

Bayesian humility*

Adam Elga

July 25, 2014

Abstract

Say that an agent is *epistemically humble* if she is less than certain that her opinions will converge to the truth, given an appropriate stream of evidence. Is such humility rationally permissible? According to *the orgulity argument* (Belot 2013): the answer is “yes,” but long-run convergence-to-the-truth theorems force Bayesians to answer “no.” That argument has no force against Bayesians who reject countable additivity as a requirement of rationality. Such Bayesians are free to count even extreme humility as rationally permissible.

1 Introduction

Presented with Bayesian confirmation theory, it is easy to feel cheated. One might have hoped for a substantive, detailed account of what sorts of evidence support what sorts of scientific hypotheses. Instead one is told how one’s evidence determines reasonable attitudes toward such hypotheses *given a prior* (an initial probability function). And one is told that different priors deliver different outputs, even for the same batch of total evidence.

One might worry that given this dependence, Bayesianism is ill-placed to explain the significant agreement observed among reasonable scientists,

*Previously titled: “Dr. Belot or: How I Learned to Stop Worrying about Bayesian Convergence Theorems and Love Finite Additivity”. For helpful comments and discussion, thanks to Andrew Bacon, Gordon Belot, Cian Dorr, Kenny Easwaran, Kevin Kelly, Bas van Fraassen, Jim Hawthorne, Brian Weatherson, and Jonathan Wright. For an extremely congenial work environment in July 2014, thanks to Mindy and Gene Stein.

or to deliver an objective account of confirmation in science.^{1,2}

In the face of this worry it is natural to seek comfort from some remarkable long-run “convergence-to-the-truth” and “washing out” theorems. These theorems show that unless priors differ radically, differences between them become negligible in the long run, under the impact of a stream of common evidence. This is sometimes thought to take the sting out of the above worry, by showing that many differences between priors don’t end up mattering.

But Belot (2013) argues that rather than helping Bayesianism, these convergence theorems are a liability to it. The argument is that the theorems preclude Bayesians from counting as rational “a reasonable modesty” about whether one’s opinions will approach the truth in the long run.

I will argue:

1. Long-run convergence theorems are no liability to *finitely additive Bayesianism*, a version of Bayesianism that rejects countable additivity as a requirement of rationality.³ Defenders of finitely additive Bayesianism are free to count any amount of humility about convergence to the truth—even extreme pessimism—as rationally permissible.⁴
2. Bayesians never needed to appeal to long-run convergence theorems in the first place. In response to worries about scientific objectivity, Bayesians can and should instead appeal to *short-run* convergence theorems (Howson and Urbach 2006, 238; Hawthorne 2014, §5). Those theorems do not require countable additivity.

¹For expressions (but not always endorsements) of this worry, see Chalmers (1999, 133) as cited in Vallinder (2012, 8), Easwaran (2011, §2.6), Earman (1992, 137), Howson and Urbach (2006, 237), Hawthorne (2014, §3.5).

²Note that for present purposes, a confirmation theory may count as Bayesian even if it imposes constraints on priors more restrictive than mere coherence. Thanks here to Cian Dorr.

³Here I apply observations from Juhl and Kelly (1994, 186) and Howson and Urbach (2006, 28–29).

⁴Weatherson (2014) convincingly argues that *unsharp* Bayesians (who hold that states of graded opinion should be represented not by probability functions, but rather by sets of probability functions) are also free to count humility about convergence to the truth as rationally permissible.

2 Long-run convergence theorems

What are the long-run convergence theorems? And why think that they get in the way of a Bayesian being sufficiently humble about whether her opinions will converge to the truth?

For simplicity, follow Belot (2013) in restricting attention to the following setup. Consider an immortal investigator whose evidence consists of successive digits from a countably infinite binary sequence (a sequence consisting of 0s and 1s). The investigator receives one digit of the sequence per day, and is interested in H , a fixed hypothesis about the whole sequence. For example, H might be the proposition that after a certain point, the sequence consist of all 1s. (For convenience I treat interchangeably a proposition about the sequence, and the corresponding set of sequences for which that proposition holds.)

Now apply Bayesian confirmation theory to this setup. In particular, suppose that the investigator starts with a prior probability function P defined over an appropriate domain that includes H ,⁵ updates by conditionalization each time she receives a digit, and is certain of all of the above.⁶

Before seeing any digits, the investigator might wonder: in the limit of seeing more and more digits, how likely is it that I will arrive at the truth about H ? In other words, how likely is it that my probability for H will converge to 1 if H is true and to 0 otherwise?

A pessimistic answer to that question is: I am unlikely to converge to the truth (about H). A more confident answer is: I will probably converge to the truth. A maximally confident answer is: my probability that I will converge to the truth equals 1.

Long-run convergence theorems entail that *if the investigator's probability function is countably additive*, then she is committed to the maximally confident answer.⁷ In other words: (countably additive) Bayesian confirmation theory entails that rationality requires investigators in the above situation

⁵We assume throughout that P is defined over at least the Borel-measurable subsets of infinite binary sequences, where the set of sequences is given the natural topology—the topology that has as a basis the set of *rectangles*, each of which is the set of sequences satisfying a finite number of constraints of the form “digit k equals b ”, where k is a natural number and b equals 0 or 1.

⁶For present purposes, we needn't decide the question of what happens when an investigator receives evidence that she had previously assigned zero probability.

⁷For a proof, see Halmos (1974, Theorem 49B, p. 213) (as cited in Schervish and Seidenfeld (1990, 410)). For an explanation emphasizing the role that countable additivity plays in a similar proof, see Kelly (1996, 325–327). For other results of this kind, see Schervish and Seidenfeld (1990). For further discussion see Earman (1992, 144–145).

to have *full* confidence that their opinions will converge to the truth.

But Belot (2013) gives an ingenious argument that rationality requires no such thing.

3 The orgulity argument

Here is a stripped-down exposition of what I shall call “the orgulity argument”—the main argument from Belot (2013).

As before, let H be a hypothesis about the infinite binary sequence that is in the domain of the investigator’s probability function. Examples of such hypotheses include: that the sequence eventually becomes periodic, that it ends with the pattern “01010101...”, that it is computed by a Turing Machine,⁸ or that it contains infinitely many 0s.⁹

Say that an investigator is **open-minded** with respect to H if for every finite batch of evidence, there is a finite extension of it that would lead her to assign probability greater than $1/2$ to H , and also a finite extension of it that would lead her to assign probability less than $1/2$ to H .^{10,11} An open-minded investigator commits to never irrevocably making up her mind about whether H or not- H is more likely. Here is the first premise of the orgulity argument:

Premise 1 It is rationally permissible to be open-minded with respect to some hypothesis.¹²

⁸I.e., that the function $d(i)$ giving successive digits of the sequence is a computable function.

⁹Note that there is no requirement here that H be countable. Cf. Belot (2013, n. 32).

¹⁰Here I adopt the suggestion from Weatherson (2014) to modify the definition of “open-minded” given in Belot (2013, 496) to introduce a pleasing symmetry. Nothing of substance hinges on this.

¹¹For an example of a countably additive prior that is open-minded with respect to a countably infinite hypothesis, see Belot (2013, n. 32). For an example of a countably additive prior P that is open-minded with respect to an uncountable hypothesis H , take H to be the set of sequences in which 1 occurs with limiting relative frequency $2/3$ and take P to be $(B_{1/3} + B_{2/3})/2$, where B_v is the *Bernoulli measure* with bias v (the probability measure that treats the digits of the sequence as if they were generated by independent tosses of a coin that has probability v of generating “1” on each toss). P is open-minded because it regards finite evidence streams consisting of almost all 1s as strongly confirming H , and streams consisting of almost all 0s as strongly disconfirming H .

¹²Here and below, all premises are assumed to concern the setup in which a Bayesian investigator successively learns digits from an infinite binary sequence.

Now take any Bayesian investigator with a countably additive prior who is open-minded about some hypothesis H , and consider the set T of sequences that get her to converge to the truth about H . (In the remainder of this section, “converge to the truth” abbreviates “converge to the truth about H ”.) We noted above that convergence theorems entail that this investigator must assign probability 1 to T . We will now see that T is in one sense a “tiny” set.

Start by defining *the Banach-Mazur game*. In this game two players generate an infinite binary sequence together, starting with the empty sequence. The players alternate moves; at each move a player extends the sequence by appending whatever finite block of digits she wishes. The goal of the player who moves second is to have the resulting infinite sequence fall outside of some fixed set G .¹³

G is said to be **meager** if there exists a winning strategy for the second player in this game—in other words, if the second player can force the generated sequence to *avoid* G . When a set of sequences is meager, it is “tiny” in one sense—it is *easy to avoid*.¹⁴ It is sometimes said that sequences “typically” have a property if the set of sequences that fail to have the property is meager.

Now for the main fact that drives the argument:

Fact (Belot 2013, 498-499) The set of sequences that get an open-minded Bayesian investigator to converge to the truth, is meager. In other words: “typical” sequences *prevent* the investigator from converging to the truth.¹⁵

Since we can see the truth of this fact, so can a reasonable open-minded investigator. She can see that typical sequences prevent her from converging to the truth. Given this, it seems permissible for her to be less than certain that she will converge to the truth. That is the next premise of the argument:

¹³This is actually the special case of the Banach-Mazur game appropriate to the present context. For a general discussion, see Oxtoby (1980).

¹⁴This is just one of several equivalent characterizations of the meager sets.

¹⁵Proof: In the Banach-Mazur game, player 2 can force the generated sequence to be one that prevents the investigator from converging to the truth about H by at each of her turns appending “a string of bits that causes P to [assign probability greater than $1/2$ to H] followed by a string of bits that causes P to [assign probability less than $1/2$ to H]. This always results in an infinite sequence that [causes P to not converge on any probability for H].” (Belot 2013, n. 41) Player 2 can implement this strategy because P is open-minded.

Premise 2 If it is rationally permissible to be open-minded about a hypothesis, then it is rationally permissible to be less than certain that one will converge to the truth about that hypothesis.

From Premises 1 and 2 we get the conclusion of the orgulity argument:

Conclusion It is rationally permissible to be less than certain that one will converge to the truth about some hypothesis.

We saw in §2 that (countably additive) Bayesianism entails the negation of this conclusion. So if the argument is sound, then that theory stands refuted.

4 Points in favor of the orgulity argument

Before suggesting a response to the orgulity argument, let me make a few points in its favor.

A defender of countably additive Bayesianism might try to reject Premise 1 by claiming that open-mindedness is irrational. That is, he might propose that investigators are rationally required to permanently make up their minds about whether H or not- H is more likely, upon receiving an appropriate finite batch of evidence.

That proposal is unappealing because (in the presence of countable additivity), it entails that investigators are rationally required to become *certain* about whether H is true upon receiving an appropriate finite batch of evidence. (For a proof, see Appendix A.)

For example, suppose that H is the claim that the sequence contains infinitely many 0s. The above proposal entails that if a Bayesian investigator is rational, some finite string of digits would get her to assign probability 0 or 1 to H . But that is absurd. It is absurd that rationality requires every investigator to count some finite string of digits as settling with certainty whether the whole string contains infinitely many 0s.

What about Premise 2? Belot (2013, 500) considers an opponent who rejects Premise 2 for the following reason: sequences of evidence digits that prevent an open-minded investigator from converging to the truth are skeptical scenarios, and the investigator may therefore reasonably assign them total probability zero. Belot responds that such sequences are not skeptical scenarios. Whatever one thinks of that response, however, an additional response is available: Even granting that the scenarios in question are skeptical scenarios, it does not immediately follow that they deserve zero probability.

As an example, consider a regularity that is very well confirmed: that gravity is an attractive. Here is a skeptical scenario: one year from now, gravity will suddenly turn repulsive. Given our evidence, that scenario deserves only a miniscule amount of probability. But that the scenario is a skeptical one does not immediately show that it deserves absolutely *no* probability. So simply calling failure-to-converge scenarios “skeptical scenarios” does not on its own make it reasonable to reject Premise 2.

Furthermore, the Fact gives at least initial support to Premise 2. It is unsettling to think that to be rational one *must* be certain that one will converge to the truth, given that “typical” sequences prevent one from doing so.¹⁶

Moral: the orgulity argument has some force as an objection to countably additive Bayesianism.

5 Finitely additive Bayesianism permits humility

Happily, the argument has no force at all against finitely additive Bayesianism, a version of Bayesianism that rejects countable additivity as an across-the-board requirement of rationality.¹⁷ That is because finitely additive Bayesians¹⁸ can comfortably accept the conclusion of the argument. They can accept that it is rationally permissible for an open-minded investigator in the sequence situation to be less than certain that she will converge to the truth.

Indeed, they can (if they wish) accept something much stronger. Let us say that an investigator in the sequence situation is *completely pessimistic* if she is certain that she will “converge to the false” — that her probability for H will converge to 0 if H is true and to 1 otherwise. It turns out that some open-minded investigators with finitely additive priors are completely pessimistic.¹⁹ (For a proof, see Appendix B.)

¹⁶Here I grant for the sake of argument that topological notions of size are relevant to what propositions it is rationally permissible to assign positive probability. One might of course flatly deny this, and so deny that Premise 2 has any appeal.

¹⁷Juhl and Kelly (1994, 185–188) and Howson and Urbach (2006, 28–29) make similar points in response to the concern that the long-run convergence theorems yield implausibly strong constraints on rationality.

¹⁸By “finitely additive Bayesians” I mean Bayesians who reject countable additivity as a requirement of rationality.

¹⁹Some may find it uncomfortable to think that such extreme pessimism is rationally permissible. That discomfort might derive from discomfort with rational failures of conglomerability in countable partitions. If so, there there is little recourse but to impose countable

So finitely additive Bayesians are free to count even complete pessimism as being rationally permissible. That is as much humility as anyone can demand. Furthermore, finitely additive Bayesianism has significant independent appeal.²⁰

Moral: the orgulity argument has no force against finitely additive Bayesianism, a viable alternative to countably additive Bayesianism.

6 Convergence in the short-run

Recall from §1 the motive given for appealing to long-run convergence theorems: The deliverances of Bayesianism depend on a choice of prior probability function. As a result, Bayesianism faces the charge of being excessively subjective, and of not sufficiently explaining agreement among reasonable scientists.

In response to those charges, it is tempting to appeal to long-run convergence theorems in order to show that differences between rational priors disappear in the long run. Bayesians who reject countable additivity cannot appeal to those theorems in this way, since those theorems depend on countable additivity. So it might seem that such Bayesians give up a valuable defense against the charge of excessive subjectivity.

But in fact, they do not. For the long-run convergence theorems were red herrings all along—they never provided an answer to the charge of excessive subjectivity. That is because that charge concerns not what happens in the infinite long run, but rather what happens in the near future (Earman 1992, 148, Howson and Urbach 2006, 238). The charge is that Bayesians cannot explain the extent to which reasonable scientists agree *now*, on the basis

additivity (Kadane et al. 1986, Schervish et al. 1984). But if the discomfort is simply with probability functions that manifest *complete* pessimism, then one might instead adopt conditions that rule out complete pessimism without ruling out more moderate pessimism.

²⁰See, for example, de Finetti (1974), Savage (1954) and Levi (1980). Works that at least take very seriously the hypothesis that countable additivity should be rejected include Seidenfeld and Schervish (1983), Dubins and Savage (1965), Kelly (1996), and Juhl and Kelly (1994).

Of course, there are objections to rejecting countable additivity as well, and an assessment of the costs and benefits of doing so is beyond the scope of this discussion. Such an assessment would need to address concerns about susceptibility to infinite Dutch Books (Bartha 2004, Seidenfeld and Schervish 1983), the possibility of paradoxical-seeming failures of conglomerability (Kadane et al. 1986, Schervish et al. 1984), the possibility of uniform distributions over countably infinite spaces (de Finetti 1974, 122), violations of intuitive comparative dominance principles (Easwaran 2013), as well as considerations of general mathematical utility (Dubins and Savage 1965).

of their actual (finite) batches of shared evidence. What happens in the limit of infinite investigation is not directly relevant.

That does not mean that Bayesians are defenseless against the charge of excessive subjectivity. Other convergence theorems—call them “short-run” convergence theorems—do help Bayesians ward off that charge. As a toy example of such a theorem, consider two Bayesian agents who are about to observe what they both regard to be independent random draws from an urn with an unknown proportion of white and black balls. Straightforward calculations show that unless the agents start out with extremely different opinions about what the proportion is, they will be confident that their opinions about the urn’s composition will become extremely similar—not just in the limit of infinite draws, but *soon* (after a small number of draws).²¹

Now, the assumption of independent sampling in the above case is admittedly quite strong. But short-run convergence results have been proven that rely on significantly weaker assumptions. For example, Hawthorne (1993, Theorem 6) can be thought of as a short-run convergence result that applies to Bayesian agents who do not regard successive batches of evidence as independent random draws. What is required instead is a much weaker condition: that the agents expect successive batches to be, on average, at least slightly informative about the hypothesis in question.²²

The bottom line is that while existing theorems do not decisively settle the matter (one can always claim that in a particular case, there is or should be more convergence to the truth than Bayesianism can account for), short-run convergence theorems at least substantially address the charge of excessive subjectivity.

Happily, none of these short-run convergence theorems rely on countable additivity. So finitely additive Bayesians can freely appeal to them.

²¹For short-run convergence results for cases of roughly this kind, see Savage (1954, §3.6), Edwards et al. (1963, 541–545), Earman (1992, 142–143), Howson and Urbach (2006, 239), and Hawthorne (2014, §4.1).

²²I have given here only the barest sketch of the wonderful convergence theorem presented in Hawthorne (1993, Thm. 6) and explained in a simplified form in Hawthorne (2014, §5).

Appendix

Notation: Unless otherwise noted, “sequence” shall mean: countably infinite binary sequence. Let S be the set of countably infinite binary sequences. When $x \in S$, we write x_i for the i^{th} digit of x , and x^i for the set of sequences whose first i digits agree with x . For any finite binary sequence s , let $[s]$ be the set of infinite binary sequences that start with s .

A Proof that countably additive priors are either closed-minded or extremely open-minded

In §4 it is claimed that under the assumption of countable additivity, if it is irrational for an investigator to be open-minded, then it is rationally required for her to be closed-minded in a certain sense. Here we prove a slightly stronger claim from which the above claim easily follows.

Say that an investigator with prior P is **closed-minded** with respect to H if some finite sequence would get her to become certain about H —in other words, if for some finite sequence s , $P(H|[s])$ equals 0 or 1.

Say that an investigator with prior P is **extremely open-minded** with respect to H if for every $\epsilon > 0$ and for every finite batch of evidence, there is a finite extension of that evidence that would lead her to assign probability less than ϵ to H , and also a finite extension of it that would lead her to assign probability greater than $1 - \epsilon$ to H .

Claim: Suppose that an investigator has a countably additive prior P , and let H be any hypothesis in the domain of P . Then the investigator is either closed-minded with respect to H or extremely open-minded with respect to it.

Proof: Assume that the investigator is not closed-minded with respect to H . We will show that she is extremely open-minded with respect to it.

Take any finite sequence s . Note that $0 < P(H|[s]) < 1$, since the investigator is not closed-minded about H . Let $P'(\cdot) = P(\cdot|[s])$ be the result of conditionalizing P on $[s]$. For any number v , let M_v be the set of sequences x for which $\lim_{i \rightarrow \infty} P'(H|x^i)$ equals v . By the long-run convergence theorem described in §2, P' is certain that P' converges to the truth about H . So $P'((H \cap M_1) \cup (\bar{H} \cap M_0)) = P'(H \cap M_1) + P'(\bar{H} \cap M_0) = 1$. But $1 > P(H|[s]) = P'(H) \geq P'(H \cap M_1)$, so $0 < P'(\bar{H} \cap M_0) \leq P'(M_0)$. So M_0 is nonempty and hence there exists a sequence x such that $\lim_{i \rightarrow \infty} P'(H|x^i) =$

0. So for any $\epsilon > 0$ there exists an n such that $P'(H|x^n) < \epsilon$. It follows that for any $\epsilon > 0$, there is a finite extension s' of s such that $P(H|[s']) < \epsilon$. A similar argument shows that for any $\epsilon > 0$ there is a finite extension s' of s such that $P(H|[s']) > 1 - \epsilon$. So the investigator is extremely open-minded. ■

B Proof of the existence of an open-minded, completely pessimistic finitely additive probability measure:

In the following definitions, v ranges over reals in the unit interval and i ranges over natural numbers.

Let L_v be the set of sequences whose limiting relative frequency of 1s equals v .

Let B_v be *Bernoulli measure with bias v* , the unique countably additive probability measure on S that treats the digits of the sequence as if they were generated by independent tosses of a coin with probability v of landing "1".

Let B_v^i be the countably additive probability measure on S that treats the first i digits of the sequence as if they were generated by independent tosses of a coin with probability v of landing "1", and the remaining digits as if they were generated by independent tosses of a coin with probability $1 - v$ of landing "1".²³

Now define P_0 and P_1 as follows. For any set K of sequences, let $P_0(K) = \text{blim}_{i=1}^{\infty} B_{.9}^i(K)$, and let $P_1(K) = \text{blim}_{i=1}^{\infty} B_{.1}^i(K)$, where *blim* is a Banach limit operator.²⁴

It is easy to check that P_0 and P_1 are finitely additive probability measures. For example, whenever K and K' are disjoint sets of sequences, $P_0(K \cup K') = \text{blim}_{i=1}^{\infty} B_{.9}^i(K \cup K') = \text{blim}_{i=1}^{\infty} (B_{.9}^i(K) + B_{.9}^i(K')) =$

²³In other words, for $v \in (0, 1)$ B_v^i is the unique countably additive measure on S such that for any finite string s , $B_v^i = v^n(1 - v)^z(1 - v)^{n'}v^{z'}$, where n and z are the number of 1s and 0s, respectively, in the first i digits of s and n' and z' are the number of 1s and 0s, respectively, in any remaining digits of s .

²⁴A Banach limit operator is a linear operator on all bounded sequences of reals, defined so that it coincides with the ordinary limit whenever that limit exists. Banach limit operators are not unique, and their existence nonconstructively follows from the axiom of choice (see Rao and Rao (1983, 39–40)). Dependence on a Banach limit is what makes the present conditions on P_0 and P_1 fail to be constructive.

$\text{blim}_{i=1}^{\infty} B_{.9}^i(K) + \text{blim}_{i=1}^{\infty} B_{.9}^i(K') = P_0(K) + P_0(K')$.

(Informally, we can think of P_0 and P_1 in the following way: P_0 treats the sequence as if a large initial segment of it is generated by tosses of a coin biased toward “1”, and the rest by a coin biased toward “0”. P_1 treats the sequence in exactly the opposite way. In each case, the initial segment is expected to be extremely long, in the following sense: for every k , however large, P_0 and P_1 treat the first k digits as if they are part of the initial segment. That is what forces P_0 and P_1 to be merely finitely additive.)

Let $P = (P_0 + P_1)/2$. P is clearly a finitely additive probability function. We will now complete the proof by showing that P is open-minded and completely pessimistic with respect to the hypothesis $L_{.9}$. Note that for any $x \in L_{.9}$ and for any i ,

$$P(L_{.9}|x^i) = \frac{P(L_{.9} \cap x^i)}{P(x^i)} = \frac{(1/2)(P_0(L_{.9} \cap x^i) + P_1(L_{.9} \cap x^i))}{(1/2)(P_0(x^i) + P_1(x^i))} \quad (1)$$

$$= \frac{(1/2)(0 + P_1(x^i))}{(1/2)(P_0(x^i) + P_1(x^i))} \quad (2)$$

$$= \frac{P_1(x^i)}{P_0(x^i) + P_1(x^i)} = \frac{1}{1 + P_0(x^i)/P_1(x^i)} \quad (3)$$

$$= \frac{1}{1 + B_{.9}(x^i)/B_{.1}(x^i)}. \quad (4)$$

(1) holds by definition. (2) holds because $P_1(L_{.9}) = 1$ and $P_0(L_{.9}) = 0$, since for each i , $B_{.1}^i(L_{.9}) = 1$ and $B_{.9}^i(L_{.9}) = 0$ by the strong law of large numbers. (3) is simple algebra. (4) holds because for any binary sequence x and any natural number i , $P_0(x^i) = B_{.9}(x^i)$ and $P_1(x^i) = B_{.1}(x^i)$. To see why, note that $P_0(x^i) = \text{blim}_{j=1}^{\infty} B_{.9}^j(x^i) = \text{blim}_{j=i}^{\infty} B_{.9}^j(x^i) = \text{blim}_{j=i}^{\infty} B_{.9}(x^i) = B_{.9}(x^i)$.

Now consider what happens to (4) as i approaches infinity: The proportion of 1s in the first i digits of x approaches $.9$ (since $x \in L_{.9}$). As a result, $B_{.9}(x^i)/B_{.1}(x^i)$ grows without bound, and hence (4) approaches 0. So when $x \in L_{.9}$, $\lim_{i \rightarrow \infty} P(L_{.9}|x^i) = 0$. A similar argument shows that when $x \in L_{.1}$, $\lim_{i \rightarrow \infty} P(L_{.1}|x^i) = 0$.

It follows that P is open-minded, since for any initial segment of digits, appending a large enough finite block consisting of 90% 1s will force P to assign a probability to $L_{.9}$ that is arbitrarily close to 1, and appending a large enough finite block consisting of 90% 0s will force P to assign a probability to $L_{.9}$ that is arbitrarily close to 0.

It also follows that P is completely pessimistic, since $P(L_{.9} \cup L_{.1}) = 1$, and the above argument shows that P converges to the wrong verdict about $L_{.9}$ for any sequence in $L_{.9} \cup L_{.1}$. ■

References

- Paul Bartha. Countable additivity and the de Finetti lottery. *British Journal for the Philosophy of Science*, 55(2):301–321, 2004.
- Gordon Belot. Bayesian orgulity. *Philosophy of Science*, 80(4):483–503, October 2013.
- Alan F. Chalmers. *What Is This Thing Called Science?* Hackett Publishing Company, Indianapolis, 3rd edition edition, 1999.
- Bruno de Finetti. *Theory of Probability*. Wiley, New York, 1974. Translated by A. Machi and A. Smith.
- L.E. Dubins and L.J. Savage. *How to gamble if you must: inequalities for stochastic processes*. McGraw-Hill series in probability and statistics. McGraw-Hill, 1965.
- John Earman. *Bayes or Bust?* MIT Press, Cambridge, Massachusetts, 1992.
- Kenny Easwaran. Bayesianism II: Applications and criticisms. *Philosophy Compass*, 6(5):321–332, 2011. ISSN 1747-9991. doi: 10.1111/j.1747-9991.2011.00398.x. URL <http://dx.doi.org/10.1111/j.1747-9991.2011.00398.x>.
- Kenny Easwaran. Why countable additivity? *Thought*, 2:53–61, 2013. doi: 10.1002/tht.60.
- W. Edwards, H. Lindman, and L. J. Savage. Bayesian statistical inference for psychological research. *Psychological Review*, 70(3):193–242, 1963.
- P.R. Halmos. *Measure Theory*. Number 9 in Graduate Texts in Mathematics. Springer, 1974.
- James Hawthorne. Bayesian induction is eliminative induction. *Philosophical Topics*, 21(1):99–138, 1993.
- James Hawthorne. Inductive logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. The Stanford Encyclopedia of Philosophy, summer 2014 edition, 2014.
- Colin Howson and Peter Urbach. *Scientific reasoning: The Bayesian approach*. Open Court, third edition edition, 2006.

- Cory Juhl and Kevin T. Kelly. Realism, convergence, and additivity. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1994:181–189, 1994.
- Joseph B. Kadane, Mark J. Schervish, and Teddy Seidenfeld. Statistical implications of finitely additive probability. In *Bayesian inference and decision techniques*, chapter 5, pages 59–76. Elsevier Science Publishers, 1986.
- Kevin T. Kelly. *The Logic of Reliable Inquiry*. Logic and computation in philosophy. Oxford University Press, 1996.
- Isaac Levi. *The enterprise of knowledge*. MIT Press, Cambridge, 1980.
- J.C. Oxtoby. *Measure and Category: A Survey of the Analogies between Topological and Measure Spaces*. Graduate Texts in Mathematics. Springer, 1980.
- K.P.S.B. Rao and B. Rao. *Theory of Charges: A Study of Finitely Additive Measures*. Pure and Applied Mathematics. Elsevier Science, 1983.
- Leonard Savage. *The foundations of statistics*. John Wiley and Sons, 1954.
- M. Schervish and T. Seidenfeld. An approach to consensus and certainty with increasing evidence. *Journal of Statistical Planning and Inference*, 25 (3):401–414, July 1990.
- M. Schervish, T. Seidenfeld, and J. Kadane. The extent of non-conglomerability of finitely additive probabilities. *Z. f. Wahrscheinlichkeitstheorie*, 66:205–226, 1984.
- Teddy Seidenfeld and Mark J. Schervish. A conflict between finite additivity and avoiding Dutch Book. *Philosophy of Science*, 50:398–412, 1983.
- Aron Vallinder. Solomonoff induction: A solution to the problem of the priors? Master’s thesis, Lund University, September 2012.
- Brian Weatherson. Belot on Bayesian orgulity. Manuscript, April 2014. URL <http://brian.weatherson.org/papers.html>.