

Bayesian Convergence to the Truth and the Metaphysics of Possible Worlds

Simon M. Huttegger*†

In a recent paper, Belot argues that Bayesians are epistemologically flawed because they believe with probability 1 that they will learn the truth about observational propositions in the limit. While Belot's considerations suggest that this result should be interpreted with some care, the concerns he raises can largely be defused by putting convergence to the truth in the context of learning from an arbitrarily large but finite number of observations.

1. Introduction. In probability theory one often deals with the infinite, as in throwing a coin infinitely often. This raises interpretive challenges for the resulting set of elementary events or “possible worlds.” Elementary events like an infinite sequence of coin flips are logically possible, but in empirical investigations their metaphysical status has to be balanced with epistemological concerns. We do not observe infinitely many coin flips. Only finite sequences are observationally accessible for us. This is not to say that there is no place for infinite sequences in probabilistic reasoning. In situations with no principled upper bound on the number of observations, they serve as idealizations that approximate large finite sequences.

These considerations are important for answering some criticisms of Bayesianism that were recently put forward by Gordon Belot (2013). Referring to convergence-to-the-truth results in probability theory, Belot draws

Received September 2014; revised March 2015.

*To contact the author, please write to: Department of Logic and Philosophy of Science, University of California, Irvine, Social Science Plaza A, Irvine, CA 92697; e-mail: shuttegg@uci.edu.

†I would like to thank Jeff Barrett, Gordon Belot, Kenny Easwaran, Teddy Seidenfeld, and Kevin Zollman for helpful comments. I am especially grateful to Jim Joyce for providing a detailed written commentary. Special thanks also go to Brian Skyrms for a finite but very large number of discussions, extending back many years, on the nuances of convergence theorems in probability theory.

Philosophy of Science, 82 (October 2015) pp. 587–601. 0031-8248/2015/8204-0003\$10.00
Copyright 2015 by the Philosophy of Science Association. All rights reserved.

a bleak conclusion: “Bayesian convergence-to-the-truth theorems tell us that Bayesian agents are forbidden to think that there is any chance that they will be fooled in the long run, even when they know that their credence function is defined on a space that includes many hypotheses that would frustrate their desire to reach the truth” (500). Bayesians, we are told, cannot help but be epistemically arrogant. Convergence to the truth is bought at the price of sweeping those scenarios under the carpet in which one does not converge to the truth, regardless of how many of those there are.

This stands in stark contrast to how convergence-to-the-truth theorems are usually viewed. As, for instance, Joyce (2010) notes for a setting similar to the one discussed by Belot, convergence to the truth is not too surprising “because the data is so incredibly informative in the limit that the subject’s prior beliefs are irrelevant to her final view as a matter of logic” (446). At the limit, we would know the truth-value of any proposition about observations—on judgment day all observations will have been made. On this view, convergence to the truth for propositions about observations seems to be a minimal desideratum for learning from experience rather than a mark of epistemic immodesty.

Belot raises several important issues that merit a more extensive discussion. In sections 4 and 5 I consider two: the notion of open-minded priors and the relationship between topology and probability theory. Both cases point to certain weaknesses in Belot’s argument, but the latter also leads to a reexamination of the convergence-to-the-truth theorem and to a new positive proposal of how to understand it (secs. 6 and 7). This proposal makes use of the *measure algebra approach* that was put forward by Kolmogorov for dealing with elementary events (Kolmogorov 1948).¹ The measure algebra is a metaphysically modest mathematical structure because instead of elementary events it takes finitely discriminable outcome propositions as basic (elementary events can be recovered only by nonconstructive means). I argue that in this setting Belot’s treatment loses its bite. To set the stage, we briefly review the convergence-to-the-truth theorem and Belot’s argument in the first two sections.

2. Martingales. Convergence to the truth is a consequence of the *martingale convergence theorem*. A martingale is an infinite sequence of random variables in which, for each n , the conditional expectation of the n th random variable given the $n - 1$ previous random variables is equal to the value of the $(n - 1)$ st random variable. A martingale can be thought of as a sequence of fair gambles. If the value of the n th random variable represents

1. Other notable mathematicians who favored the latter approach are Halmos (1944) and Carathéodory (1956). For more information, see Skyrms (1995).

the total funds of a gambler at time n , then the gambler does not expect to win or lose. The martingale convergence theorem says, roughly speaking, that a martingale which meets some technical requirements converges with probability 1 (for details see, e.g., Ash 2000).

We follow Belot in specializing learning situations to the set of all infinite binary sequences. You can think of an experiment in which a coin is flipped infinitely often but also of any other kind of learning situation in which we observe whether an event is present. In order to fix ideas, we will usually refer to coin flips.

The set of all infinite binary sequences can be equipped with the topology of pointwise convergence (i.e., *Cantor space*). Now, consider some prior P over Cantor space. The conditional probability $P_n(A)$ for a measurable set A given the first n observed digits is a random variable (i.e., a measurable function from Cantor space to the reals). It is well known that the infinite sequence of conditional probabilities $P_1(A), P_2(A), \dots$ is a martingale. The sequence therefore converges with prior probability 1. Moreover, the limit is with probability 1 equal to the indicator of A , that is, the random variable that takes the value 1 on infinite sequences that are in A and the value 0 otherwise.² The indicator of A can be thought of as its truth-value.

Note that the limit is equal to the indicator only with probability 1—this is the point of departure for Belot's argument. In general there is a nonempty "exceptional" set of infinite sequences in which the conditional probabilities for A do not converge to the indicator. Let us call the set of infinite binary sequences on which the sequence of conditional probabilities converges to the indicator of A the *success set* and its complement the *failure set*. Then the result above implies that the failure set is assigned probability 0 by P while the success set is assigned probability 1.³

The martingale convergence theorem assumes that the prior probability measure is countably additive. There is a martingale convergence theorem for certain kinds of finitely additive probability measures due to Purves and Sudderth (1976) that is relevant for convergence to the truth (see Zabell 2002), although the martingale convergence theorem does not hold in general for probability measures that are only finitely additive.⁴ Moreover, it

2. The latter fact depends on A being a measurable subset of Cantor space and conditional probabilities P_n being taken relative to the first n digits of the sequence. Otherwise, conditional probabilities would converge almost surely but not necessarily to the indicator function (and thus not to the truth). For more information, see the discussion of the martingale convergence theorem and its history in Schervish and Seidenfeld (1990).

3. Consider the two sequences $P_1(A), P_2(A), \dots$ and $P_1(B), P_2(B), \dots$ of conditional probabilities for two distinct events A and B . Then the success set and the failure set for A and B need not be the same.

4. An extreme case of this is a result by Elga which says that a merely finitely additive prior can believe that it will not converge to the truth with probability 1 (Elga 2015).

should be emphasized that the convergence-to-the-truth theorem is only true under the special circumstances set out above. For instance, if the truth-value of a proposition is not determined by observations (not even by infinitely many), then convergence to the truth is not guaranteed; while the martingale convergence theorem guarantees that conditional probabilities converge, they need not converge to 0 or 1 in this case. Thus, Bayesians by no means think that they will always converge to the truth.

3. Immodest Bayesians? Bayesians do tend to think of the martingale convergence theorem as reassuring. It shows that evidence triumphs over prior opinions under the appropriate circumstances.⁵ Belot, however, invites us to view it as an Achilles heel of the Bayesian approach. His argument starts with the observation that the failure set for any A usually is nonempty. Belot rightly points out that some sequences are in the failure set because the agent has a “closed mind.” An agent may simply assign probability 0 to particular open sets of binary sequences. This kind of closed-mindedness may or may not be justifiable, depending on whether the agent has strong evidence for thinking that the true sequence is not in some open set. But Bayesians as well as anyone else can reject this kind of closed-mindedness whenever it is unjustified.

The kind of closed-mindedness Belot is after is of a different type, however. What he means to show is that there are failure sets that bear witness to a deep and unavoidable type of closed-mindedness that applies to a Bayesian even if she thinks of herself as having an open-minded prior.

In order to make this precise, we have to be clear about the meaning of an open-minded prior. One kind of open-minded prior assigns positive probability to any finite initial segment of binary sequences. An agent with such a prior does not rule out any finite sequence of evidence that she might observe. This type of open-mindedness is consistent with what may seem fairly closed-minded priors. For example, think of the set of all sequences that eventually become constantly zero. This set is countable and dense in Cantor space. If a prior assigns positive probability to each of its members and probability 0 to its complement, then every open set of Cantor space has positive probability, while the prior is closed-minded with respect to the possibility of observing infinitely many ones.

For this reason Belot (2013, 496) introduces another type of open-minded prior. This new concept of open-mindedness refers to a measurable set R of infinite binary sequences. A prior is *open-minded with respect to R* if for all data sets (finite initial segment of observations) there is an extension such that the conditional probability of R given the data set plus the extension is

5. The merging-of-opinions theorem by Blackwell and Dubins (1962) is a deeper expression of this idea.

less than $1/2$ and another extension such that the conditional probability of R given the data set plus that extension is greater than $1/2$. Such a prior exists whenever R is a countable dense subset of Cantor space. An agent who is open-minded with respect to R never fully makes up her mind as to the question whether an infinite sequence is in R .

Suppose now that R is a countable dense subset of Cantor space, and consider an open-minded prior with respect to R . Belot (2013, 497–99) develops a clever argument which shows that the failure set of an open-minded prior is residual in the space of infinite binary sequences. Its complement—the success set—is thus meagre. (The notions of meagre and residual sets are used in topology. A set is meagre if it is the countable union of nowhere dense sets. The complement of a meagre set is residual. Elements of a meagre set are atypical from a topological point of view.) Thus, relative to the topology of Cantor space, the failure set is topologically significant and the success set topologically negligible. But, despite this, the prior probability of the former is 0 and that of the latter is 1 (since this result holds regardless of the prior).

So here we have the case of a failure set that should for topological reasons not be ignored but which is essentially ignored by a Bayesian agent. Our Bayesian ignores a topologically large part of the space of sequences where she fails and focuses on the small part where she succeeds. But what is even worse, our agent is forced to have such beliefs by the formal apparatus of probability theory; even if she wanted to she cannot have a consistent prior in which the failure set has positive probability. Belot concludes that Bayesianism is epistemically flawed.

In the following sections I consider this argument and its presuppositions in three steps. In the first place, the argument rests on Belot's notion of open-mindedness. Taking a closer look at this notion does not lead to a decisive blow against Belot's conclusion, but there are reasons to doubt whether this kind of open-mindedness is something generally desirable. Second, one of the presuppositions of Belot's argument is that probability measures should be constrained by the topology of the underlying space. There is some truth in this but plausibly not enough so as to make the argument work. Finally, and most importantly, I am going to say more about convergence to the truth with arbitrarily large but finite information.

4. Open-Minded Priors. One important aspect of Belot's argument is the assumption of having an open-minded prior with respect to a measurable set R . It should be observed that this concept of open-mindedness is not as open-minded as it might appear on first inspection. The relativization—being open-minded with respect to a measurable set R —is actually important. For there cannot be a probability measure that is open-minded with respect to any measurable subset of Cantor space. Such a prior would need to assign positive

probability to each measurable set, in particular each singleton (set containing one infinite sequence); otherwise, there are sets of prior probability 0, and since the posterior probability of such a set will remain 0 forever, the prior cannot be open-minded with respect to those sets. However, a prior that assigns positive probability to each singleton does not exist. This follows from a well-known result which says that in any probability space there are at most countably many singletons with positive probability and because Cantor space is uncountable. Thus, Belot's relative notion of open-mindedness does not extend to open-mindedness tout court. One has to choose salient measurable sets with respect to which one wishes to be open-minded.

This is important for two reasons: (i) because of the role open-mindedness plays in Belot's argument but also (ii) because of the broader question when open-mindedness is a reasonable assumption. As to i, much of the effect of Belot's argument rests on the idea that one should be open-minded with respect to some set R in at least some situations. Indeed, because of a result that is due to Adam Elga, one might think this to be generally desirable (see Elga 2015). Elga shows that if a prior is not open-minded with respect to R , then there is some finite binary sequence such that, upon observing it, the posterior of R will be equal to 0 or 1, meaning that the agent becomes certain of whether R is true on the basis of a finite batch of evidence only. As we have seen, there always are sets with respect to which a prior is not going to be open-minded. Hence, there always are sets concerning which a Bayesian irrevocably makes up her mind after finitely many observations.

This may sound devastating at first: How can one rationally be certain after having made only finitely many observations that, for instance, the full sequence will not be constant from some point onward? Surely we have to observe the full sequence in order to make up our minds concerning this question. However, there are many situations in which it is perfectly reasonable to make up one's mind on the basis of finite observations with regard to hypotheses such as sequences eventually going constant. Consider a sequence of independent and identically distributed coin tosses with unknown bias p . Suppose that my prior assigns 0 probability to irrational values p and positive probability to each rational p . This prior is not open-minded with respect to the hypothesis that the sequence is eventually constant. To see this, note first that my prior assigns positive probability to $p = 1$ (the coin is two-headed) and $p = 0$ (the coin is two-tailed). The infinite sequences corresponding to these biases are the only ones that are constant; in fact, whenever $0 < p < 1$ the observed sequence will not be eventually constant with probability 1. Thus, I initially think that I might observe a constant sequence, but after observing at least one 0 and at least one 1 my posterior probability that the sequence is constant will be equal to 0. So I am not open-minded with respect to the hypothesis that the sequence is eventually going constant. But

there is nothing wrong with this unless there is something wrong with my prior, and my prior seems to be perfectly respectable.⁶

These considerations point to the broader issue of when we should be open-minded. A plausible response here is that we should be open-minded whenever we cannot rule out any possibilities (i.e., whenever we know nothing about the process that generates the sequence). But if we really know nothing about the process, then why should we think that finite batches of evidence are relevant for the probability of R as is required by open-mindedness relative to R ? In fact, why should one have a definite prior at all and not move to imprecise probabilities in order to be as noncommittal as possible?

This is certainly not the place to try to answer all of these questions. What I hope to have demonstrated is that Belot's concept of open-mindedness is more nuanced than one might think. The open-mindedness of a prior with respect to a set is not a maximally open state of mind that does not rule out any possibilities; it rather represents a state of mind that is committed to some possibilities at the expense of others. Open-mindedness with respect to one set implies closed-mindedness with respect to others. Often, this kind of closed-mindedness is reasonable. A Bayesian agent is closed-minded with respect to the failure set (it has probability 0), but closed-mindedness with respect to the failure set is not unavoidably unreasonable or irrational.

5. Topology and Measure. Even if you grant this point, there are two reasons for why you might still feel disturbed by Belot's result. First, the failure set is topologically large, so there appears to be an independent reason not to ignore it. Second, a Bayesian must be closed-minded relative to the failure set. I discuss these two objections in turn, starting with the first one in this section.

While a failure set can be topologically large, a Bayesian might insist that probability theory is not topology. Belot himself refers to various results that show how topological notions and measure-theoretic notions can come apart (Belot 2013, sec. 3). Meagre sets can have probability 1. Residual sets can have probability 0. The epistemic freedom of an agent even allows her to assign probability 1 to a denumerable set or a finite set or a singleton. The only constraint is that degrees of belief be consistent.

Furthermore, the mathematical structure of measure theory is very different from the mathematical structure of topology. Topological notions are invariant under homeomorphisms (a homeomorphism is a continuous map from one topological space into another). But measure-theoretic notions do not in general exhibit this invariance. Taking all of this together suggests

6. Thanks to Jim Joyce for the example and for raising several of the points mentioned in the following paragraphs.

that topological and probabilistic concepts are fairly independent of each other and that results about the topology of a space do not prescribe specific probability distributions for that space. From a Bayesian perspective, this makes a lot of sense. Topology is a mathematical theory of concepts like closeness and limit point, whereas probability is a mathematical theory of rational degrees of belief. The two theories have very different domains, and so there is no reason to suppose that there are any general principles connecting the two in the way required by Belot's argument, which appears to be something along the lines of "if a set is residual in the topology, then it should have positive probability."

Although this response seems to be correct, it will probably not convert any doubter. There is more to say, though. The role of topology points to a deeper presupposition of Belot's argument—its reliance on the infinite. If we restricted ourselves to finite sequences of data, convergence to the truth for propositions about those data would be completely uncontroversial; in the finite realm you know the truth-value of every proposition about observations after having made all the observations. The martingale convergence theorem shows that this carries over, in a certain sense, to the case of infinite sequences. But in this case (unlike the finite one) we have to deal with the problem of nonempty failure sets that no Bayesian agent can avoid. In the next two sections I show how we can deal with this problem by applying one plausible way of finitist thinking.

6. Modest Metaphysics. Belot clearly thinks of infinite binary sequences as genuine epistemic possibilities. For example, in the context of convergence to the truth he states that "there is a rich infinite family of sequences the agent could be shown that would prevent convergence to the truth" (Belot 2013, 484).

But are infinite sequences something that can be learned? This question is important if we take seriously some very general epistemic constraints, such as our own epistemic finitude. Consider the paradigm examples of inductive learning Belot mentions (Belot 2013, 493): tossing coins, measuring the successive bits of the binary expansion of a constant of nature, or determining whether there is more gold in India or in China, minute by minute. In the context of learning from experience, those sequences of observations can be certainly thought of as, possibly very large, finite binary sequences. If there is no upper bound to observations, it is convenient to work with infinite binary sequences in order to approximate arbitrarily large finite sequences. But it is essential to interpret limiting results very carefully when agents do not actually have access to infinite observations. Our motivation for treating infinite binary sequences as idealized objects is thus empirical: infinite sequences make distinctions between events that cannot be made by finite observations or measurements, regardless of how precise they are.

This perspective calls for a metaphysics that is more modest than the metaphysics of standard probability theory. The mathematical superstructure of standard probability theory allows degrees of belief to refer to all kinds of infinitary objects. Within that superstructure, infinite sequences are indeed epistemic possibilities—that is, something one might coherently suppose (in the indicative mode). An agent might suppose, just as Belot suggests, that she is shown a sequence from the failure set. According to the martingale convergence theorem this is a probability 0 event, yet it is an epistemic possibility. Such an epistemic possibility does not need to be something that one can learn, however. The question now is which parts of the mathematical superstructure are relevant for learning.⁷

Let us start by taking a closer look at the success set and the failure set. So far we have only seen that the failure set may be residual and the success set meagre. But it is also important to observe how the two sets relate to each other in the topology. Because the failure set in Belot's example is residual, it is uncountable and dense in the space of infinite binary sequences. Belot also shows that the failure set is dense for any prior over Cantor space. It follows that every sequence in the success set can be approximated arbitrarily closely by a sequence in the failure set: if x is a sequence in the success set, then for any n there exists a sequence y in the failure set that agrees with x in the first n elements. In other words, any open set containing x also contains a sequence that is in the failure set.

Under very plausible assumptions, the success set is also dense in the space of infinite binary sequences. We only need to assume that the prior is open-minded in the sense of assigning positive probability to any open set.⁸ Suppose that the success set is not dense. Then there exists an open set B such that all sequences in B are in the failure set. But because of the martingale convergence theorem B must have prior probability 0. This contradicts the assumption that all open sets have positive prior probability. Hence, the success set is dense.

We get the following important result:

Empirical indistinguishability. The success set and the failure set of an open-minded prior are both dense in Cantor space. Thus, any sequence in the failure set can be approximated arbitrarily closely by a sequence in the

7. Of course, I do not mean to imply that the standard framework should be abandoned. Apart from its mathematical fruitfulness, standard probability theory might also be useful for many epistemic questions. What I wish to point out is that there are epistemic constraints on the superstructure once we put it in the context of learning from experience.

8. The prior constructed in Belot (2013), n. 37, is an example of a prior that is open-minded both in this and in Belot's sense. The result reported here could be reformulated appropriately for any prior.

success set and vice versa. Sequences cannot be identified as belonging to the success set or the failure set by arbitrarily precise finite observations.

The success set and the failure set cannot be distinguished observationally if we only have an arbitrary finite number of observations. This indicates that the existence of a failure set may not be a significant threat to Bayesian convergence to the truth with increasing but finite batches of evidence. For any finite time, each sequence in the failure set can be associated with at least one sequence in the success set. There is convergence to the truth for any proposition whose truth-value depends only on a finite number of observations for the success set. These propositions approximate all other events. So, in a sense, the Bayesian converges to the truth in terms of having degrees of beliefs that get closer to the indicator without necessarily ever reaching it, since the number of observations is finite.

There is no failure set on this view. The failure set ceases to be relevant once we stop making distinctions that can only be made by being infinitely precise. In the next section I outline how this informal idea can be made precise.

7. Measure Algebras. Despite using classical mathematics in his famous (1933) monograph, Kolmogorov is a champion of the finite. Later work by Kolmogorov can be used to turn the idea of a modest metaphysics discussed in the previous section into a substantial theory (1948; trans. in Kolmogorov 1995). For Kolmogorov, one of the drawbacks of his 1933 theory of probability is that “the notion of an elementary event is an artificial superstructure imposed on the concrete notion of an event. In reality, events are not composed of elementary events, but elementary events originate in the dismemberment of composite events” (1995, 61). Elementary events are possible worlds, for example, the infinite binary sequences of Cantor space. What Kolmogorov is suggesting is to take outcome propositions (such as “the first three digits are 110”) as basic and view possible worlds as artifacts deriving from outcome propositions. He then goes on to show that this idea is captured mathematically by *metric Boolean algebras*.

Mathematical structures like Cantor space make more discriminations than we should ascribe to reality. Consider, for instance, the open set of all infinite sequences starting with 110, and suppose that we remove from it the sequence 1100000 . . . (110 followed by zeroes). The alleged difference between this set and the original one is smaller than any finite discrimination (any number of zeroes you observe after the third trial is compatible with both propositions). The outcome described by both sets really expresses something about the first three observations. The additional distinctions that are being made are irrelevant to this outcome. A metric Boolean algebra takes the elementary events out of Cantor space by identifying these two sets, and

similar ones, with each other. It does this by factoring out sets of probability 0 for prior probability measures that assign (i) positive probability to every open set and (ii) 0 probability to each particular infinite sequence. Such a prior can be thought of as an open-minded, antimetaphysical prior—it is antimetaphysical since assumption ii expresses the belief that no individual infinite sequence is the true one.⁹

For such a prior, two measurable sets are said to be of the same *metric type* if their symmetric difference has probability 0. (The symmetric difference of two sets A and B is the set of sequences that are in A but not in B or vice versa.) Being of the same metric type is an equivalence relation, so we may identify all measurable sets that have the same metric type.¹⁰ By identifying all sets of the same metric type, we cast out infinite sequences since each of them is of the same metric type as the empty set. Metric types are the basic elements Kolmogorov wanted to have—composite events that do not depend on the concept of an elementary event.

The quotient construction through metric types yields a Boolean algebra by transferring the Boolean operations from the original space to the new class of sets in the natural way. The original probability can likewise be used for the quotient algebra by requiring that the probability of a metric type is equal to the probability of an event of that metric type. The resulting structure is a metric Boolean algebra, that is, a Boolean algebra with a (in general only finitely additive) probability measure that assigns 0 probability only to the null element of the Boolean algebra and probability 1 only to its unit element. (The null element corresponds to the empty set and all sets of probability 0 in the original space, and the unit element to all sets of probability 1.)

Taking the distance between two metric types to be the probability of their symmetric difference defines a metric. Since the probability measure in our original space was assumed to be countably additive, the Boolean metric space is in fact a complete metric space (Kolmogorov 1995, 60). The complete metric space is a Boolean σ -algebra for which countable additivity holds automatically (since convergence for metric types is defined as the symmetric difference going to 0; see Kolmogorov 1995, 62–63). If we do not have countable additivity at the outset, it can easily be introduced by completing the metric Boolean algebra. The Boolean σ -algebra together with its probability measure is called a ‘measure algebra’ (Halmos 1944).

In the metric Boolean algebra there are only outcomes and no possible worlds (infinite sequences). For any metric Boolean algebra, possible worlds can be recovered through the representation theorem of Stone. According to

9. For the underlying metaphysics, cf. Skyrms (1993).

10. That is, we are forming a quotient algebra by taking the σ -algebra modulo the σ -ideal of sets of probability 0.

the isomorphism between Boolean algebras and fields of sets given in Stone (1936), an outcome corresponds to the set of possible worlds where the outcome occurs. Possible worlds are maximally specific outcomes (they are the prime ideals of the Boolean algebra; see Łoś 1955). Since Stone's theorem uses the axiom of choice, possible worlds are cognitively remote, highly idealized entities.¹¹

Our measure algebra is thus a fairly satisfying representation of that part of a probability space that is accessible to finite observations. Now, returning to our original question, what does convergence of conditional probabilities mean in the new framework? The short answer is that in Cantor space the failure set has probability 0; hence, it is associated with the null element of the corresponding metric Boolean σ -algebra. The success set, however, corresponds to the unit element of the metric Boolean σ -algebra because its probability is 1. Thus, convergence to the truth holds without exceptions.

Let us look at this in a bit more detail. The conditional probability $P_n(A)$ is a random variable. Recall that a random variable is a measurable function that assigns a real number to each possible world. That it is measurable means that its inverse maps each Borel set B to a set in the σ -algebra. The Borel sets are countable unions and intersections of open intervals of the real line. Thus, a measurable function does not exceed the standard conceptual resources of the real numbers.

Since there are no possible worlds in the measure algebra but only outcomes, random variables cannot be defined in the measure algebra. Instead, random variables are associated with σ -homomorphisms (Łoś 1955). The idea is simple: every random variable X from Cantor space to the reals generates a map from Borel sets to the measurable subsets of Cantor space by mapping each Borel set B to the set of elementary events that X maps into B . The map associated with a random variable can be used in measure algebras. A σ -homomorphism is a map from the Borel sets to the Boolean σ -algebra that preserves countable unions and complementation. Two random variables induce the same σ -homomorphism from the Borel sets to the Boolean σ -algebra if they agree almost surely. This makes it possible to define the integral of a σ -homomorphism over the measure space as the integral of an inducing random variable over the probability space (see Sikorski [1949] for details). Thus, from a probabilistic perspective, random variables and induced σ -homomorphisms are essentially the same.

Each $P_n(A)$ induces a σ -homomorphism f_n in the following way: the function f_n maps each Borel set B to the metric type of the set of all infi-

11. The analogue of the Stone theorem can fail for Boolean σ -algebras. However, for every Boolean σ -algebra B there is a σ -field of sets F and a σ -ideal I such that B is isomorphic to the quotient algebra F/I . This is the representation theorem of Loomis (1947) and Sikorski (1948).

nite sequences to which $P_n(A)$ assigns a value in B . This means that for each value between 0 and 1 f_n identifies those outcomes in the Boolean σ -algebra where conditional probabilities after n observations take on that value (modulo observationally irrelevant distinctions). Then the sequence of σ -homomorphisms f_1, f_2, \dots converges to the σ -homomorphism f that is induced by the indicator of A (Sikorski 1949). Since the indicator of A is equal to 0 or 1, $f(\{0, 1\})$ is the unit element of the Boolean σ -algebra. Moreover, since f preserves Boolean operations, the unit element of the Boolean σ -algebra is the union of the outcome $f(\{1\})$ where the metric type of A is true and the outcome $f(\{0\})$ where the metric type of A is false. This is the sense in which we have convergence to the truth in the measure algebra.

It should be noted that $f(\{1\})$ and $f(\{0\})$ may themselves be idealized elements of the Boolean σ -algebra if A is an infinitary event (e.g., the limiting relative frequency of 1s is one-half). If we only allow outcome propositions that correspond to finite binary sequences, we would have a metric Boolean algebra instead of a metric Boolean σ -algebra. In the metric Boolean algebra there are outcomes that are arbitrarily close to the convergence-to-the-truth outcomes. By completing this metric Boolean algebra with respect to the metric, we get a Boolean σ -algebra in which sets such as those in question arise as limiting elements, while the elements of the metric Boolean algebra are dense in the metric Boolean σ -algebra.

As noted above, a consequence of these considerations is that the failure set gets absorbed into the null element of the metric Boolean σ -algebra. This may seem pretty ad hoc. Looking only at this conclusion might suggest that we did nothing more than sweep the failure set deeper under the carpet. However, eliminating the failure set is the result of the main idea of Kolmogorov's approach—to identify events that cannot be finitely discriminated. This is the reason why the failure set is not part of the measure algebra. Far from being ad hoc, our main conclusion is firmly grounded on a plausible epistemic constraint.

Let us now reconsider Belot's example in the context of an antimetaphysical open-minded prior. If R is a countable dense subset of Cantor space, then R has probability 0 according to our prior. Thus, the null element of the measure algebra is its metric type, and it gets a probability of 0 throughout the process of learning from experience. This reflects our choice of prior. Some may find this prior to be too radical, especially because it excludes many priors that are open-minded in Belot's sense. What if one wants to assign positive probability to especially salient countable dense sets such as the binary expansions of rationals or computable reals? Here one can also apply the measure algebra framework. Suppose that our prior assigns positive probability to each element of a countable dense subset R of Cantor space. The prior is metaphysical since it thinks that each element of R can be true with positive probability. For simplicity, we assume again that the prior is also

open-minded in the sense of assigning positive probability to each open set. (The prior is thus open-minded with respect to R .) By the same reasoning as above, convergence to the truth holds without any qualification by a failure set. Even though the prior is metaphysical, it factors out many differences that cannot be discriminated by finite means. As a result, we again have convergence to the truth without qualification by a failure set.

The main difference between a metaphysical prior and an antimetaphysical prior is that their measure algebras include different outcomes. The elements of a measure algebra depend on the prior since different priors can have different sets of measure 0. Thus, the difference in opinion between agents becomes amplified when we move from the standard framework to the measure algebra. At the same time, two agents may hold the same beliefs in the measure algebra but have slightly different beliefs when we look at the more fine-grained level of the standard framework. For these reasons one might think that the measure algebra framework is not a good substitute for the standard measure-theoretic framework. I do not suggest to always use measure algebras instead of the standard approach, but I think that each approach has its virtues and vices. For convergence to the truth using the measure algebra is particularly apt since it allows one to analyze increasingly large but finite sequences of observations. This does not mean that the measure algebra is the correct framework for all questions regarding degrees of belief or that the classical measure-theoretic framework is mathematically flawed.

8. Conclusion. I have shown that infinite sequences are not necessary for Bayesian learning from experience and that they can be viewed as artifacts of an idealization. This result defuses Belot's main argument. However, I agree with Belot and others that the value of convergence-to-the-truth theorems and merging-of-opinions results should not be overstated. They make substantive assumptions about a learning situation. What they do show is that in certain learning situations the influence of individual priors vanishes and that posterior probabilities correctly reflect increasing information.

REFERENCES

- Ash, Robert B. 2000. *Probability and Measure Theory*. San Diego, CA: Academic Press.
- Belot, Gordon. 2013. "Bayesian Orgulity." *Philosophy of Science* 80:483–503.
- Blackwell, David, and Lester Dubins. 1962. "Merging of Opinions with Increasing Information." *Annals of Mathematical Statistics* 33:882–86.
- Carathéodory, Constantin. 1956. *Mass und Integral und ihre Algebraisierung*. Stuttgart: Birkhäuser.
- Elga, Adam. 2015. "Bayesian Humility." Unpublished manuscript, Princeton University.
- Halmos, Paul R. 1944. "The Foundations of Probability." *American Mathematical Monthly* 51:497–510.
- Joyce, James M. 2010. "The Development of Subjective Bayesianism." In *Handbook of the History of Logic*, vol. 10, *Inductive Logic*, ed. Dov M. Gabbay, Stephan Hartmann, and John Woods, 415–76. Amsterdam: Elsevier.

- Kolmogorov, Andrey N. 1933. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Berlin: Springer.
- . 1948. “Algèbres de Boole métriques complètes.” *Zjazd Matematyków Polskich* 20:21–30. Trans. Richard Jeffrey in Kolmogorov 1995.
- . 1995. “Complete Metric Boolean Algebras.” *Philosophical Studies* 77:57–66.
- Loomis, Lynn H. 1947. “On the Representation of a σ Complete Boolean Algebra.” *Bulletin of the American Mathematical Society* 53:757–60.
- Łoś, Jerzy. 1955. “On the Axiomatic Treatment of Probability.” *Colloquium Mathematicum* 3:125–37.
- Purves, Roger A., and William D. Sudderth. 1976. “Some Finitely Additive Probability.” *Annals of Probability* 4:259–76.
- Schervish, Mark J., and Teddy Seidenfeld. 1990. “An Approach to Consensus and Certainty with Increasing Information.” *Journal of Statistical Planning and Inference* 25:401–14.
- Sikorski, R. 1948. “On the Representations of Boolean Algebras as Fields of Sets.” *Fundamenta Mathematica* 35:247–56.
- . 1949. “The Integral in a Boolean Algebra.” *Colloquium Mathematicum* 2:20–26.
- Skyrms, Brian. 1993. “Logical Atoms and Combinatorial Possibility.” *Journal of Philosophy* 90:219–32.
- . 1995. “Strict Coherence, Sigma Coherence and the Metaphysics of Quantity.” *Philosophical Studies* 77:39–55.
- Stone, Marshall H. 1936. “The Theory of Representations for Boolean Algebras.” *Transactions of the American Mathematical Society* 40:37–111.
- Zabell, Sandy L. 2002. “It All Adds Up: The Dynamic Coherence of Radical Probabilism.” *Philosophy of Science* 69:98–103.