Mathematics and transferable knowledge

By: Lo Min Choong Julian | 22 April 2024

Written for the partial completion of NTU's HY3018 Epistemology

---

In this essay, I outline an account of knowledge which seems most plausible to me, i.e. knowledge is transferable true doxastic states, see precisely why this account fails (ordinary propositions are not evidence in and of itself), and show how it compares to traditional accounts, thus demonstrating why it is worth consideration.

In *Probabilistic Proofs and Transferability* (2008), Kenny Easwaran notes the debate within the mathematical-philosophical community regarding the acceptability of certain kinds of proofs. Some argue incomplete and computerised proofs are acceptable, while probabilistic proofs are not. Others argue that probabilistic proofs are acceptable. Don Fallis, a proponent of the acceptability of probabilistic proofs argues that "given the proofs that mathematicians do accept, there is no epistemic reason for them to reject 'probabilistic proofs'" (p. 342). This is what Easwaran responded to. Ideally, mathematicians would only submit and publish full formal proofs in symbolic form so that every step is present in the paper. However, in actual publications, incomplete proofs where "many steps are only gestured at, or even left out completely, perhaps with a note that 'It is easy to see that…'" show that mathematicians choose not to provide complete formal proofs because they believe a reader with sufficient knowledge and expertise will be able to fill in the blanks (p. 342). Thus, mathematicians believe it is reasonable to omit steps whenever such steps are not central to the proof at hand, precisely because readers are expected to have pre-existing knowledge, or minimally be familiar with the omitted proof/lemma/steps, thus be able to "fill in the blanks". For example, if a paper uses the triangle inequality, the author may choose not to list the steps which involve the inequality merely because the author expects the readers to be familiar with it and its applications, precisely because it is basic knowledge (for mathematicians). That is to say, if the steps are trivial or obvious such that any reader will be able to write the steps out themselves without the author's help, then the author has reasonable grounds to omit said steps from the proof. On the other hand, if someone claimed to have proven the Riemann hypothesis, then it is expected that all essential steps should be listed in full, as such a proof is not obvious or trivial, nor can it be left to the reader as an exercise.

Regarding computerised proofs, Easwaran highlights the Four Colour Theorem proof, which contains steps which can "only be carried out by computer calculation and will not fit into the published version" (p. 342). This may sound like a problem for transferability, but we will see why computerised proofs are acceptable. For context, the Four Colour Theorem proof was controversial during its submission because it was the first computerised proof to be accepted for publication. In the interest of space, such steps can be omitted from the journal publication, although it is expected that the full proof be made available elsewhere. The controversial nature stems from the inability and infeasibility for humans to actually do or check what computers did. In principle, a human could do some of the calculations that a computer can do, e.g. checking to exhaustion every possible four-coloured map. But given the intensity of such a task is too demanding—it leaves no room for error. Checking every permutation (a billion cases; 1,834 reducible maps) by hand is infeasible (Thomas, 2017). This is where the point of contention lies. Given that something is infeasible to check by hand, should we accept computer-assisted proofs? The problem lies in trust, or verification. Thomas says "an argument can be made that our `proof' is not a proof in the traditional sense, because it contains steps that can never be verified by humans." (Thomas, 2017). He is talking about the computation itself, specifically the unknowable accuracy of the hardware and software which ultimately produced the proof. We don't know (can't be certain) whether the compiler was error-free. You could run the program again, and see if the results agree (match), but that says nothing about the accuracy of the program itself. The Four Colour Theorem proof was controversial (i.e. not accepted by a majority of mathematicians initially) precisely because of its inability to be verified. Tymoczko dubbed the Four Colour Theorem proof a "non-surveyable proof", arguing that since we are unable to (feasibly) verify it by hand, we have grounds to reject it (Tymoczko, 1979). Therein lies the problem with knowledge and probabilistic proofs: verification. The Four Colour Theorem only gained some support after "Robertson, Sanders, Seymour, and Thomas published a more streamlined proof of the theorem [in 1995]" (Gonthier, 2005). The theorem essentially attained its current status after a 2005 paper, in which Gonthier says, "we have written a formal proof script that covers both the mathematical and computational parts of the proof. We have run this script through the Coq proof checking system [13,9], which mechanically verified its correctness in all respects. Hence, even though the correctness of our proof still depends on the correct operation of several computer hardware and software components (the processor, its operating system, the Coq proof checker, and the Ocaml compiler that compiled it), none of these components are specific to the proof of the Four Colour Theorem." (Gonthier, 2005). Thus, Gonthier demonstrates that a computer proof can be verified in a manner which allows all doubts about the proof itself to be adequately addressed. The Four Colour Theorem is significant in this discussion precisely because it was the first major proof of its kind. This explains the mathematical community's resistance towards it. They did not want room for doubt or error. Hence, they demanded a proof that is acceptable by the community. The question is: what exactly makes a proof acceptable? What criteria distinguishes proofs from non-proofs?

In *Probabilistic Proofs and Transferability* (2008), Kenny Easwaran proposes that the criteria of acceptable mathematical proofs is transferability. He says "the basic idea is that a proof must be such that a relevant expert will become convinced of the truth of the conclusion of the proof just by consideration of each of the steps in the proof" (p. 343). For example, Easwaran notes that the proof of Fermat's Last Theorem is transferable in principle. However, in practice, mathematicians rely on the testimony of others as the proof is incredibly long and inaccessible to the majority of the population (both the general public and mathematicians themselves). Thus, the acceptability of the proof of Fermat's Last Theorem rests on its transferability. Consequently, Easwaran says probabilistic proofs are unacceptable because it is unable to divorce itself from the testimony of the author. For example, in a test of primality, "Miller’s 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 non-prime n. The author can be convinced, because she selects the values to check ‘at random’ [...] But the reader just has to trust that the author has not cherry-picked the [data]. Thus, she cannot convince herself without relying in some sense on the testimony of the author." (p. 357). In short, there exists proofs in mathematics which are reliant on the testimony of the authors, and thus are non-transferable proofs. This non-transferability is what makes such proofs undesirable, hence unacceptable. Even if the author publishes the list of numbers used as witnesses to a candidate's primality, one can never be certain that the list of numbers were indeed generated randomly and fairly. One could repeat the proof for themselves by generating their own, independent list of witnesses and seeing for themself if the candidate passes the primality test. This does generate knowledge via verification, but it does not change the non-transferability of the submitted proof. Such attempts allow one to "peer review" knowledge. The problem is, if such proofs are taken to be "public knowledge" or "canon", it ultimately rests on the author's testimony. The status of probabilistic proofs, and whether it genuinely is non-transferable, is up for further discussion and consideration by the community (which is outside the scope of this essay).

Recall my earlier mention of Fermat's Last Theorem: the proof is about a hundred pages long; it requires specialised expertise to follow the proof. Thus, while the proof is transferable in principle, most people still attain knowledge via testimony. Easwaran says this is permissible, as the proof is feasibly transferrable. Consider a stronger example—In *Testimony and Autonomy in Mathematics* (2011), Easwaran makes a distinction between an individual epistemic agent and the fictional collective when discussing who actually possesses knowledge. He notes that the proof of the classification of finite simple groups is transferable in principle, but most still learn it via testimony. What makes this proof unique in a sense is its length: it took a hundred authors and tens of thousands of journal pages to lay out the proof. It can be argued that no one person has followed the proof in its entirety, given its length and technical inaccessibility. Thus, who actually possesses knowledge of the proof? If you take strict transferability, then no one does, since no one has understood every premise, lemma and step independently (i.e. by themselves, as a rational epistemic agent). Then, technically speaking, several epistemic agents have partial knowledge of the proof (because people wrote and read it). Under strict transferability, could you argue that the mathematical community possesses knowledge of the proof? What would that even mean? Can the fictional collective have knowledge? Is the fictional collective an epistemic entity? These are open (perhaps unimportant) questions, which will remain outside the scope of this essay. These questions serve as primers to the role of epistemic agents within epistemology.

In his discussion about the role of transferability in mathematics, Easwaran says "The author can actually give the evidence to the reader" (Easwaran, 2011, p. 4). He means formal proofs are pieces of evidence that are literally transferable to the readers, so that one may independently and rationally realise the proof's soundness. Thus, Easwaran operates under the individualised epistemic agent train, rather than the fictional collective camp. This is inline with traditional accounts of knowledge such as justified true belief and gettier counters. Thus, while the reader may begin knowing through testimony, it is ultimately transferable knowledge that underpins the justification of one's belief. Easwaran believes that this is specific and unique to mathematics, as he says in *Probabilistic Proofs* that this position is "obviously untenable in one’s ordinary life. If I did not 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" (Easwaran, 2008, p. 353). He also notes that transferability is a "social epistemic virtue" (p. 343). Easwaran believes statements such as "road work ahead" is evidence that is meaningfully distinct from mathematical propositions. He says "road work ahead" is merely a proposition which is true, where no evidence is presented to the reader such that one may independently arrive at the conclusion that there is indeed road work ahead. After all, if one were to disregard all testimonial evidence, one will only gain knowledge when one sees for themselves the road works as they pass it by. Thus, ordinary life will never have transferable proofs as statements will never be transferable the same way mathematical propositions are. It really does seem that mathematics is unique in this sense.

Easwaran's analysis of mathematical knowledge is appreciated precisely because it is thorough in covering several aspects of knowledge. He considered the role of the epistemic agent in his account. In *Testimony and Autonomy*, he says "One is an expert on anything that one knows" (p. 2), showing that one gains some form of power or ability via knowing. In *Dr. Truthlove, or How I Learned to Stop Worrying and Love Bayesian Probabilities* (2014), Easwaran proposes an account of beliefs under a Bayesian system, analysing whether it is tenable or not. Easwaran's Bayesian system is beyond the scope of this essay. Instead, I would like to highlight certain key concepts used. He says "the appropriate way to think of her doxastic state is in terms of attitudes that come in degrees from 0 to 1. These “credences” are the things that matter." (p. 2). He also says "I will call the objects of belief “propositions”, and I will assume that these propositions are characterized by sets of possible “situations”. [...] [T]hese situations represent the uncertainties for the agent." (p. 4).

Thus, while the truth values of propositions are discrete, our beliefs are not, as our beliefs are doxastic states which are represented by a sort of credences. As Easwaran notes, the Bayesian system and the axioms of probability are merely tools in analysing the nature of operation of such a belief system, thus one should not be committed in saying our beliefs are bound by the axioms of probability, &c (Easwaran, 2014, p. 19). This meets the anti-luck epistemologists half way, spotlighting the role of absolute certainty within epistemology, as it is technically attainable under this proposed system, but not set as the goal, simply being one of possible states. Furthermore, it more accurately corresponds with reality, in which uncertainty is present in many situations, including ordinary settings.

My proposed account of knowledge, transferable true doxastic states, parallels justified true belief because it makes each component more precise (except truth). This precision is useful because the defeaters of this account points us in a direction which advances epistemological exploration. Transferability is a stronger form of justification, as both probabilistic and nonprobabilistic proofs are forms of justification, except one is acceptable, thus "better" than the other. I note that it is possible to commit ourselves to social virtue epistemology, given Easwaran's earlier comment, but that is noncentral. The problem with transferability is its inability to be the sort of justification that we want to have in ordinary settings. Take the original gettier case: there is a sheep in the field, and a wolf that looks indistinguishably similar to a sheep (Pritchard, 2016, p. 6). The proof in this case is the presence of the sheep or the wolf, and in the misidentification of the animal. The proof itself is transferable in a sense, just not in the mathematical sense. Person A may misidentify the wolf and come to the incidentally true belief that there is a sheep in the field. Person B, independent of person A, may form the same belief, or note that there is a wolf and the sheep in the field. The point is, in the absence of any testimony, an epistemic agent has the opportunity and ability to arrive at the same erroneous conclusion, or the correct (i.e. non-misidentification) conclusion. This is a result of conditions, i.e. logic and the availability of evidence, not of the nature of the evidence itself (which Easwaran pinpoints in his discussion of transferable mathematical propositions). Thus, transferability in ordinary life is not mathematical, it's practical. Practical transferability refers to a scenario where epistemic agents are able to arrive at the truth independently. In a sense, it is like being in a courtroom setting, where evidence and arguments are laid out to be judged. For example, Chuck, the chicken-sexer, is able to determine the sex of a chicken via his reliable process, whatever that may be (Prichard, 2016, pp. 14, 47). The case that Chuck has knowledge can be made by saying the knowledge in question is transferable in principle. For example, we could determine the sex of a chicken via an ultrasound machine, or an X-ray scan, or something else that will give one information of a chicken's sex. The counter to this is that the means of generating knowledge is materially different. We are unable to verify or understand Chuck's perspective because his epistemic generative process is inaccessible to us. We will never know what it is like to be Chuck because we will never have his chicken-sexing intuitive abilities, thus we will never know for certain via verification whether Chuck has knowledge. The rejoinder to this argument is we have arrived at the same true belief, regardless of how we know. Thus, knowledge is transferable in ordinary life. This argument works even if there is only one person who know that p, or if p is obstructively inaccessible to every epistemic agent, given our discussion of Fermat's Last Theorem and the classification of finite simple groups (because p is transferable in principle, therefore p is knowledge). Trust and verification dictates what is acceptable knowledge. Transferability aims to eliminate the former and produce the later.

Regarding doxastic states, consider the gettier counter once more: perhaps you realise that you are uncertain about what type of animal is in the field. Then you can say your doxastic state is some non-discrete credence. This is an imprecise description of beliefs, thus unsatisfactory, but it hints at something deeper. The discussion about beliefs, doxastic states and credences attempts to address one of the five problems of knowledge outlined by Williams: problem 5, the problem of value (Williams, 2001, p. 2). What beliefs are is not central in assigned epistemology readings because it is not currently a priority for the field and the classroom, as we have yet to settle the other problems of knowledge which currently take precedence. Problem 3, the problem of method, has been the highlight of this essay within mathematical epistemology, which took the backseat in our discussion of ordinary knowledge, precisely because it is controversial to say knowledge is transferable true doxastic states. We have yet to nail down the problem of method because knowledge itself resisted our attempts of description. Perhaps if we accepted uncertainty and attempted to incorporate it into our accounts of knowledge, we may discover more about knowledge itself. Thus, I have demonstrated what transferability is, and the significance of epistemic agents arriving at conclusions independently. While transferability is an unsatisfying account of knowledge, it points us to consider our standards and criteria of knowledge.

References

Easwaran, K. (2008, November 6). Probabilistic Proofs and Transferability. Philosophia Mathematica, 17(3), pp. 341–362. https://doi.org/10.1093/philmat/nkn032.

Easwaran, K. (2011). Testimony and Autonomy in Mathematics. Unpublished draft.

Easwaran, K. (2014, September 1 & 2016, December). Dr. Truthlove, or How I Learned to Stop Worrying and Love Bayesian Probabilities. Noûs, 50(4), pp. 816-853. https://www.jstor.org/stable/26631430.

Gonthier, G (2005). A computer-checked proof of the Four Colour Theorem. https://api.semanticscholar.org/CorpusID:18529792.

Pritchard, D. (2016). *Epistemology*. Palgrave Macmillan. ISBN: 9781137526908. https://philpapers.org/rec/PRIE-25.

Thomas, R (2017, July 19). The Four Color Theorem. *Website*. https://thomas.math.gatech.edu/FC/fourcolor.html.

Tymoczko, T (1979). The Four-color Problem and Its Philosophical Significance. *The Journal of Philosophy* 76(2), pp. 57-83. https://doi.org/10.2307/2025976.

Williams, M. (2001). *Problems of knowledge: a critical introduction to epistemology*. Oxford University Press. ISBN: 0192892568. https://philosophy.jhu.edu/faculty-books/problems-of-knowledge-a-critical-introduction-to-epistemology/.