

Probabilistic Proofs, Lottery Propositions, and Mathematical Knowledge

Yacin Hamami*

To appear in the *Philosophical Quarterly*

Abstract

In mathematics, any form of probabilistic proof obtained through the application of a probabilistic method is not considered as a legitimate way of gaining mathematical knowledge. In a series of papers, Don Fallis has defended the thesis that there are no epistemic reasons justifying mathematicians' rejection of probabilistic proofs. The present paper identifies such an epistemic reason. More specifically, it is argued here that if one adopts a conception of mathematical knowledge in which an epistemic subject can know a mathematical proposition based solely on a probabilistic proof, one is then forced to admit that such an epistemic subject can know several lottery propositions based solely on probabilistic evidence. Insofar as knowledge of lottery propositions on the basis of probabilistic evidence alone is denied by the vast majority of epistemologists, it is concluded that this constitutes an epistemic reason for rejecting probabilistic proofs as a means of acquiring mathematical knowledge.

*Centre for Logic and Philosophy of Science, Vrije Universiteit Brussel, Brussels, Belgium. Email: yacin.hamami@gmail.com. Website: www.yacinhamami.com.

In mathematics, the only accepted means to acquire knowledge of a mathematical proposition other than an axiom is through a *deductive proof* of it. In particular, any form of *probabilistic proof* obtained through the application of a probabilistic method is *not* considered as a legitimate way of gaining mathematical knowledge. In a series of papers,¹ Don Fallis has defended the thesis that “mathematicians do not have good grounds for their rejection of probabilistic methods” (Fallis, 1997, p. 165), where a ‘good’ ground is, for Fallis, an *epistemic* one.² In other words, Fallis holds that there are no epistemic reasons to deny knowledge of a mathematical proposition based solely on a probabilistic proof of it. In this paper, I will argue that if one adopts a conception of mathematical knowledge in which an epistemic subject S can know a mathematical proposition based solely on a probabilistic proof, one is then forced to admit that S can know several *lottery propositions*³ based solely on probabilistic evidence.⁴ Insofar as knowledge of lottery propositions on the basis of probabilistic evidence alone is denied by the vast majority of epistemologists for reasons to be recalled below,⁵ I conclude that this constitutes an epistemic ground for rejecting probabilistic proofs as a means of acquiring mathematical knowledge.⁶

Before engaging in this discussion, it is important to get clear on what probabilistic proofs are. In the context of Fallis’ thesis, one is said to possess a *probabilistic proof* for a mathematical proposition Φ whenever: (1) one has run a probabilistic algorithm designed to decide whether Φ , (2) the algorithm has yielded ‘YES’ as a response, (3) there is a non-zero probability that the algorithm is mistaken when it yields ‘YES’ as a response, and (4) one possesses a bound on the probability that such a mistake occurs.⁷ A *probabilistic* (also called *randomized*) *algorithm* is an algorithm which appeals to one or several *random choices*—i.e., *lotteries*—during its execution.⁸ Because a probabilistic proof results from the run of an algorithm that has a non-

¹See Fallis (1997, 2000, 2002, 2011).

²Fallis recognizes that mathematicians might have reasons of various kinds to reject probabilistic proofs (see Fallis, 1997, p. 166). His thesis only concerns the *epistemic status* of probabilistic proofs, as he puts it: “I am only claiming that [mathematicians] do not have good *epistemic* reasons” (Fallis, 1997, p. 166).

³A lottery proposition is a proposition expressing the outcome of a lottery, the paradigmatic example being “ticket t is a loser” while talking about one of the tickets of a given lottery. The term ‘lottery proposition’ is originally due to Vogel (1990).

⁴I will develop the argument in the paradigmatic case discussed in this literature, namely the probabilistic proofs produced by the *Miller-Rabin primality test*, and will indicate how the argument generalizes.

⁵The epistemic status of lottery propositions has received substantial attention in contemporary epistemology in discussions relative to the *lottery paradox* in its ‘knowledge’ version—to be distinguished from the lottery paradox due to Kyburg (1961) which does not concern knowledge but rational acceptance. The ‘knowledge’ version of the lottery paradox is originally due to Harman (1968, 1973), and has since then generated an extensive literature. See, among others, Cohen (1988), Vogel (1990), DeRose (1996), Lewis (1996), Nelkin (2000), Williamson (2000), Hawthorne (2004), Douven (2007), Kvanvig (2009), and Smith (2010, 2016).

⁶A different epistemic reason for rejecting probabilistic proofs has been advanced by Easwaran (2009) who identifies a property that he calls ‘transferability’ and that, according to him, deductive proofs possess and probabilistic proofs lack. This proposal has been critically discussed by Jackson (2009) and Fallis (2011). As Easwaran’s proposal is orthogonal to the one to be developed here, I will not discuss it in this paper. The interested reader is invited to consult the references just mentioned.

⁷I agree with Jackson (2009) that the use of the term ‘probabilistic proof’ might be confusing here, insofar as the term has been used in mathematics to refer to deductive proofs relying on the *probabilistic method*, a method described by Alon and Spencer (2015) as a way to “prove the existence of a combinatorial structure with certain properties” by constructing “an appropriate probability space and show that a randomly chosen element in this space has the desired properties with positive probability” (Alon and Spencer, 2015, p. xiii). Jackson suggests to switch to the term ‘randomized argument’, but I choose to stick to the term ‘probabilistic proof’ in order to be consistent with the terminology previously used in this philosophical discussion.

⁸Motwani and Raghavan (1995) define a probabilistic algorithm as “an algorithm that is allowed access to a source of independent, unbiased, random bits; it is then permitted to use these random bits to influence its computation” (Motwani and Raghavan, 1995, p. 6). Notice that a probabilistic algorithm does not necessarily yield an answer that runs a risk to be incorrect, that is, some probabilistic algorithms do give the correct answer all the time. In this case, the random choices or bits can lead to variations in the behavior of the algorithm, in particular regarding its running time performances.

zero probability of being mistaken, a probabilistic proof for a mathematical proposition Φ only provides *probabilistic evidence* for Φ .

The example that has received most attention in previous philosophical discussions of Fallis' thesis is the one of the probabilistic proofs obtained through the probabilistic algorithm developed by Michael O. Rabin (1980) for deciding whether a number n is prime.⁹ The algorithm exploits a certain condition originally identified by Miller (1976) for a number b to be a *witness to the compositeness of n* , whose interest lies in the fact that one can easily show that if there exists a witness to the compositeness of n , then n is composite, and, by contraposition, that if n is prime, then there are no witnesses to the compositeness of n .¹⁰ Rabin's algorithm is, furthermore, based on a fundamental theorem proved by Rabin (1980, p. 130) which states that if a number $n > 4$ is composite, then more than $3/4$ of the numbers $1 \leq b < n$ are witnesses to the compositeness of n . The algorithm works as follows: first, it chooses randomly and independently k numbers $1 \leq b_1, \dots, b_k < n$;¹¹ second, it evaluates whether each b_i is a witness to the compositeness of n ; finally, if at least one of the b_i is a witness to the compositeness of n , the algorithm outputs that n is composite, otherwise the algorithm outputs that n is prime. The algorithm's output is always correct when it says that n is composite. The algorithm can, however, be mistaken when it says that n is prime, the reason being that the random choices of the k numbers $1 \leq b_1, \dots, b_k < n$ might happen to pick only nonwitnesses of the compositeness of n when n is composite. Thanks to the fundamental theorem proved by Rabin, we know that the probability of picking randomly and independently k nonwitnesses to the compositeness of n when n is composite is smaller than $1/4^k$. Thus, we know that the probability that the algorithm is mistaken when it says that n is prime is smaller than $1/4^k$, and so that the probability that the algorithm is correct when saying that n is prime is greater than $(4^k - 1)/4^k$. Insofar as Rabin's algorithm can be mistaken when saying that a given number is prime, a probabilistic proof for the proposition that " n is prime" obtained through this algorithm only provides probabilistic evidence for it.

I will now argue, in the paradigmatic case of the probabilistic proofs produced by Rabin's probabilistic algorithm, that if one accepts that an epistemic subject S can know a mathematical proposition based solely on a probabilistic proof, one is then forced to admit that S can know several lottery propositions based solely on probabilistic evidence. To this end, assume that S can know the mathematical proposition " n is prime" based solely on a probabilistic proof produced by Rabin's algorithm, that is, based on a given run of the algorithm with input n and output " n is prime". Assume, furthermore, that S has indeed run Rabin's algorithm with input n , obtained as output that " n is prime", and that on this basis:

S knows that n is prime,

where the probabilistic evidence S has for the proposition " n is prime" is greater than $(4^k - 1)/4^k$. Now, we can also assume that S knows the mathematical proposition "if n is prime, then there are no witnesses to the compositeness of n ", since S can easily acquire knowledge of this proposition by simply consulting either Miller (1976) or Rabin (1980). Since S knows that n is prime, and S knows that if n is prime, then there are no witnesses to the compositeness of n , S can come to know by some straightforward deductions all the mathematical propositions of the form " b is not a witness to the compositeness of n " for any number b such that $1 \leq b < n$. If we assume that S has competently carried out those deductions, we have that, for any number b such that $1 \leq b < n$:

S knows that b is not a witness to the compositeness of n .

⁹This method is now known as the *Miller-Rabin primality test*.

¹⁰The algorithmic interest of this property comes from the fact that checking whether a number b is a witness to the compositeness of a number n can be done at a low computational cost (see Rabin, 1980).

¹¹Allowing possible repetitions, that is, some of the picked numbers might end up being equal.

Now, the mathematical propositions of the form “ b is not a witness to the compositeness of n ” are lottery propositions for the particular lottery consisting in drawing a number between 1 and $n - 1$ and saying that one wins whenever the drawn number is a witness to the compositeness of n . Furthermore, S only has probabilistic evidence for those lottery propositions, since S has deduced them from the proposition “ n is prime” for which S only had probabilistic evidence. This means that S knows those lottery propositions solely based on probabilistic evidence. Thus, the previous argument establishes, in the particular case of the probabilistic proofs obtained through Rabin’s probabilistic algorithm, that if S can know a mathematical proposition based solely on a probabilistic proof, S can know several lottery propositions based solely on probabilistic evidence. To put it differently, the argument shows that if one accepts that S can know the mathematical proposition “ n is prime” based on a probabilistic proof obtained through Rabin’s probabilistic algorithm, one has as a consequence that, based only on the knowledge of a single draw of k numbers $1 \leq b_1, \dots, b_k < n$ where none of them are witnesses to the compositeness of n , one can know that any other number drawn in this interval will not be a witness to the compositeness of n . This argument can be applied to any situation where the output of the probabilistic algorithm entails a proposition stating a property of the outcome of one or more of the lotteries present in the considered probabilistic algorithm.¹²

To see more distinctively why lottery propositions occur in the context of the probabilistic proofs produced by Rabin’s algorithm, it is useful to consider the following epistemic situation analogous to the previous one. Imagine an urn with $N - 1$ balls for which S knows that either all the balls in the urn are white, or more than $3/4$ of the balls in the urn are black. Suppose that k balls are randomly and independently drawn from the urn.¹³ If one of the balls is black, S can deduce that more than $3/4$ of the balls in the urn are black, and in this case there are no issues in saying that S *knows* that more than $3/4$ of the balls are black. If none of the balls are black, the chances that all the balls in the urn are white are very high, since the chances to pick k times in a row a white ball in the urn in the second situation where more than $3/4$ of the balls in the urn are black is very low (as we said earlier, the probability that this happens is less than $1/4^k$).¹⁴ Now, suppose that, in this situation, we would say that:

S knows that all the balls in the urn are white.

By a straightforward deduction, S can then reach an epistemic state in which for any ball b in the urn:

S knows that b is not black.

As a matter of fact, S has strong probabilistic evidence for each of these propositions, the probabilistic evidence for each proposition being greater than $(4^k - 1)/4^k$. Yet, despite such strong probabilistic evidence, if another ball b is drawn randomly from the urn, would you be willing to say prior to the drawing that S *knows* that b won’t be black?

If you have answered negatively to the previous question, you might have followed a general inclination to deny knowledge of lottery propositions based solely on probabilistic evidence. Most contemporary epistemologists reject the possibility that an epistemic subject can know a lottery proposition based solely on probabilistic evidence¹⁵—paradigmatically that an epistemic subject can know that a given ticket is a loser based solely on the probabilistic evidence

¹²Needless to say, this property should be such that it cannot be determined prior to the drawing.

¹³Allowing possible repetitions, that is, each drawn ball is put back into the urn after it has been picked.

¹⁴If k is very small, say equal to 1 or 2, one might not say that a probability of $1/4^k$ is ‘very low’. Since in the concrete applications of Rabin’s algorithm k is relatively large, we shall assume here that k is sufficiently large to meaningfully say that $1/4^k$ is ‘very low’.

¹⁵Douven (2007) qualifies this possibility as “downright absurd” (Douven, 2007, p. 327). All the authors listed in footnote 5 reject knowledge of lottery propositions as well. To my knowledge, only two epistemologists have embraced the possibility of knowing lottery propositions: Morillo (1984) and Reed (2010).

that can be computed from the considered lottery setting—and this has been taken as a datum for most contemporary philosophical theorizing about knowledge. This general inclination is particularly manifest when one considers the connection between knowledge and assertion, as well as between knowledge and practical reasoning. For, if S knew that her ticket in a given lottery was a loser, S would have no reservations about asserting flat-out that her ticket is a loser, and yet there seems to be strong disinclination to flat-out assert a lottery proposition such as “ticket t is a loser” at any time prior to the drawing of the lottery, or at least prior to obtaining by other means information on which ticket has been drawn. Furthermore, if S knew that her ticket was a loser, it would be hard to make sense of why S bought the ticket in the first place (assuming S bought the ticket), and why S would intend to keep the ticket until the drawing, instead of (say) throwing it away or giving it to someone else. As mentioned earlier, the epistemic situations described in the two previous paragraphs can be construed as lottery situations in which one wins whenever the drawn number is a witness to the compositeness of n , or whenever the drawn ball is black. Thus, whatever reason there is to deny that an epistemic subject S can know a lottery proposition based on probabilistic evidence also counts as a reason to deny that S can know that b is not a witness to the compositeness of n for any number b such that $1 \leq b < n$, or that S can know that b is not black for any ball in the urn, in the epistemic situations previously described.

It could be objected that, in these situations, it might happen that there is no winning ticket. This, however, is not an issue since, as DeRose (1996) has shown, our disinclination to grant knowledge of lottery propositions still holds in lottery situations where there is no winning ticket:

[W]ith many lotteries, there is no winning ticket. Many of the big state lotteries, for example, *usually* have no winner. Still, it seems, you don’t know you’ve lost. In case you think that is because the jackpot is carried over to the next month’s drawing, so we think of the whole process as one giant lottery which will eventually have a winner, note that our ignorance of losing seems to survive the absence of that feature. Suppose a billionaire holds a one-time lottery, and you are one of the 1 million people who have received a numbered ticket. A number has been drawn at random from among 100 million numbers. If the number drawn matches that on one of the 1 million tickets, the lucky holder of that ticket wins a fabulous fortune; otherwise, nobody receives any money. The chances that you’ve won are 1 in 100 million; the chances that somebody or other has won are 1 in 100. In all likelihood, then, there is no winner. You certainly don’t believe there’s an actual winner. Do you know you are a loser? Can you flat-out assert you are a loser? No, it still seems. Here, the mere chance of being a winner—with nothing remotely like an assurance that there actually is a winner—does seem to destroy knowledge of your being a loser. (DeRose, 1996, p. 571)

To my knowledge, this point has been widely accepted in the subsequent epistemological literature on lottery propositions (see, e.g., Nelkin, 2000, p. 389; Williamson, 2000, p. 248; Hawthorne, 2004, p. 8). Still, it may be objected further that, in the lottery situations I have considered, it is *very* unlikely that there will be a winning ticket. But it would be very strange if the capacity of an epistemic subject to know that her ticket is a loser depends on the actual probability that there will be a winning ticket. For consider again the case of big state lotteries which may not have a winner. The probability that there will be a winning ticket in a draw depends on the number of tickets that have been sold. Now, the probability that your ticket is a loser does not depend at all on the number of sold tickets, and so the evidence you possess in favor of believing that your ticket is a loser does not depend on the probability that there will be a winning ticket. It would seem very odd to imagine, for instance, that you would suddenly know that your ticket is a loser because you have learnt that only a small amount of tickets

have been sold.

Easwaran (2009) has rightly pointed out that probabilistic methods such as the Miller-Rabin primality test differ from standard lottery situations in that they *track* the truth of the considered propositions. He then observed that: “Although obeying these probabilistic tracking conditions may or may not be either necessary or sufficient for knowledge (see Roush, 2005), they certainly make the situation epistemically better than in a lottery case” (Easwaran, 2009, p. 348). I agree with Easwaran that this feature makes the situation epistemically better for probabilistic methods as compared to standard lottery situations. The key question is whether this is *sufficient* for such probabilistic methods to yield *knowledge* of the mathematical propositions they purport to establish.¹⁶ I will now argue that it is not.¹⁷

First of all, it should be noted that the lottery propositions I identify in the above argument are *not* the mathematical propositions established by Rabin’s probabilistic algorithm—which are propositions of the form “ n is prime”—but other mathematical propositions that can be *deduced* from them, namely mathematical propositions of the form “ b is not a witness to the compositeness of n ” for any number b such that $1 \leq b < n$. It should also be noted that in the epistemic situation I consider in which the epistemic subject S has run Rabin’s algorithm with input n and obtained as output that “ n is prime”, S does *not* track these latter propositions: if the number b being drawn turns out to be a witness to the compositeness of n (in the very unlikely case in which n is composite, and all the numbers randomly picked by the algorithm turned out to be non-witnesses to the compositeness of n), S would still believe that “ b is not a witness to the compositeness of n ” prior to the draw. The tracking view of knowledge would then classify these propositions on a par with standard lottery propositions, and together with propositions admitting very unlikely exceptional cases that the epistemic subject cannot track such as “the ice cubes have melted” in the famous example proposed by Vogel (1987).¹⁸ In a discussion of this family of epistemic situations, Roush (2005) wrote the following:

Knowledge I have about there being a large number of eligible tickets and that the drawing is fair give me knowledge that my ticket will very probably not win, but it is generally agreed that I do not know that my ticket will not win. Either tracking view has a neat explanation of this fact since if my ticket were going to win I might, and probably would, still believe that it was not, because I am isolated from any indication that it will win even if it will. It is similar with the ice cubes not melting

¹⁶As witnessed by the sentence just quoted, Easwaran (2009) does not take a stand on this issue. In a defense of the thesis that non-deductive methods—including probabilistic methods—can yield knowledge of mathematical propositions, Paseau (2014, p. 788) uses this distinguishing feature to dismiss any connection between probabilistic methods and lottery situations. The argument presented in this paper shows that such a connection cannot be so easily dismissed.

¹⁷In this discussion, I will take as a representative of the tracking view of knowledge the *recursive tracking view* developed by Roush (2005) because this version of the tracking view is formulated in terms of conditional probabilities which makes it easy to apply it to the Miller-Rabin primality test, but also because Roush (2005) provides a detailed discussion of lottery cases and how they are handled by her view. Roush’s view is based on Nozick’s tracking account of knowledge (Nozick, 1981). In addition to the traditional requirements of truth and belief, Nozick’s account proposes two conditions for knowing a proposition p : the ‘variation’ condition which says that if p were not true, then the subject would not believe that p ; the ‘adherence’ condition which says that if p were true, then the subject would believe that p . Roush (2005) implements two key revisions to Nozick’s account. The first one is to replace the use of subjunctive conditionals with conditional probabilities: the variation condition now says that $P(\neg b(p) \mid \neg p) > t$, and the adherence condition now says that $P(b(p) \mid p) > t$, where $b(p)$ means “subject S believes p ” and t is a suitable high threshold. The second one is to add a closure condition on knowledge under known implication which is implemented through a recursive clause (Roush, 2005, p. 47). For a detailed description of the recursive tracking view see Roush (2005, chapter 2). For another updated version of the tracking view that relies on dispositions instead of subjunctive conditionals for the variation and adherence conditions, see Briggs and Nolan (2012).

¹⁸In this example, one has left ice cubes outside for a few hours on a very hot day, preferring to go back inside the house to avoid the heat. The question is whether one knows in this case that “the ice cubes have melted” although one has not gone outside to check whether the ice cubes have melted.

and the other cases here. The subject is isolated from any indication of things not going in the expected way, and so does not know that they definitely have not, only that they very probably have not. (Roush, 2005, p. 67)

In the epistemic situation I consider, the epistemic subject S has no indications that Rabin's algorithm has not gone in the expected way, i.e., that she is not in one of those exceptional and very unlikely cases in which the algorithm would have picked only non-witnesses to the compositeness of n although n was indeed composite. In this family of epistemic situations, Roush (2005, pp. 65–67) considers that one does not know that “the ice cubes have melted”, “ticket t is a loser”, and (I would add) “ b is not a witness to the compositeness of n ”, although one does know that “the ice cubes have *very probably* melted”, “ticket t is *very probably* a loser”, and (I would add) “ b is *very probably* not a witness to the compositeness of n ”.

Now, because Rabin's algorithm does track the truth of propositions of the form “ n is prime”, the tracking view will still attribute knowledge of a mathematical proposition of the form “ n is prime” based on a probabilistic proof produced by Rabin's probabilistic algorithm. And through the recursion clause (Roush, 2005, p. 47)—i.e., through some straightforward deductions—knowledge of the mathematical propositions “ b is not a witness to the compositeness of n ” for any number b such that $1 \leq b < n$. What are we to do with that? My assessment is that the (recursive) tracking view is getting things wrong here. The problem is that, based solely on a run of Rabin's algorithm with input n , the epistemic subject can come to know through deduction that she is not in one of those exceptional cases in which Rabin's algorithm would have picked only non-witnesses to the compositeness of n although n was composite. This seems absurd, and a similar conclusion has been rejected by Roush in the ice cubes case and similar ones:

Notice that if we attributed knowledge that the ice cubes melted, and not merely that they probably melted, to the person in the house, then by closure our subject would, if he or she were sufficiently reflective, thereby have knowledge that this instance was not one of those exceptional cases. This is because this follows deductively from the generalization and the instance, both of which our subject knows. But that this was not an exceptional case is precisely what our subject does not know because of failure to track the instance. (Roush, 2005, p. 66)

So what is going wrong here? I believe that the answer is to be found in the deductions by which the epistemic subject can come to know all the mathematical propositions of the form “ b is not a witness to the compositeness of n ” for any number b such that $1 \leq b < n$ when she knows that “ n is prime”. These deductions exploit the mathematical theorem that “if n is prime, then there are no witnesses to the compositeness of n ”, and this theorem is precisely what allows the epistemic subject to conclude that she is not in one of those exceptional cases in which Rabin's algorithm would have gone astray. Interestingly, such deductions are not possible in the ice cubes case since what one knows and tracks there is the empirical generalization that “ice cubes left in high temperatures generally melt”, and there it is the “generally” that blocks the possibility to deduce from this generalization that “the ice cubes have melted”, although one can still deduce that “the ice cubes have *very probably* melted”. As Roush notices, the trick in these empirical generalizations is that “knowledge that exceptional events have probably not occurred can be had indirectly through knowledge of the generalizations to which those events would be exceptions” (Roush, 2005, p. 65). This trick is not available in the epistemic situation considered in the above argument, since when the epistemic subject knows that “ n is prime” she can come to know that “ b is not a witness to the compositeness of n ” for *all* numbers b such that $1 \leq b < n$ (without exception!). I will leave it to the defenders of the tracking view to figure out what is the best way to accommodate this example in their theory of knowledge. At any rate, I believe that the right conclusion in this case is *not* to attribute knowledge of the

propositions that “ n is prime” and that “ b is not a witness to the compositeness of n ” based solely on a probabilistic proof provided by a run of Rabin’s algorithm with input n , but rather knowledge of the propositions that “ n is *very probably* prime” and that “ b is *very probably* not a witness to the compositeness of n ”.

It may be objected that the evidence provided by a run of Rabin’s algorithm cannot be classified as merely probabilistic because it is produced by an *inductive method*. Since it is commonly held that inductive methods can produce knowledge—e.g., we may know that all ravens are black based on a finite number of observations of black ravens—one may wonder why Rabin’s algorithm, conceived as an inductive method, could not yield knowledge of the mathematical propositions at stake. There is, however, an important difference between Rabin’s algorithm and standard forms of induction which has to do with the subject’s *background knowledge*. More specifically, an epistemic subject using Rabin’s algorithm *knows*, prior to any run of the algorithm, that there *will* be cases in which the algorithm will mistakenly say that “ n is prime” because the algorithm would have only picked nonwitnesses to the compositeness of n . This is because the subject possesses a certain amount of mathematical knowledge before running the algorithm: she knows that (1) there are infinitely many composite numbers, and so that some of the natural numbers to be tested will be composite, and (2) for some composite number n , it will be the case that some of the numbers strictly smaller than n are not witnesses to the compositeness of n . By contrast, in a standard case of induction, the agent does not have any background knowledge on what may or may not happen in the course of her observations. Interestingly, these observations are in direct line with the arguments developed by Ryan (1996) as to why probabilistic evidence can yield knowledge in standard cases of induction, but not in lottery cases. A key observation of Ryan is that, in lottery cases, the epistemic subject also has *counterevidence* for the proposition “ticket t is a loser” for she *knows* that there will be a winning ticket¹⁹—for lotteries where it is assured that there will be a winning ticket—or that there may be a winning ticket—for lotteries where there may not be a winning ticket.²⁰ When evaluating the epistemic situation of a subject, one must then consider the *total evidence* available to the subject for a given proposition, that is, both her positive and *negative* evidence. This is why, according to Ryan, the overwhelming probabilistic evidence for the proposition “ticket t is a loser” is not sufficient for knowledge due to the presence of counterevidence, while it may be sufficient in cases where no counterevidence is present as in standard cases of induction. From this perspective, the probabilistic evidence produced by Rabin’s algorithm is closer to lottery cases than to standard cases of induction. This is because, in the case of Rabin’s algorithm, the subject always possesses counterevidence for the proposition “ n is prime”—the subject knows that there *will* be cases in which the algorithm will mistakenly say that “ n is prime” because the algorithm would have only picked nonwitnesses to the compositeness of n . This background knowledge then plays a role similar to the knowledge that there will be, or there may be, a winning ticket in a lottery situation. These considerations highlight then an important insight for the debate on the epistemic status of probabilistic proofs, namely that the *total* evidence possessed by the epistemic subject must be taken into consideration, that is, not only the overwhelming probabilistic evidence that may result from a probabilistic proof, but also the counterevidence that may be present in the subject’s background knowledge.

One potentially fruitful line of inquiry to better understand the epistemic status of probabilistic proofs is to see whether epistemological diagnostics as to what is blocking knowledge in

¹⁹In a similar vein, Pollock (1983, p. 237) has pointed out that, in lottery cases, the subject has “statistical grounds both for accepting and for rejecting the conclusion that any given ticket will lose”. It is such “conflicting considerations” that, according to Pollock, block the knowledge of lottery propositions. Pollock accepts that, in the absence of conflicting considerations, a subject can know a proposition on the basis of (high) statistical evidence.

²⁰Ryan (1996) focuses on lotteries where there is a guaranteed winner, but her point also holds for lotteries where there may not be a winning ticket since the fact that there may be a winning ticket constitutes as well counterevidence for the proposition “ticket t is a loser”.

lottery cases also apply to probabilistic proofs. As we have just discussed, the diagnostic proposed by Pollock (1983) and Ryan (1996) in terms of the presence of conflicting considerations or counterevidence also applies to the case of Rabin’s algorithm. Another relevant diagnostic is the one proposed by Smith (2010, 2016) in terms of *normic support*. Smith has argued that, in lottery cases, the epistemic subject lacks justification for believing lottery propositions because the subject’s evidence about the considered lottery does not normically support propositions of the form “ticket t is a loser” (Smith, 2010, p. 20). The idea is that, given the subject’s evidence, although it would be very unlikely that ticket t is a winner, it would not at all be *abnormal* if ticket t turns out to be a winning ticket, and this situation would not call for any particular *explanation*. Interestingly, this diagnostic also applies to the case of Rabin’s algorithm. More specifically, in a situation where an epistemic subject has come to believe the proposition “ n is prime” on the basis of a run of Rabin’s algorithm with input n and output “ n is prime”, it would not at all be abnormal, and it would not call for any particular explanation, if n turns out to be composite. The reason is that the subject *knows* that Rabin’s algorithm may turn out to pick only nonwitnesses to the compositeness of n when n is composite, in the same way that, in lottery cases, the subject *knows* that ticket t may turn out to be a winning ticket. The key idea of Smith is that normic support is a necessary condition for justification, and hence for knowledge. This idea can then be recruited to explain why one cannot know a lottery proposition on the sole basis of probabilistic evidence, and similarly why one cannot know a mathematical proposition on the sole basis of a probabilistic proof.

Finally, one may worry that the argument developed in this paper relies or promotes a general skepticism about knowledge on the basis of probabilistic or statistical evidence. Whether rejecting knowledge of lottery propositions leads to a form of skepticism, or at least forces us to renounce to large chunks of knowledge that we presumably have, is a well-known issue that has been extensively discussed in epistemology. However, most epistemologists commonly agree that this is not the case; the dominant view is that knowledge of lottery propositions should be rejected while most of the knowledge we commonly have should be preserved, including knowledge on the basis of high probabilistic evidence. The goal of the epistemological literature on the lottery paradox in its epistemic version is precisely to explain why an epistemic subject cannot know lottery propositions while being in a position to know various propositions that may be less likely to be true given the subject’s evidence. The argument developed in this paper exploits the rejection of knowledge of lottery propositions to argue that one cannot know a mathematical proposition based solely on a probabilistic proof. Insofar as rejecting knowledge of lottery propositions does not lead to skepticism, this argument does not rely or promote a general skepticism about knowledge on the basis of probabilistic or statistical evidence.

Fallis has defended the thesis that mathematicians do not have good epistemic grounds for rejecting probabilistic proofs as a means of acquiring mathematical knowledge. In this paper, I have argued that if one accepts that an epistemic subject S can know a mathematical proposition based solely on a probabilistic proof, one is then forced to admit that S can know several lottery propositions based solely on probabilistic evidence.²¹ There are, however, strong epistemic reasons to deny knowledge of lottery propositions based solely on probabilistic evidence. Taken together, this constitutes an epistemic ground for rejecting probabilistic proofs as a means of acquiring mathematical knowledge.²²

²¹The argument has been developed in the paradigmatic case of the probabilistic proofs produced by Rabin’s probabilistic algorithm. As we saw, the argument generalizes insofar as it can be applied to any probabilistic proof obtained through a probabilistic algorithm whose output entails a proposition stating a property of the outcome of one or more of the lotteries present in the algorithm.

²²It is important to notice that the argument provided here is primarily about *knowledge*. It does not prevent the possibility to *rationally accept* or *justifiably believe* a mathematical proposition based solely on a probabilistic proof.

References

- Noga Alon and Joel H. Spencer. *The Probabilistic Method (Fourth Edition)*. John Wiley & Sons, Inc., Hoboken, NJ, 2015.
- Rachael Briggs and Daniel Nolan. Mad, bad and dangerous to know. *Analysis*, 72(2):314–316, 2012.
- Stewart Cohen. How to be a fallibilist. *Philosophical Perspectives*, 2:91–123, 1988.
- Keith DeRose. Knowledge, assertion and lotteries. *Australasian Journal of Philosophy*, 74(4): 568–580, 1996.
- Igor Douven. A pragmatic dissolution of Harman’s paradox. *Philosophy and Phenomenological Research*, 74(2):326–345, 2007.
- Kenny Easwaran. Probabilistic proofs and transferability. *Philosophia Mathematica*, 17(3): 341–362, 2009.
- Don Fallis. The epistemic status of probabilistic proof. *The Journal of Philosophy*, 94(4): 165–186, 1997.
- Don Fallis. The reliability of randomized algorithms. *The British Journal for the Philosophy of Science*, 51(2):255–271, 2000.
- Don Fallis. What do mathematicians want? Probabilistic proofs and the epistemic goals of mathematicians. *Logique & Analyse*, 179–180:373–388, 2002.
- Don Fallis. Probabilistic proofs and the collective epistemic goals of mathematicians. In Hans Bernard Schmid, Marcel Weber, and Daniel Sirtes, editors, *Collective Epistemology*, pages 157–175. Ontos Verlag, Heusenstamm, 2011.
- Gilbert Harman. Knowledge, inference, and explanation. *American Philosophical Quarterly*, 5 (3):164–173, 1968.
- Gilbert Harman. *Thought*. Princeton University Press, Princeton, 1973.
- John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, Oxford, 2004.
- Jeffrey C. Jackson. Randomized arguments are transferable. *Philosophia Mathematica*, 17(3): 363–368, 2009.
- Jonathan Kvanvig. Assertion, knowledge, and lotteries. In Patrick Greenough and Duncan Pritchard, editors, *Williamson on Knowledge*. Oxford University Press, Oxford, 2009.
- Henry Ely Kyburg. *Probability and the Logic of Rational Belief*. Wesleyan University Press, Middletown, CT, 1961.
- David Lewis. Elusive knowledge. *Australasian Journal of Philosophy*, 74(4):549–567, 1996.
- Gary L. Miller. Riemann’s hypothesis and tests for primality. *Journal of Computer and System Sciences*, 13(3):300–317, 1976.
- Carolyn R. Morillo. Epistemic luck, naturalistic epistemology and the ecology of knowledge or what the frog should have told Dretske. *Philosophical Studies*, 46(1):109–129, 1984.
- Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, Cambridge, 1995.

- Dana K. Nelkin. The lottery paradox, knowledge, and rationality. *The Philosophical Review*, 109(3):373–409, 2000.
- Robert Nozick. *Philosophical Explanations*. Harvard University Press, Cambridge, MA, 1981.
- Alexander Paseau. Knowledge of mathematics without proof. *The British Journal for the Philosophy of Science*, 66(4):775–799, 2014.
- John L. Pollock. Epistemology and probability. *Synthese*, 55(2):231–252, 1983.
- Michael O. Rabin. Probabilistic algorithm for testing primality. *Journal of Number Theory*, 12(1):128–138, 1980.
- Baron Reed. A defense of stable invariantism. *Noûs*, 44(2):224–244, 2010.
- Sherrilyn Roush. *Tracking Truth: Knowledge, Evidence, and Science*. Oxford University Press, Oxford, 2005.
- Sharon Ryan. The epistemic virtues of consistency. *Synthese*, 109(2):121–141, 1996.
- Martin Smith. What else justification could be. *Noûs*, 44(1):10–31, 2010.
- Martin Smith. *Between Probability and Certainty: What Justifies Belief*. Oxford University Press, Oxford, 2016.
- Jonathan Vogel. Tracking, closure, and inductive knowledge. In Steven Luper-Foy, editor, *The Possibility of Knowledge: Nozick and His Critics*, pages 197–215. Rowman & Littlefield, Totowa, NJ, 1987.
- Jonathan Vogel. Are there counterexamples to the closure principle? In Michael D. Roth and Glenn Ross, editors, *Doubling: Contemporary Perspectives on Skepticism*, pages 13–27. Kluwer Academic Publishers, Dordrecht, 1990.
- Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, Oxford, 2000.