate, the reader can use the communication to come to believe the result without having to refer to the intention of the author at any step in the process, so the Gricean condition still fails.

5 Non-transferability of Probabilistic Proofs

Returning to the probabilistic proofs described earlier, I suggest that in each case, though the test satisfies Fallis' criteria, it doesn't satisfy this criterion of transferability. In the case of probabilistic DNA proof, the published paper can say that the test was run, and say what the results were, but the reader has no independent way of convincing herself that this is actually true. She must rely on the testimony of the author.

In the case of the primality test, there is a bit more that can be done, because the author can publish the sequence of integers less than n that are checked, and the details of the calculations showing for each that it is not a witness of non-primality. However, unless the reader believes that these numbers were selected in a manner independent of the primality of n, she has no reason to be convinced. Miller and Rabin's initial proof only shows that at most $n/4$ such integers fail to be witnesses - so a sequence of 100 non-witnesses can often be found, even for composite n. The author can be convinced, because she selects the values to check "at random" (that is, in some manner independent of her selection of n as the number to consider).[4] But the reader just has to trust that

[4] [Fallis, 2000] spends a long time discussing worries about the impossibility of using a random number generator that produces integers in the relevant range with equal chance. However, the relevant notion of probability here isn't chance, but rather uncertainty. Sampling the least-significant bits of the system clock at the moment the author decides to run the

the author hasn't cherry-picked the sequence of k's to fit the n, or cherry-picked the n to fit the sequence of k's. Thus, she can't convince herself without relying in some sense on the testimony of the author. The proof is non-transferable, and the communication requires the Gricean condition.

One might seek to make the proof transferable by fixing some canonical list of "random" values to check for each problem. If such a list is "random enough" it can serve the same purpose as randomly selected numbers, and because it is canonical, the reader can be sure that the author hasn't manipulated things. One way to do so might be to use something like the successive strings of appropriate length chosen from the decimal expansion of π. Of course, if the particular number whose primality is being checked is connected to π in some way, we can no longer be sure of the randomness here, so we may have to do something more convoluted, like using some particular published table of the distances between pairs of stars in our galaxy, in alphabetical order.

However, this won't always work for all purposes, because once the sequence of potential witnesses to check is fixed, one might be able to cherry-pick the number n to be checked for primality. The claim that some particular number n is prime is almost never publishable (exceptions generally occur only for values larger than any currently known prime). However, one may be able to prove (say) that a certain equation has no solutions, provided that a particular number is prime. If there is only one relevant number, then the canonical list of witnesses to check may work, provided that the list is long enough (and the reader's prior betting odds on the primality of this number are no worse than the author's). But if the author just needs to prove that *at least one member* of some set is prime, and she doesn't say how she picked the one whose primality to demonstrate, then she could search for one in this set that might be "easier" to claim to be prime, given the particular canonical sequence.

Thus, there seems to be no way to use a canonical list of "random" witnesses to check. Instead, mathematicians will have to generate this list themselves each time, just ensuring that the process used to generate potential witnesses is independent of the process used to generate the number whose primality is being

---

algorithm may not have a well-defined chance of producing every number in the relevant range (especially if we assume that the author's decision-making process is deterministic), but it still seems rational for the author to believe to an equal degree the proposition that this process will produce any number as opposed to any other. Thus, when updating her subjective certainties in light of the outcome of the process, she will be using equal probabilities, regardless of what the underlying chances actually are.

If it turns out that there is in fact a systematic bias in this process, so that it leads to false diagnoses of primality relatively often, then although the author has been misled, I think she has still been perfectly rational. Only once she has reason to believe that something like this is likely would she seem irrational for assigning equal degree of belief to each number being picked by this process.

Fallis claims, "unpredictability is not the important issue for our purposes" - but I think this is not quite correct. What is important is that the mathematician rationally have an extremely low degree of belief that a series of non-witnesses will be picked given that the number is composite, and certainty that such a series will be picked given that the number is prime. Then, updating by Bayes' theorem guarantees that if a series of non-witnesses is picked, her degree of belief in primality will greatly increase, and if at least one number is a witness, then she will become certain of compositeness.

tested. Because this list must be generated each time, the reader must trust the author's testimony as to the source (and independence) of the list, and must share the author's prior credence of primality, in order for the proof to transfer. To ensure independence, this source will most likely have to be something like the separation in nanoseconds between consecutive keystrokes on the author's computer, rather than some pseudo-random number generating algorithm that can be directly transferred to the reader. The author can describe the process used to generate the numbers, but the reader will not be able to verify that this process does in fact produce these particular numbers.

Another attempt to make these proofs transferable is to compare both DNA and Miller-Rabin proofs to deterministic computer proofs. The author can publish the method used, and allow the reader to repeat the computation herself. Presumably the reader will assign the same conditional credences as the author in the case of DNA calculation, and can also find some source she trusts to generate integers uniformly at random in some range for the Miller-Rabin test. However, in both cases, the actual series of steps generated by the reader will be different from the ones generated by the author. In deductive proofs, proof sketches, and deterministic computer proofs by contrast, the sequence of steps is the same. Thus, although there is some way to transfer something, it is not the same proof that gets transferred.

6 Is Transferability Desirable?

Thus, I claim that Fallis' conclusion was not quite right - there is one epistemic property relevant to establishing the truth of mathematical claims that traditional methods have but probabilistic methods lack. However, there still remains a question as to whether this property should actually be a consideration in the acceptability of mathematical proofs.

Transferability of this sort is clearly not a criterion for publication in scientific journals in general. As long as the conclusions depend at least in part on the results of some experiment, the reader must rely on the author's (and perhaps referee's) testimony that the experiment worked as reported, and also that the author really performed the experiment exactly as claimed. This practice is standard in the physical sciences, and could easily be adopted by mathematicians as well, as suggested at the end of [Fallis, 2000]. It is important to note though that scientists always take great care to include in their papers a description of the methodology of the experiment, so that the reader can attempt to replicate them. While experimental results aren't transferable in my sense, they are certainly very useful things to publish for the reason that they are replicable. It seems that probabilistic proofs are just as good in this sense.

This transfer requires a bit more effort on the part of the reader than standard proofs, but in many cases the reader must already put in a lot of effort to understand the reasoning that makes each step in a proof or proof sketch plausible. In standard proofs and proof sketches, the reader's credence boost can be seen as independent of actually believing the author - the author's role

is just to cause the reader to entertain the analogy, which either convinces her or shows her how to convince herself. In the case of probabilistic proof however, the reader *must* go through the work to convince herself. She can't be presumed to have any credence that the sequence of numbers displayed is independent of the number to be checked, given that both come from the same author.

Although the situation of probabilistic proof is analogous to the situation in the physical sciences, it is plausible that the situation of currently accepted mathematical proofs is analogous to the situation in the humanities, and especially philosophy. If the content of a paper is just some claims, together with some arguments establishing these claims, then the reader doesn't need to rely on the intentions of the author in forming her beliefs.[5] (In fact, some of the best philosophical papers will convince their readers of a conclusion even if the reader thinks that the author's intention is generally a reason *not* to believe the conclusion!) Thus, as Schiffer pointed out, Grice's own paper was a counterexample to his thesis. The goal of having transferability of this sort may explain why many philosophers are more willing to accept arguments based on the author's own intuitions than "experimental philosophy" based on actual data about intuitions - when the author mentions her intuitions, the reader can check that she has the same intuitions, and doesn't need to just trust the author as she does with the results of surveys.

In mathematics a further question arises when considering extremely complicated proofs that make use of specialist knowledge in multiple areas of mathematics. If these proofs are very broad-ranging, then they may not in practice be fully transferable to anyone. This issue seems to have become more relevant with the increasing commonality of multi-author papers. In many cases, no single author understands the complete proof - so it would be surprising if any single reader could. Whether such proofs can really count as transferable, and whether this sort of transferability is important, may be important questions.

To work out the details of probabilistic proof, we will need a more well-developed account of Bayesianism in mathematics (or perhaps some other sort of notion of partial belief, which will presumably interact appropriately with the probabilistic methods mentioned here). There is a difference between probabilistic proofs and traditional mathematical proofs, and it may well be a distinction that mathematicians care about, but other scientists have gotten over it. Thus, although I disagree with Fallis about the particular point about the non-existence of a property to distinguish probabilistic proofs from traditional ones, I agree with him about the broader picture. Mathematicians will surely need to embark at some point on this debate about whether non-transferable proofs that are easily replicated should be allowed. The Miller-Rabin test is more easily replicated (with greater choice in the source of randomness) than probabilistic DNA proof. The question is just whether mathematicians (and philosophers) should be willing to act more like scientists in this way, rather

[5]Of course, in philosophical papers, a greater emphasis is placed on describing historical dialectics and asserting the importance of one's claims than in mathematics. This part of the discussion may well require trusting the author, and therefore relying on her intentions in forming one's beliefs. But there is still a large part of the paper that doesn't.

than humanists.

References

[Corfield, 2003] Corfield, D. (2003). *Towards a Philosophy of Real Mathematics*. Cambridge University Press.

[Fallis, 1997] Fallis, D. (1997). The epistemic status of probabilistic proof. *The Journal of Philosophy*, 44(4):165–186.

[Fallis, 2000] Fallis, D. (2000). The reliability of randomized algorithms. *British Journal for the Philosophy of Science*, pages 255–271.

[Fallis, 2003] Fallis, D. (2003). Intentional gaps in mathematical proofs. *Synthese*, 134:45–69.

[Gaifman, 2004] Gaifman, H. (2004). Reasoning with limited resources and assigning probabilities to arithmetical statements. *Synthese*, 140:97–119.

[Garber, 1983] Garber, D. (1983). Old evidence and logical omniscience in Bayesian confirmation theory. In Earman, J., editor, *Testing Scientific Theories*, volume 10. Minnesota Studies in the Philosophy of Science.

[Grice, 1957] Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66(3):377–388.

[Schiffer, 1972] Schiffer, S. (1972). *Meaning*. Oxford.