



Approaching deterministic and probabilistic truth: a unified account

Gustavo Cevolani¹ · Roberto Festa²

Received: 3 September 2020 / Accepted: 1 July 2021
© The Author(s) 2021

Abstract

The basic problem of a theory of truth approximation is defining when a theory is “close to the truth” about some relevant domain. Existing accounts of truthlikeness or verisimilitude address this problem, but are usually limited to the problem of approaching a “deterministic” truth by means of deterministic theories. A general theory of truth approximation, however, should arguably cover also cases where either the relevant theories, or “the truth”, or both, are “probabilistic” in nature. As a step forward in this direction, we first present a general characterization of both deterministic and probabilistic truth approximation; then, we introduce a new account of verisimilitude which provides a simple formal framework to deal with such issue in a unified way. The connections of our account with some other proposals in the literature are also briefly discussed.

Keywords Truthlikeness · Verisimilitude · Truth approximation · Probability · Similarity · Information · (Dis)agreement · Divergence

Contents

1	Introduction	1
2	Deterministic and probabilistic truth approximation	2
3	A probability-based measure of truthlikeness	3
3.1	Truthlikeness and (logical) probability	3.1
3.2	Truthlikeness as agreement with the truth	3.2

This article belongs to the topical collection “Approaching Probabilistic Truths”, edited by Theo Kuipers, Ilkka Niiniluoto, and Gustavo Cevolani.

✉ Gustavo Cevolani
gustavo.cevolani@imtlucca.it
Roberto Festa
festa@units.it

¹ IMT School for Advanced Studies Lucca, Lucca, Italy

² Department of Humanistic Studies, University of Trieste, Trieste, Italy

4	Comparison with other accounts
5	Toward a unified account
5.1	A unified approach to deterministic and probabilistic truth approximation
5.2	Truthlikeness and distance between probabilities
6	Concluding remarks and future work
	Proofs
	References

1 Introduction

Popper introduced the notion of the truthlikeness (or verisimilitude) of a scientific theory or hypothesis in order to make sense of the idea that the goal of inquiry is an increasing approximation to “the whole truth” about the relevant domain. Explicating this intuition is the shared problem of all post-Popperian accounts of verisimilitude, which not only differ from each other on the way to solve it, but also disagree on some crucial features of the notion of truthlikeness. Despite their disagreements, these accounts share an important assumption, i.e., that both the theories and the truth are “deterministic” or “categorical”. This means that both the theory and the truth are commonly construed as descriptions, in some suitable language, of the relevant facts in the domain (or of the relevant “nomic” possibilities characterizing it), and truthlikeness is a matter of closeness or similarity between the corresponding propositions in that language. With few exceptions, no mention is made of other relevant cases, where either the theory or the truth, or both, are instead “probabilistic”: for instance, when theories are construed as probabilistic doxastic states, or the truth is the objective, statistical, or probabilistic distribution of relevant features in the domain.

In this paper, we address this issue and put forward a unified account of both deterministic and probabilistic truth approximation. More precisely, we define a measure of the truthlikeness of both deterministic and probabilistic theories in terms of the probability they assign to the “basic features” of the relevant domain. As we argue, this measure (actually, a continuum of measures) fits well within the so called basic feature (BF) approach to verisimilitude, and shares interesting connections with other accounts proposed in the literature. Moreover, it can be applied quite naturally to all four different possible cases of truth approximation, construed as increasing closeness of a deterministic or probabilistic theory to a deterministic or probabilistic truth.

We proceed as follows. Section 2 outlines the main problem and motivation for the paper, and describes the four cases of truth approximation, mentioning some possible applications and some relevant references for each of them. In Sect. 3, we introduce a new, probability-based measure of truthlikeness, following the idea that a theory is the closer to the truth the more it “agrees” with it on the basic features of the domain; as we argue, this account provides a common, abstract framework for studying both deterministic and probabilistic truth approximation. Section 4 shows how our measure generalizes other existing accounts in the literature, and how deterministic truthlikeness can be construed as a special case of probabilistic truthlikeness. In Sect. 5, we show how the proposed measure provides a common foundation for both deterministic and probabilistic truthlikeness, and hence a unified treatment of deterministic and probabilistic truth approximation; we also discuss its relations with two standard mea-

asures of the distance between probability distributions. Finally, Sect. 6 contains some concluding remarks and directions for future work. The proofs of the main claims in the paper appear in the final appendix.

2 Deterministic and probabilistic truth approximation

From an abstract point of view, one can think of the issue of truth approximation in science as follows. A researcher (or a community of researchers) is interested in investigating some domain of natural or social phenomena (states, situations, facts, systems, and so on). To this purpose, the researcher adopts some conceptual framework, and, in particular, some (more or less formal) language \mathcal{L} . Within \mathcal{L} , one can formulate a number of “theories”; but once \mathcal{L} and the conceptual frame are fixed, one of such theories, denoted t , will represent “the whole truth”, i.e., the most complete and correct characterization of the relevant domain given the expressive resources of \mathcal{L} .

At the beginning of the investigation, t is usually unknown, and often remains so: it represents, so to speak, the final, possibly ideal, target of inquiry. At this point, the issue of truth approximation arises: it makes sense, from both a theoretical and a practical point of view, to ask whether some given theory h “approximates” the target t , i.e., whether it is sufficiently close or similar to it; or, whether h is closer or more similar to t than some alternative theory g . (Fig. 1 schematizes what said so far.) This raises in turn two different issues. The first is the “logical” or “semantic” problem of truth approximation: assuming that t is known, specify when h is actually a good approximation of t , i.e., define a suitable notion of “closeness” or “similarity” between theories and the truth. The second is the “epistemic” or “methodological” problem of truth approximation: given the available evidence (data, observations, experimental results, etc.), specify how one can rationally estimate how well h approximates t , or whether h approximates t better than g .

So far, we have left the notion of “theory” undefined. Depending on the specific scientific discipline or domain, a theory h could be a proper axiomatic system, a (statistical) model, a scientific hypothesis, a numerical interval, a plain proposition or even a “belief state” of some (more or less rational) agent, including artificial ones (like, say, a neural network). However, an account of truth approximation developed at a

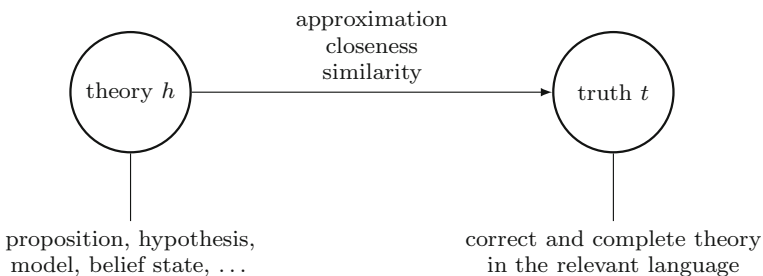


Fig. 1 A schematic representation of the problem of truth approximation

sufficient level of generality should be able to deal with many of the most relevant cases, by uncovering the underlying “logic of truth approximation”. In the philosophy of science, this project was started by Popper (1963), and has been systematically pursued within the post-Popperian research program on “verisimilitude” or “truthlikeness” (for surveys, see Niiniluoto 1998; Oddie 2016). The results of this ongoing exploration are mixed: on the one hand, we have now different accounts of truth approximation available, and have learned much about many crucial aspects of this notion; on the other hand, no such account has gained widespread acceptance in the community, nor has proved able to establish a common ground for further research (see, e.g., Cevolani 2017).

While the multiplicity of truthlikeness theories remains an open problem, here we are concerned with a different issue. As diverse from each other as they are, most of the existing accounts of truth approximation agree on the “deterministic” nature of the truth to be approached. In other words, one usually assumes that the truth t is either the descriptive or factual truth about the relevant domain or the “nomic” truth about what is, e.g., physically or biologically possible in that domain. Most often, also the theories considered as candidate approximations of the truth are thought of in such deterministic terms. This is a relevant limitation, since it excludes the important cases where either the truth, or the theory, or both, are probabilistic in nature.

To illustrate, consider the following toy examples. Suppose that the (unknown) truth about tomorrow’s weather in some location is that it will be hot, rainy, and windy. If Adam says that it is hot and rainy, and Eve says that it is hot and dry, it seems clear that Adam’s “theory” should be judged to be closer to the truth than Eve’s, since she is incorrect on one aspect of the truth (i.e., whether it will rain or not), while he is not. All available accounts of truthlikeness can provide this kind of assessments. However, they are not well equipped to deal with cases like the following.

Suppose that Adam thinks that the probability of rain tomorrow is 80%, while Eve assesses such probability at 30%. If it turns out that the actual weather is rainy, it seems clear that Adam’s estimate was more accurate than Eve’s, in the sense of being closer to the truth. Or, suppose again that Adam believes that it will rain tomorrow, whereas Eve thinks it will not. Moreover, assume that the meteorological record tells that the actual frequency of rainy days in the relevant location is 90%. Now, one would be inclined to say that Adam’s beliefs are closer to the (objective, probabilistic) truth than Eve’s, since it is much more probable than not that it will rain tomorrow. In order to make sense of these intuitive judgments, however, one needs a notion of truthlikeness or verisimilitude which is applicable to both deterministic and probabilistic contexts, a notion that most accounts on the market do not provide.

In connection with the last two examples above, note that the former deals with a case of probabilistic “theories” (beliefs, estimates) approaching a deterministic truth; the latter case, vice versa, is one of deterministic theories approaching a probabilistic (statistical) truth. In general, one can think of four relevant possibilities, as shown in Table 1, that we can briefly describe as follows.

DD: *Deterministic theories approaching a deterministic truth.* Both theory h and the truth t are construed as deterministic, i.e., as either qualitative or quantitative propositions (or sets of propositions, possibly logically closed). The weather

example used above is a case in point: the truth is a description like “hot, rainy, and windy”, and theories are propositions like “hot and rainy”, “hot and dry”, and so on. Another relevant case, ubiquitous in the sciences, is the estimation of the numerical value of some quantity or parameter: here the truth is a point in some relevant space, and theories are represented by (closed) intervals.

- DP: *Deterministic theories approaching a probabilistic truth.* Theory h is deterministic, while the truth t about the domain is probabilistic. This happens, for instance, when h is a “tendency hypothesis”—like “Italians tend to be Catholic” or “Italians love pizza”, or “Birds fly” — and t specifies the precise probability distributions of the relevant traits in the relevant populations (like the proportion of pizza-loving Italians, and so on). Arguably, this kind of hypotheses, and the corresponding notion of truth approximation, play an important role in many social sciences, for instance in sociology or political science.
- PD: *Probabilistic theories approaching a deterministic truth.* Theory h is a probabilistic or statistical theory (model, distribution), while the truth t is deterministic. The standard case may be assessing the truthlikeness of “doxastic” or belief states expressed as probabilistic credences, given an objective but unknown state of the matter. A typical example is medical diagnosis, where the physician estimates the likelihood the patient is ill or not. This is relevant to Bayesian epistemology and other approaches to uncertain reasoning in different fields.
- PP: *Probabilistic theories approaching a probabilistic truth.* Finally, both theory h and the truth t may be construed as either probabilistic or statistical distributions or models. This case includes again assessing the truthlikeness of belief states when the truth is not deterministic, as well as many scientific applications where one needs to estimate the accuracy of different probabilistic predictions with respect to some underlying objective distribution.

In the philosophical literature, the four cases above have attracted very different amounts of attention. As already mentioned, all main approaches within the post-Popperian research program on verisimilitude or truthlikeness deal with DD as their most relevant application: these include so-called consequence approaches by Popper himself, Schurz and Weingartner (2010) and others (cf. Cevolani and Festa 2020), the content approaches of Kuipers (2000, 2019) and Miller (1978), and the similarity approaches defended by Oddie (1986, 2013) and Niiniluoto (1987, 2003). In all such accounts, deterministic truth approximation is straightforwardly construed as a “game of excluding falsity and preserving truth” (Niiniluoto 1999, p. 73), either in terms of balancing the true and false information provided by h (as in content- and consequence-based approaches, including Popper’s original one) or in assessing the overall closeness of h to the truth (as in similarity-based accounts). Such ideas have been applied to a number of relevant scientific issues, including for instance statistical estimation problems (Niiniluoto 1987; Festa 1993) and problems relative to machine learning in AI (Niiniluoto 2005).

As for case DP, it has been virtually ignored, even if Festa (2012, 2007a, b) has emphasized its relevance for a methodological analysis of interesting issues in the social sciences, pointing to so-called prediction logic developed by Hildebrand et al. (1977) as an important case-study in econometric analysis.

Table 1 Deterministic and probabilistic truth approximation: four cases

		Truth	
		Deterministic	Probabilistic
Theory	Deterministic	DD	DP
	Probabilistic	PD	PP

Finally, cases PD and PP are often discussed together, with the implicit understanding that the former may be considered as a special case of the latter (a point also emphasized in the present paper). Here, one can elaborate on a wealth of formal measures studied by mathematicians, statisticians, and scientists working in a variety of fields, from economics to biology and artificial intelligence (cf., e.g., Crupi et al. 2018). A canonical example is the so called Brier score (Brier 1950), which measures the accuracy of (meteorological) probabilistic forecasts of mutually exclusive possible outcomes (like “rain” vs. “no rain”) against the final outcome actually observed (and hence can be construed as an instance of case PD above). Another is the so called Kullback-Leibler divergence (Kullback and Leibler 1951), used for instance by Rosenkrantz (1980) and Niiniluoto (1987) in their analysis of “legisimilitude” (nomic verisimilitude) in terms of the distance of a probabilistic law from a probabilistic truth (case PP above). On the other hand, philosophers like Pettigrew (2016) working on so called accuracy-first epistemology (i.e., on accuracy-based justifications of probabilism as a norm for credence in the line of Bruno de Finetti and others) study the closeness of probabilistic beliefs to deterministic or probabilistic states, thus exploring our cases PD and PP (see Pettigrew (2019) for a survey). Perhaps surprisingly, this latter line of research has developed quite in parallel to that on truthlikeness, even if interesting connections exist and start being explored (Oddie 2017; Schoenfeld 2019).

In this paper, we take a step back and propose to look at the issue of truth approximation from a more abstract and unified perspective. In particular, we start again from the basic ideas underlying truthlikeness theories and propose a simple framework for reasoning about both deterministic and probabilistic truth approximation.

3 A probability-based measure of truthlikeness

Intuitively, a theory (proposition, hypothesis, etc.) h is verisimilar when it is, in some suitably defined sense, “close” or “similar to the truth”. To spell out this intuition in detail, we introduce here a simple formal framework, borrowed from the so called basic feature (BF) approach to truthlikeness (Cevolani and Festa 2020), to be further discussed in the next section. In this and the next section, we focus on the formal properties of the resulting truthlikeness measures, assuming as customary both a deterministic theory and a deterministic truth (case DD from Table 1). In Sect. 5, we shall argue that the proposed measures apply equally well to the remaining cases of truth approximation.

3.1 Truthlikeness and (logical) probability

We work within a finite propositional language \mathcal{L}_n with n atomic propositions. Such atomic propositions, together with their negations, form the $2n$ “basic” or “elementary” propositions of \mathcal{L}_n , also called its “literals”. Intuitively, a basic proposition describes a possible elementary fact in the relevant domain (like “it’s hot” and “it doesn’t rain”); in particular, the n true basic propositions of \mathcal{L}_n , denoted t_1, \dots, t_n , describe the “basic features” of the world. For the moment, we can avoid specifying what these basic features amount to; as we shall see in the following, different interpretations are possible.

Within \mathcal{L}_n , one can express 2^n non-logically-equivalent propositions, including the tautological and the contradictory ones, denoted \top and \perp respectively. A so called constituent (or “state description”) of \mathcal{L}_n is a consistent conjunction of n basic propositions; intuitively, it completely describes a possible state of affairs of the relevant domain (a “possible world”). There are 2^n constituents, which are the logically strongest factual propositions of \mathcal{L}_n . The set $R(h)$ of constituents entailing h (or, equivalently, the class of possible worlds in which h is true) is called the range of h . By definition, each constituent is logically incompatible with any other, and only one of them is true; this is denoted t and is the strongest true statement expressible in \mathcal{L}_n . Thus, t can be construed as “the (whole) truth” in \mathcal{L}_n , i.e., as the complete, true description of the actual world; of course, t entails all true basic statements of \mathcal{L}_n , t_1, \dots, t_n , thus offering a complete and true description of the basic features of the world.

A probability distribution on \mathcal{L}_n assigns a degree of probability to all constituents, and hence to all propositions of \mathcal{L}_n . The “logical” probability measure m assigns the same degree of probability to each constituent; since there are 2^n constituents, each of them has probability $1/2^n$ (Kemeny 1953). Moreover, the probability of h is simply the proportion of constituents entailing h out of the total number of constituents: $m(h) = |R(h)|/2^n$. It follows that if b is a basic proposition, $m(b) = 2^{n-1}/2^n = 1/2$. Assuming that h is consistent (as we shall always assume in the following), the conditional logical probability of g given h is defined as usual, i.e., $m(g|h) = m(h \wedge g)/m(h)$. This means that $m(g|h)$ is the proportion of the cases (i.e., constituents) in which g is true out of the total number of cases in which h is true.

Here, we shall define the truthlikeness of h in terms of the (logical) probability that h assigns to each of the basic features of the world, as follows. Let us consider a consistent set of n basic propositions b_1, \dots, b_n of \mathcal{L}_n . (The set of the atomic propositions of \mathcal{L}_n and the set of the true literals t_1, \dots, t_n entailed by t are two examples of such a set). Given a generic proposition h , we consider its associated “basic vector” $m(b_1|h), \dots, m(b_n|h)$, i.e., the vector of the conditional logical probabilities of each basic proposition b_i given h . Intuitively, the basic vector of h specifies how probable each b_i is, assuming h is true; in other words, it specifies the probability distribution on the chosen basic propositions “from the standpoint” of h . In particular, if h is tautological the corresponding basic vector is “flat”, in the sense that all b_i are assigned probability $m(b_i|\top) = \frac{1}{2}$. At the other extreme, the basic vector of a constituent w is “opinionated”, in the sense that, for any b_i , $m(b_i|w)$ is either 1 or 0.

In general, different statements h and g will have different basic vectors associated with them; i.e., they will assign to each b_i different probabilities. Intuitively, the more different these probabilities are, i.e., the farther from each other are the basic vectors of h and g , the less similar or close will h and g be. This suggests to define similarity (or closeness) between propositions in terms of the “basic distance” between the conditional probabilities they assign to each b_i . Given a single basic proposition b_i , the simplest measure d of the distance between h and g , as evaluated with respect to b_i , is arguably:

$$d^i(h, g) \stackrel{df}{=} |m(b_i|h) - m(b_i|g)|, \quad (1)$$

which varies between 0 and 1. Accordingly, the basic closeness (or similarity) between h and g , again as evaluated with respect to b_i , can be defined as:

$$\begin{aligned} s^i(h, g) &\stackrel{df}{=} 1 - d^i(h, g) \\ &= 1 - |m(b_i|h) - m(b_i|g)|. \end{aligned} \quad (2)$$

It is then natural to define the closeness (or similarity) between h and g as the normalized sum of their basic similarities:

$$sim(h, g) \stackrel{df}{=} \sum_1^n \frac{s^i(h, g)}{n}, \quad (3)$$

Note that $sim(h, g)$ also varies between 0 and 1; it is maximal when h and g exactly agree on the probability value of each basic proposition (i.e., when $d^i(h, g) = 0$ for each b_i) and it is minimal when h and g “maximally disagree” on such values (i.e., when $d^i(h, g) = 1$ for each b_i). For example, one can easily check that the distance between a tautology and a constituent is always $\frac{1}{2}$; or that the maximal closeness between two constituents (disagreeing on just one basic proposition) is $\frac{n-1}{n}$; or that the distance between two (different) basic propositions is always $1/n$.

Finally, we can apply Eq. (3) above to define a measure of the truthlikeness of h in terms of its similarity to the truth t :

$$vs(h) = sim(h, t) = \sum_1^n \frac{s^i(h, t)}{n} \quad (4)$$

The truthlikeness of h is thus construed as the sum of the basic similarities between h and t , i.e., in terms of the distance between their associated basic vectors.

As anticipated, the above definition of truthlikeness is limited to deterministic theories (represented by their logical probability vectors) approaching a deterministic truth. However, nothing prevents to apply it to the case of genuinely probabilistic theories, as far as they can be associated to a vector of “basic probabilities”. This amounts to consider, instead of the logical probability $m(b_i|h)$, a statistical or epistemic probability $p_h(b_i)$ as given by the relevant distribution associated to the probabilistic theory

Table 2 Cases of “qualitative” agreement and disagreement between (the basic logical probability vectors associated to) two propositions h and g

		g	
		$m(b g) > \frac{1}{2}$	$m(b g) \leq \frac{1}{2}$
h	$m(b h) > \frac{1}{2}$	Agreement	Disagreement
	$m(b h) \leq \frac{1}{2}$	Disagreement	Agreement

h under examination. The above definitions can be then applied to define the truthlikeness of probabilistic theories, as we shall better see in Sect. 5. For the moment, we discuss an extension of the truthlikeness measure v_S introduced above, which allows for richer assessments of similarity between theories.

3.2 Truthlikeness as agreement with the truth

Equation (3) defines the similarity between two arbitrary propositions h and g of \mathcal{L}_n in terms of how much h and g agree (or, conversely, disagree) on the (logical) probabilities assigned to each of the basic feature of the world. The closer the probabilities they assign to each basic proposition are, the closer the two propositions are overall. However, quantitative agreement of this kind is not always the most important fact to assess the distance between two probabilistic estimates: in some cognitive contexts, “qualitative” (dis)agreement can well be more relevant.

To clarify this point, let us go back to our previous example, with Adam and Eve trying to estimate the probability of rain tomorrow. We can contrast two different scenarios. In the first, Adam thinks that the probability of rain tomorrow is, say, 90%, while Eve assesses such probability at 60%. In the second, their estimates for the probability of rain are instead 65% for Adam and 35% for Eve. Now, it seems clear that Adam and Eve disagree in their estimates in both scenarios. Moreover, the distance between their estimates is the same in both cases, if evaluated by their plain difference as for measure d above. However, it seems also clear that, in the second situation, their disagreement is deeper than in the former. This is because, if asked “Is rain more probable than not tomorrow?”, their answers would coincide in the first scenario (“yes”) but would be opposite in the second (“yes” for Adam, “no” for Eve). This is what we call “qualitative” disagreement, as opposed to “(merely) quantitative” disagreement (which can be considered a sort of “qualitative agreement”). As we argue, qualitative (dis)agreement can often be more important than mere quantitative (dis)agreement, especially in those contexts where precise numerical predictions or measurements are not available or impractical.¹

To deal with this issue, we introduce the following notions and terminology (see Table 2). Given two propositions h and g and a basic proposition b , we shall say that h and g “(qualitatively) agree” on b iff $m(b|h)$ and $m(b|g)$ are either both strictly higher or both lower than $\frac{1}{2}$; and that h and g “(qualitatively) disagree” on b otherwise, i.e., iff $m(b|h)$ is higher than $\frac{1}{2}$ while $m(b|g)$ is equal to or lower than $\frac{1}{2}$, or the other

¹ This may often happen in many “soft” sciences (cf. Festa 2012). In non-numerical contexts, the notions of qualitative vs. quantitative (dis)agreement just introduced may be linked to the notions of strong vs. weak disagreement as defined in Cevolani (2014).

way round. In short, h and g agree on b if they both deem b as probable (i.e., more probable than not) or improbable (i.e., less probable than not); they disagree if b is probable for one of them and improbable for the other. Of course, strict, quantitative agreement between h and g (to the effect that $m(b|h) = m(b|g)$) is possible only when h and g agree at least qualitatively; but it is possible that h and g agree qualitatively on b while quantitatively disagreeing on its precise probability (as in the first scenario of the example above). In other words, qualitative disagreement implies quantitative disagreement, but not vice versa.

We can now define a measure of the closeness or similarity between h and g , which generalizes definition (3) by taking into account the distinction between qualitative and quantitative (dis)agreement. First, we (re)define the basic distance between h and g , as evaluated with respect to b_i , as follows:

$$s_{\eta}^i(h, g) \stackrel{df}{=} \begin{cases} s^i(h, g) & \text{if } h \text{ and } g \text{ agree on } b_i \\ \eta \times s^i(h, g) & \text{if } h \text{ and } g \text{ disagree on } b_i \end{cases} \quad (5)$$

with parameter η ($0 \leq \eta \leq 1$) balancing the relative importance of qualitative (dis)agreement in assessing similarity. Second, we replace Eq. (3) with the following:

$$sim_{\eta}(h, g) \stackrel{df}{=} \sum_1^n \frac{s_{\eta}^i(h, g)}{n} \quad (6)$$

Note that the above definition is quite general, since it covers different (extreme) cases. Of course, if one takes $\eta = 1$, then $sim_{\eta}(h, g)$ is just the old similarity measure $sim(h, g)$, where similarity assessments are insensitive to qualitative (dis)agreement. At the opposite extreme, if $\eta = 0$, then h and g are similar to some degree only in so far they (qualitatively) agree on some basic propositions; otherwise, even close quantitative agreement on basic propositions on which h and g disagree does not count at all toward their overall similarity. (This is a rather extreme case: if for instance Adam's and Eve's forecasts for rain are, respectively, 51% and 49%, according to measure $sim_0(h, g)$ the closeness between their estimates concerning rain is zero.) For intermediate values of η ($0 < \eta < 1$), quantitative (dis)agreement does count in assessing the similarity between h and g , but less than their qualitative (dis)agreement.

Finally, we can apply Eq. (6) above to define a measure of the truthlikeness of h in terms of its agreement with the truth t :

$$vs_{\eta}(h) = sim_{\eta}(h, t) = \sum_1^n \frac{s_{\eta}^i(h, t)}{n} \quad (7)$$

Note that $vs_{\eta}(h)$ measures the truthlikeness of h in terms of how close the probability assigned by h to each of the true basic propositions t_i is to the "real" one as given by t , independently of whether t is deterministic or probabilistic itself. (In the former case, t assigns only extreme values to the probabilities of t_i , either 1 or 0; in the latter, intermediate values are possible.) Thus, Eq. (7) provides a simple definition of

the truthlikeness of propositional theories as based on the amount of (dis)agreement between a theory and the truth, measured in terms of the distance between their corresponding (logico-)probability vectors. In Sect. 6, we shall better discuss the application of this measure to the problem of deterministic and probabilistic truth approximation; before doing this, however, we point out some interesting relations that this measure shares with other definitions proposed in the literature.

4 Comparison with other accounts

In the foregoing section, we defined a measure of truthlikeness in terms of the (logical) probability that theory h assigns to the basic features of the world. In this section, we show that, in the particular case of deterministic truth approximation, our measure boils down to some well-known measures proposed in the literature. In such sense, our account is sufficiently general to subsume these other accounts in a single, unified framework.

We shall start with the BF account of verisimilitude which, as mentioned before, provides the general background for the definition of truthlikeness as agreement with the truth discussed in Sect. 3. Within this approach, the truthlikeness of theory h depends on a balance of (the amount of) true and false information it provides about the truth (Cevolani et al. 2011, 2013; Cevolani and Festa 2020).² To make sense of this idea, in previous work we introduced the notion of the “partial information” provided by h about the basic features of the world. Here, we use this notion to define a new measure of deterministic truthlikeness, which we will later compare with our older measure.

A proposition h (“fully”) entails another proposition g when g is true in all cases h is, or, equivalently, when $R(h) \subseteq R(g)$, i.e., when the range of h is included in that of g . Instead, h “partially” entails g when g is true in most (but not necessarily all) the cases h is; more precisely, h partially entails g when h is “positively relevant” for g , i.e., when $m(g|h) > m(g)$. Of course, if h (fully) entails g , then h also partially entails g , since $m(g|h)$ takes its maximum value 1; but not vice versa.³

Let us now consider an arbitrary basic proposition b of \mathcal{L}_n . Since $m(b) = 1/2$, h will partially entail b , by definition, iff $m(b|h) > m(b)$ iff $m(b|h) > 1/2$. Intuitively, the amount of information provided by h on b is the greater, the higher $m(b|h)$ raises over $m(b) = 1/2$. Accordingly, we can define the amount of information provided

² This intuition was arguably at the core of Popper’s original definition of verisimilitude (cf. Popper 1963, p. 237), which read roughly as follows: the more true consequences and the less false consequences a theory or proposition h has, the greater its truthlikeness. Popper hoped to show that false theories can be sometimes closer to the truth than other theories, thus defending his realist view of cognitive and scientific progress as a succession of likely false, but increasingly verisimilar, theories. Unfortunately, Popper’s definition failed to accomplish this, and it is now widely regarded as untenable as a general definition of truthlikeness; see Niiniluoto (1998) for an instructive reconstruction of the earlier history of truthlikeness, and Fine (2019) for a recent, qualified defense of Popper’s original account. Despite the failure of such definition, Popper’s central intuition — that the truthlikeness of h balances the amount of true information and of false information entailed by h — can be arguably retained in order to develop an adequate account of verisimilitude (Schurz and Weingartner 2010; Cevolani and Festa 2020).

³ The theory of partial entailment employed here comes from some now classical work by Carnap (1950) and Salmon (1969); for recent developments, see for instance Crupi and Tentori (2013) and Roche (2018)[Sec. 5].

by h on b as a simple linear transformation of the plain difference between these two values:

$$inf_m(h, b) \stackrel{df}{=} (m(b|h) - m(b)) + \frac{1}{2} = m(b|h). \tag{8}$$

Following the intuition underlying the BF approach, h will be close to the truth when h provides much information about the basic features of the world, as described by the true literals t_1, \dots, t_n of \mathcal{L}_n . As a consequence, it is natural to define its truthlikeness as the amount of (partial) information provided by h about such true basic propositions:

$$vs_m(h) \stackrel{df}{=} \sum_1^n \frac{inf_m(h, t_i)}{n} = \sum_1^n \frac{m(t_i|h)}{n} \tag{9}$$

Note that $vs_m(h)$ varies between 1 (if h is the “whole truth” t) and 0 (if h is the “complete falsehood” f , i.e., the conjunction of the negations of all t_i), with the truthlikeness $vs_m(\top) = 1/2$ of a tautology as a middle point. This “partial definition” measure vs_m is interesting because it defines truthlikeness as the plain (normalized) sum of the information provided by h on each of the basic features of the world, measured in turn by the plain logical probability of the corresponding basic proposition t_i given h . Intuitively, this makes much sense and vindicates (at least in part) Popper’s original intuition: h is close to the truth when it provides much information about it.

However, a point is worth noting here. According to definition (8), h provides information about b both when h “supports” or “confirms” b (i.e., when $m(b|h) > m(b)$) and when h “undermines” or “disconfirms” b (i.e., when $m(b|h) < m(b)$), as well as when h is “neutral” for b (i.e., when $m(b|h) = m(b)$). In particular, h provides information about the basic features of the world independently of whether h supports or undermines (or it is neutral for) the true basic propositions t_i . This is as it should be, since “information” is different from “supporting” or “confirming” information. However, one can reasonably argue that, as far as one is interested in defining truthlikeness, supporting information about the truth counts more, toward increasing verisimilitude, than undermining or neutral information. Following this intuition, we define the amount of information provided by h on b as different depending on whether h supports or undermines (or is neutral to) b :

$$inf_\theta(h, b) \stackrel{df}{=} \begin{cases} inf_m(h, b) & \text{if } h \text{ supports } b, \text{ i.e.,} \\ & \text{if } m(b|h) > m(b) \\ \theta \times inf_m(h, b) & \text{if } h \text{ undermines or is neutral to } b, \text{ i.e.,} \\ & \text{if } m(b|h) \leq m(b) \end{cases}$$

with parameter θ ($0 \leq \theta \leq 1$) balancing the relative importance of supporting vs. undermining/neutral information in assessing truthlikeness. Then, we replace Eq. (9) with the following as a (generalized) definition of the truthlikeness of h :

$$vs_\theta(h) \stackrel{df}{=} \sum_1^n \frac{inf_\theta(h, t_i)}{n} \tag{10}$$

Note that vs_{θ} reduces to measure vs_m if $\theta = 1$. At the opposite extreme, if $\theta = 0$, then only supporting information about the truth is relevant when assessing the truthlikeness of h . In all other cases, when $0 < \theta < 1$, both supporting and undermining/neutral information provided by h about the truth is relevant, but the former more than the latter. In other words, the information provided by h about a true basic proposition t_i “fully” increases the truthlikeness of h when t_i is supported by h , but does only “partially” so if h undermines or is neutral to t_i .

The definition just proposed offers a simple and intuitively compelling account of deterministic truthlikeness, i.e., of the closeness of h to a deterministic truth t (“deterministic” meaning, to recall, that the basic vector associated to t is “opinionated”, i.e., all true literals t_i are certainly true, given that, of course, $m(t_i|t) = 1$ for any t_i). As a consequence, it has strict connections with similar definitions of truthlikeness discussed within the BF approach, as well as within other accounts of verisimilitude. To see this, let us start with our previous proposal in Cevolani and Festa (2020).

Following the idea of defining truthlikeness as partial information about the truth, in that paper we defined the amount of information provided by h on b as:

$$inf(h, b) \stackrel{df}{=} 2 \times (m(b|h) - m(b)) = 2 \times \left(m(b|h) - \frac{1}{2} \right),$$

which varies between -1 (if $h \models \neg b$) and 1 (if $h \models b$). Note that this definition above is simply a different normalization of the plain difference between logical probabilities than the one we used in Eq. (8) above. Following Popper, we defined $PB_T(h)$ and $PB_F(h)$, respectively, as the class of basic truths (true basic propositions) and of basic falsehoods (false basic propositions) partially entailed by h . Moreover, $inf_T(h)$ and $inf_F(h)$ were defined, respectively, as follows:

$$inf_T(h) = \frac{1}{n} \times \sum_{b \in PB_T(h)} inf(h, b) \text{ and } inf_F(h) = \frac{1}{n} \times \sum_{b \in PB_F(h)} inf(h, b)$$

i.e., as the normalized amount of information provided about the basic truths (resp. falsehoods) partially entailed by h . Finally, the truthlikeness of h is defined as the difference between the amounts of partial true and false information provided by h about the basic features of the world:

$$vs_{inf}(h) = inf_T(h) - inf_F(h) \tag{11}$$

Note that $vs_{inf}(h)$ varies between 1 (if h is the “whole truth” t) and -1 (if h is the complete falsehood), with the truthlikeness $vs_{inf}(T) = 0$ of a tautology as a natural middle point.

Measure vs_{inf} has a number of interesting properties, detailed in Cevolani and Festa (2020). Interestingly, there we also show that it is essentially identical to another well-known measure of truthlikeness, the so called Tichý–Oddie “average” measure (Oddie 1986, 2013). Oddie’s account (and the so called similarity approach to truthlikeness in general) starts by defining a measure $\lambda(w, t)$ of the likeness or closeness of an arbitrary constituent to the true constituent t . In our framework, one can simply define

$\lambda(w, t)$ as the number of atomic propositions on which w and t agree, divided by n (i.e., in terms of the so called Clifford distance between constituents, see Niiniluoto (1987)). In this way, one immediately obtains that $\lambda(w, t) = 1$ iff w is the truth itself, and that the complete falsehood f is maximally distant from t , since $\lambda(t, f) = 0$. Then, truthlikeness is defined as the average closeness to the truth of the constituents in the range $R(h)$ of h :

$$vs_{av}(h) \stackrel{df}{=} \frac{\sum_{w \models h} \lambda(w, t)}{|R(h)|} \quad (12)$$

where $|R(h)|$ is the number of constituents entailing h . This measure varies between $vs_{av}(t) = 1$ and $vs_{av}(f) = 0$; a tautology has an intermediate degree of truthlikeness $vs_{av}(\top) = \frac{1}{2}$.

While being based on deeply different intuitions about truthlikeness — which is defined by vs_{inf} as the balance of true and false partial information provided by h and by vs_{av} as the average closeness to the actual world of the worlds in which h is true — these two measures turn out to be essentially identical (i.e., ordinally equivalent), as proved in Cevolani and Festa (2020) [see theorem 1 on p. 1637]:

Observation 1 (Cevolani and Festa (2020)) For any h , $vs_{inf}(h) = 2 \times vs_{av}(h) - 1$.

The above result is interesting in its own, and has a number of relevant implications for ongoing debates in the truthlikeness literature, including the one on the classification of different accounts of verisimilitude, as well the discussion of the relationships between truthlikeness and logical strength (for discussion, see Cevolani and Festa 2020). Such debates are relatively immaterial for our present purposes, so we won't go into the details.

What is interesting is instead considering the relations among the four measures presented in this section, i.e. vs_{inf} and vs_{av} on the one hand, and our measures vs_m and vs_θ , on the other hand. In this connection, one can prove that:

Observation 2 For any h , if $\theta = 1$ then $vs_\theta(h) = vs_m(h) = vs_{av}(h) = \frac{1}{2}(vs_{inf}(h) + 1)$.

In words, in the special case where no difference is made between supporting and undermining/neutral information, measure vs_θ is exactly the same as Tichý-Oddie's average measure of truthlikeness. In view of observation 1, this also clarifies that definition (9) is essentially a more direct and compact version than definition (11): vs_{inf} is just a linear transformation of vs_m , and hence the two definitions boil down to essentially the same measure.

Observation 2 shows that the “generalized” partial information measure vs_θ proposed in this section subsumes different other measures of deterministic truthlikeness, including Tichý-Oddie's average measure. More interestingly, measure vs_θ is in turn a special case of the agreement measure vs_η introduced in Sect. 3 in terms of basic distance between probability vectors. Indeed, one can prove that:

Observation 3 If the truth is deterministic, then, for any h , $vs_\eta(h) = vs_\theta(h)$ iff $\eta = \theta$.

It also immediately follows from observation 2 that, if $\eta = \theta = 1$, then all the definitions discussed so far single out exactly the same measure, since in that case $vs_\eta(h) = vs_\theta(h) = vs_m(h) = vs_{av}(h) = vs(h)$.

The above results is interesting because it shows that, within the account proposed here, deterministic truth approximation can be construed as a special case of probabilistic truth approximation. Indeed, the agreement measure vs_η provides a common (logico-)probabilistic treatment of both deterministic and probabilistic theories and the truth; when the latter is deterministic (opinionated), however, the proposed measure boils down to measure $vs_\theta(h)$ of deterministic truthlikeness, which in turns generalizes other similar measures in the literature. In the next, final section, we discuss the implications of the results obtained so far for the general theme of the paper.

5 Toward a unified account

In this paper, we introduced a new measure vs_η of the truthlikeness of h (more precisely, a continuum of such measures), construed as its “agreement” with the truth, as measured by the distance between the logical probability assigned by h to each of the basic features of the world and their actual probability value as given by the truth itself. We also showed how measure vs_η integrates within the BF approach by generalizing previous measures proposed in the literature, including Tichý-Oddie’s average measure. In this final section, we go back to the general problem of analyzing deterministic and probabilistic truth approximation in view of the proposed account.

5.1 A unified approach to deterministic and probabilistic truth approximation

Let us start by emphasizing a couple of interesting features of our proposal.

First, the generalized BF approach as based on measure vs_η allows for an interesting outlook on the logical problem of truthlikeness in general. The reason is that vs_η is based on a measure of inter-theoretic distance, i.e., a measure of the similarity between arbitrary theories (cf. Definition 6). Accordingly, truthlikeness is defined as a special case of theory-theory similarity, i.e., that of theory-truth similarity, where “the truth” is conceived as the most complete, correct theory in the domain (in line with the general idea schematized in Fig. 1). However, the proposed account allows in general for the comparison, in terms of quantitative or qualitative (dis)agreements, of arbitrary theories, not necessarily complete in this sense. This is arguably a welcome extension of the BF account, and an important feature for a theory of truthlikeness in general.

Second, our account provides a dual foundation for deterministic truthlikeness, one in terms of the agreement of a theory with the truth (vs_η), one in terms of the information provided by the theory about the truth (vs_θ). At first sight, these two conceptualizations of truthlikeness are clearly different, but they turn out to be in fact equivalent (in view of observation 3). This leaves open the possibility to employ indifferently one or the other of our measures, choosing the one intuitively more relevant for the application at hand. In particular, the idea of truthlikeness as partial information discussed in Sect. 4 suggests interesting links between the analysis of

verisimilitude and that of (Bayesian) confirmation: in a nutshell, a theory is the closer to the truth the more it confirms each of the basic truths about the domain.

Third, and finally, the proposed account allows for a unified treatment of both deterministic and probabilistic truth approximation, since it provides a common probabilistic foundation for the truthlikeness of both deterministic and probabilistic theories. Indeed, as mentioned while introducing measure vs_η , assessing the similarity between theory h and the truth t boils down to comparing the “logico-probability vectors” associated to h and t . This means that both theories and the truth are formally represented as probability distributions over the basic propositions of \mathcal{L}_n , leaving open the possibility to interpret such probabilities in different ways without changing the relevant definitions. In turn, such flexibility in the interpretation of our continuum of measures allows one to cover all four cases of truth approximation discussed in Sect. 2, as follows.

As far as theory h is concerned, one can give both an objective and a subjective interpretation of its claims on the probability of the basic truths t_1, \dots, t_n of \mathcal{L}_n . The former case has been already discussed in Sect. 3: it amounts to considering the basic (logical probability) vector $m(t_1|h), \dots, m(t_n|h)$ associated to h with respect to each t_i . Then, the similarity between h and t is evaluated as the normalized sum of the basic distances between the two basic vectors associated to h and t , as explained there. However, the same machinery works even if theory h is given a subjective interpretation in terms of a probabilistic credences or degrees of belief. In this case, one models the theory h of a rational agent as an epistemic probability distribution p_h defined over the propositions of \mathcal{L}_n , including its basic propositions. As a result, one can consider the basic (epistemic probability) vector $p_h(t_1), \dots, p_h(t_n)$ associated to h with respect to each t_i . Again, the same definitions of distance and similarity used above are applied to evaluate the truthlikeness of h , here construed as a doxastic state (see the next subsection for a toy example).

As for the truth t , it can also be interpreted in either deterministic or probabilistic terms, even if, of course, it doesn’t make sense to give a subjective interpretation of it. In the former case, the basic vector associated to t will be an “opinionated” (logical) probability distribution, assigning only “extreme” values (i.e., 0 or 1) to the probabilities of the basic propositions; in other words, the truth represents, so to speak, the “certainty vector”, assigning to each of the basic propositions its plain truth-value (i.e., 0 or 1). In the latter case, the basic vector associated to t will be a standard probability distribution assigning also intermediate probability values to the basic truths t_1, \dots, t_n ; such probabilities will represent the objective, statistical frequencies of the corresponding basic propositions.

With such alternative interpretations of the basic vectors associated to h and t in mind, one can see how the four possible cases of truth approximation are dealt with. Case DD has been already discussed in Sect. 3: it amounts to representing both theory h and the truth t in terms of their associated logical probability vectors. As for case DP, the truth is instead represented as the “true” statistical probability distribution over the basic propositions, representing for instance the relative frequency of the different kinds of individuals in the domain. A deterministic theory like “All Italians love pizza” (assigning logical probability 1 to the corresponding kind of individual) will then be the closer to the truth the highest the actual frequency of pizza-loving

Italians in the relevant population. Finally, for both cases PD and PP, h is interpreted as a basic epistemic probability vector associated to a doxastic state of some rational agent. The only difference is that such credences may be construed as approaching a deterministic (opinionated) truth (as in PD) or a probabilistic (statistical) truth (as in PP).

In any case, measure vs_η provides a straightforward measure of the closeness of both deterministic and probabilistic theories to the truth, interpreted in turn both in a deterministic or probabilistic way. In this sense, the proposed account provides a unified treatment of both deterministic and probabilistic truth approximation.

5.2 Truthlikeness and distance between probabilities

As far as the comparison between probabilistic theories is concerned (case PP above), it is interesting to study the relation between the measure of truthlikeness proposed here and some of the many measures of distance (and, implicitly, of closeness) between probability distributions proposed by mathematicians in various contexts (we thank an anonymous reviewer for prompting us doing this). As mentioned in Sect. 2, there is a plethora of such measures, the Brier score (Brier 1950) and the directed divergence (also known as “relative entropy”) proposed by Kullback and Leibler (1951) being just two prominent examples. Let us focus on these two measures, in order to compare them with our account.

We denote w_1, \dots, w_q (with $q = 2^n$) the (mutually exclusive and jointly exhaustive) constituents or possible worlds of \mathcal{L}_n and p_h , and p_g two arbitrary probability distributions defined on them (intuitively associated to probabilistic theories h and g , respectively). Then we can define the Brier score as follows:

$$B(p_h, p_g) \stackrel{df}{=} \sum_{w_i} \frac{(p_h(w_i) - p_g(w_i))^2}{q} \tag{13}$$

Note that B is a distance measure between probability distributions, i.e., the smaller $B(p_h, p_g)$ the closer p_h and p_g .

As for the Kullback-Leibler divergence, it is also a measure of the distance between distributions, defined as follows (assuming the logarithm having base 2):

$$KL(p_h, p_g) \stackrel{df}{=} \sum_{w_i} p_h(w_i) \times \log \frac{p_h(w_i)}{p_g(w_i)} \tag{14}$$

In words, KL is the expected logarithm distance between p_h and p_g , with respect to p_h .

One should note that both B and KL (as well as other measures of this kind) are defined between “complete” probability distributions, relative to all the constituents of \mathcal{L}_n . On the contrary, our truthlikeness measure vs_η is defined in terms of the “basic distance” d (see Eq. 1) relative to the basic statements of \mathcal{L}_n only. Of course, every “complete” distribution also entails a “basic” distribution for the probabilities of basic propositions (while the converse is not true). Thus, it is interesting to ask whether the

Table 3 A toy example: three probability distributions on the four constituents of the hot-rainy language, with the resulting probabilities for the basic propositions “hot” and “rainy”

Hot	Rainy		p_t	p_h	p_g
1	1	Hot & rainy	0.4	0.2	0.2
1	0	Hot & not rainy	0.3	0.1	0.1
0	1	Not hot & rainy	0.2	0.5	0.4
0	0	Not hot & not rainy	0.1	0.2	0.3
		Hot	0.7	0.3	0.3
		Rainy	0.6	0.7	0.6

assessment of the distance between two probabilistic theories h and g , as provided by $sim_\eta(h, g)$, agrees with the one provided by measures like B or KL . As the following example shows, no univocal answer can be given to this question.

Let us consider the simplest possible case, where \mathcal{L}_2 contains only two atomic propositions, say “hot” and “rainy”; accordingly, there are only four constituents or possible worlds (and only four basic propositions). We consider three different probability distributions on such constituents, as shown in Table 3, corresponding to two different probabilistic theories h and g , and one “target” theory t construed as the “true” distribution.

We can then compute, for each of the three measures B , KL , and vs_η , the distance of both h and g from the truth t . With the necessary calculations we obtain:

$$B(p_h, p_t) = 0.045 > 0.04 = B(p_g, p_t)$$

and

$$KL(p_h, p_t) = 0.502 < 0.516 = KL(p_g, p_t).$$

As for vs_η , focusing for simplicity on the special case $\eta = 1$, we have:

$$vs(h) = 0.75 < 0.8 = vs(g)$$

since h and g are equally (in)accurate on the probability of “hot”, but h is less accurate than g on the probability of “rainy”.

Summing up, we have that h is closer to t than g according to the Kullback-Leibler divergence, whereas g is closer to t than h according to both the Brier score and our truthlikeness measure. This opposite assessment by different measures is not surprising, in view of the fact that such measures are in general not ordinally equivalent. This means that, given a target probabilistic theory t (or probability distribution p_t), and two other theories or distributions h and g , assessments of relative closeness to t of h and g can be reversed by employing different measures. This issue of “measure sensitivity” has been repeatedly noted in different contexts, even if it remains often implicit in discussions regarding specific applications.⁴

⁴ For discussion, see for instance Cevolani (2017) for the measure sensitivity of truthlikeness assessments, and Crupi et al. (2018) for a comprehensive survey and comparison of different measures of entropy and information.

A consequence of this fact is that the assessment of closeness between probabilistic theories given by our truthlikeness measure vs_η can both agree and disagree with the assessments based on different measures like the Brier score and the Kullback-Leibler divergence, depending on the specific measure considered. As for the Brier score, however, the agreement is granted, at least for the basic case where $\eta = 1$, by the following result (proved in the Appendix):

Observation 4 If $\eta = 1$, then, for any h and g , $vs(h) > vs(g)$ iff $B(p_h, p_t) < B(p_g, p_t)$.

In other words, our basic measure vs and the Brier score B are ordinally equivalent, in the sense that they deliver the same ordering of truthlikeness among probabilistic theories. This result is interesting for at least two reasons. First, because it shows how, by generalizing the basic feature account of truthlikeness as proposed in Sect. 3, one arrives at defining a verisimilitude measure that is equivalent to a well-known scoring rule originally meant to measure the accuracy of probabilistic predictions. In doing so, this result highlights interesting connections between the foundations of two important research programs which proceeded virtually in parallel. Similarly, given the central role that the Brier score plays in recent work on arguments for probabilism based on the notion of epistemic utility (Pettigrew 2019), our result also sheds new light on recent work bridging the gap between this latter approach and the theory of truthlikeness (Oddie 2017; Schoenfield 2019). Second, Observation 4, along with the example considered above (see Table 3), shows that our proposed measure may well disagree with other approaches based on the Kullback-Leibler divergence (or on measures ordinally equivalent to it). Since one such approach, first proposed by Rosenkrantz (1980) and further developed by Niiniluoto (1987), provides arguably the most advanced account of the truthlikeness for probabilistic theories in the literature, further work is needed to better understand and characterize the disagreement between the two approaches. While we have to leave this exploration for the future, we suggest below a possible way to tackle the issue.

6 Concluding remarks and future work

Before concluding, let us briefly discuss two issues that need to be addressed in order to improve on our present account, and that we have to leave as prospects for future research.

First, what we discussed in the present paper is mainly an abstract formal framework, based on a simple propositional framework, for reasoning about the foundations of deterministic and probabilistic truth approximation. In order to bring our approach closer to potential applications, further work is needed. As an example in this direction, our account can be straightforwardly generalized to monadic predicate logic with both qualitative and quantitative predicates, within which one can reconstruct a number of scientific theories. Moreover, one should consider different concrete instances of our framework, starting from specifying the precise interpretation for the basic propositions of \mathcal{L}_η : e.g., as kinds of individuals, as nomic possibilities, as credences, as frequencies, etc. This would also allow for comparisons with other existing approaches

in the literature, some of which were already mentioned at the end of Sect. 2: in particular, with the analysis of the verisimilitude of tendency hypotheses (Festa 2012) and with ongoing work on the accuracy of probabilistic belief states (Pettigrew 2019).

Second, further work is needed in order to fully understand the relations between the approach to probabilistic truth approximation proposed here and the more standard one based on measures of distance between probability distributions, like the Brier score or the Kullback-Leibler divergence. As a step forward in this direction, we proved in Sect. 5 that our truthlikeness measure agrees with the Brier score on the ordering of probabilistic theories according to their relative closeness to a target distribution, while it may disagree with the Kullback-Leibler measure. In view of the plethora of measures of this kind, however, a more systematic approach is needed to disentangle their relations and their implications for the issue of truth approximation. In this connection, one may build on some more or less recent work on so-called convex or conjunctive propositions, i.e., conjunctions of basic propositions, each describing a basic feature of the domain. As argued in that literature (cfr., e.g., Cevolani and Festa 2020; Cevolani 2020), assessments of the truthlikeness of conjunctive propositions can provide a “core ordering” of truthlikeness that all possible verisimilitude measures should respect. Similarly, one may suggest that respecting the intuitive assessment of the truthlikeness of probabilistic theories provided by our ‘basic’ measure *vs* may be an adequacy condition for distances between probability distributions, separating those that (like the Brier score) agree on that ordering from those (like the Kullback-Leibler divergence) that don’t.

Third, and finally, in this paper we only discussed the logical problem of truth approximation, leaving aside the epistemic problem. This amounts to specify how one can rationally estimate the truthlikeness of theory *h*, when the truth *t* is unknown but some relevant evidence *e* is available. An answer to this problem can be given as follows: assuming that an epistemic probability distribution is given, representing the rational degrees of belief of an agent in the propositions of \mathcal{L}_n , define a measure of the expected truthlikeness as the average truthlikeness of *h* in each possible world, weighted by the corresponding probability on *e* (cf. Niiniluoto 1987). Such a measure is easily definable in the case DD of deterministic truth approximation; however, it is not so clear how to apply it to the other three cases, where either a statistical or epistemic probability distribution is involved in the very representation of *h* and *t*, as explained above. In other words, it is not so clear what the implications are, for the proposed measures of truthlikeness, of the interplay between the probabilities defining the basic vectors associated to *h* and *t*, and those representing the rational degrees of belief of the agent estimating their relative similarity. We leave the exploration of this open problem for future research.

Acknowledgements We would like to thank Theo Kuipers and Ilkka Niiniluoto for detailed comments on a previous draft; Vincenzo Crupi and Luca Tambolo for useful discussions on the topics of this paper; and two anonymous reviewers for detailed critical comments on the manuscript. Gustavo Cevolani acknowledges financial support from the Italian Ministry of Education, Universities and Research (MIUR) through the grant n. 201743F9YE (PRIN 2017 project “From models to decisions”).

Funding Open access funding provided by Scuola IMT Alti Studi Lucca within the CRUI-CARE Agreement.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Proofs

In the following, we prove the two main claims from Sect. 4, i.e., observations 2 and 3, as well as observation 4 in Sect. 5.

Proof of Observation 2 Observation 2 amounts to the claim that, for any h , if $\theta = 1$ then $vs_\theta(h) = vs_m(h) = vs_{av}(h) = \frac{1}{2}(vs_{inf}(h) + 1)$. The first equality immediately follows from the definition of vs_θ in Eq. (10); and the third one is just a reformulation of observation 1. It remains to be proved that $vs_m(h) = vs_{av}(h)$, i.e., the partial information measure defined in Eq. (9) is equivalent to Tichý-Oddie’s average measure of truthlikeness. In fact, this follows immediately from the proof of observation 1 as given by Cevolani and Festa (2020)[see eq. 2 on p. 1644], where we show that (adapting the notation to the one of the present paper):

$$vs_{av}(h) = \sum_1^n \frac{m(t_i|h)}{n}$$

By comparing the above equation with Eq. 9 in the present paper, one immediately see that $vs_m(h)$ is identical, by definition, to $vs_{av}(h)$. □

Proof of Observation 3 As for Observation 3, it says that, if the truth is deterministic, then, for any h , $vs_\eta(h) = vs_\theta(h)$ iff $\eta = \theta$, i.e., that, in such case, the agreement measure defined in (7) is essentially the same as the generalized partial information measure of truthlikeness defined in (10). To see why, let’s first note that, if the truth is deterministic, then $m(t_i|t) = 1$ for any true literal t_i . In then follows from definition (1) that, since $0 \leq m(t_i|h) \leq 1$:

$$d(h_i, t_i) = |m(t_i|h) - m(t_i|t)| = 1 - m(t_i|h)$$

and hence that:

$$s(h_i, t_i) = 1 - d(h_i, t_i) = m(t_i|h).$$

It also follows, by definition, that the basic distance between h and a true literal reduces to:

$$s_\eta(h_i, t_i) = \begin{cases} m(t_i|h) & \text{if } h \text{ and } g \text{ agree on } b_i \\ \eta \times m(t_i|h) & \text{if } h \text{ and } g \text{ disagree on } b_i \end{cases}$$

The above equality should be compared with the following one, which is immediately obtained by the definition of inf_{θ} :

$$inf_{\theta}(h, t_i) = \begin{cases} m(t_i|h) & \text{if } h \text{ supports } t_i, \text{ i.e.,} \\ & \text{if } m(t_i|h) > m(t_i) \\ \theta \times m(t_i|h) & \text{if } h \text{ undermines or is neutral to } t_i, \text{ i.e.,} \\ & \text{if } m(t_i|h) \leq m(t_i) \end{cases}$$

By recalling the definition of agreement and disagreement from Table 2, one can see that, in the case of a deterministic truth, the notions of agreement and support, on the one hand, and of disagreement and undermining/neutrality, on the other hand, collapse on each other. In fact, since $m(t_i|t) = 1 > m(t_i)$:

- if $m(t_i|h) > m(t_i)$, then h supports t_i and h agrees with t on t_i ; moreover, $s_{\eta}(h_i, t_i) = inf_{\theta}(h, t_i) = m(t_i|h)$;
- if $m(t_i|h) \leq m(t_i)$, then h undermines or is neutral to t_i and h disagrees with t on t_i ; moreover, $s_{\eta}(h_i, t_i) = \eta m(t_i|h)$ and $inf_{\theta}(h, t_i) = \theta m(t_i|h)$.

Thus, $s_{\eta}(h_i, t_i) = inf_{\theta}(h, t_i)$ iff $\eta = \theta$. Now, it follows from definitions (7) and (10) that:

$$vs_{\eta}(h) = \sum_1^k \frac{m(t_i|h)}{n} + \sum_{k+1}^n \frac{\eta m(t_i|h)}{n}$$

and

$$vs_{\theta}(h) = \sum_1^k \frac{m(t_i|h)}{n} + \sum_{k+1}^n \frac{\theta m(t_i|h)}{n}$$

where k is the number of true literals that h supports or, equivalently, agrees upon with t ; in other words, the two measure differ only on the weights they give to the other $n - k$ addends. It follows that $vs_{\eta}(h) = vs_{\theta}(h)$ iff $\eta = \theta$. □

Proof of Observation 4 Observation 4 says that our basic truthlikeness measure defined in Eq. 4 and the Brier score defined in Eq. 13 are ordinally equivalent, meaning that for any two probabilistic theories h and g , with associated probability distributions p_h and p_g on the constituents of \mathcal{L}_n , one has that $vs(h) > vs(g)$ iff $B(p_h, p_i) < B(p_g, p_i)$.

This can be easily seen by noting that both vs and B are defined in terms of the (absolute value of) the plain difference between probabilities (cfr. Eqs. 1 and 13). Now consider two probability distributions p_h and p_g on the constituents w_i of \mathcal{L}_n and let be

$$d(p_h(w), p_g(w)) \stackrel{df}{=} |p_h(w) - p_g(w)|$$

their distance as calculated on a single arbitrary constituent w . The Brier score then defines the distance between p_h and p_g as follows:

$$B(p_h, p_g) = \sum_{w_i} \frac{d(p_h(w_i), p_g(w_i))^2}{2^n},$$

where 2^n is the total number of constituents. Instead, *sim* (on which *vs* is based, see Eqs. 4 and 3) is defined in terms of a distance measure Δ as follows:

$$sim(h, g) = 1 - \Delta(h, g),$$

with

$$\Delta(h, g) \stackrel{df}{=} \sum_{b_i} \frac{d(p_h(b_i), p_g(b_i))}{n}$$

where the b_i are the basic propositions of \mathcal{L}_n and n is their total number. Now, for any probability distribution p_h :

$$p_h(b_i) = \sum_{w_j \models b_i} p_h(w_j);$$

i.e., the probability of a basic proposition is just the sum of the probabilities of the constituents entailing it. It follows that:

$$\begin{aligned} d(p_h(b_i), p_g(b_i)) &= d\left(\sum_{w_j \models b_i} p_h(w_j), \sum_{w_j \models b_i} p_g(w_j)\right) \\ &= \left| \sum_{w_j \models b_i} p_h(w_j) - \sum_{w_j \models b_i} p_g(w_j) \right| \\ &= \sum_{w_j \models b_i} |p_h(w_j) - p_g(w_j)| \\ &= \sum_{w_j \models b_i} d(p_h(w_j), p_g(w_j)) \end{aligned}$$

and thus:

$$\Delta(h, g) = \sum_{b_i} \sum_{w_j \models b_i} \frac{d(p_h(w_j), p_g(w_j))}{n}.$$

It is now clear that B and Δ are both increasing functions of the same quantity $d(p_h(w_j), p_g(w_j))$, and hence vary together. As a consequence, B and *vs* agree on the ordering of arbitrary probabilistic theories h and g . □

References

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1), 1–3.

- Carnap, R. (1950). *Logical foundations of probability*. Chicago: University of Chicago Press.
- Cevolani, G. (2014). Truth approximation, belief merging, and peer disagreement. *Synthese*, *191*(11), 2383–2401.
- Cevolani, G. (2017). Truthlikeness and the problem of measure sensitivity. In M. Massimi, J.-W. Romeijn, & G. Schurz (Eds.), *EPSA15 selected papers* (pp. 257–271). Berlin: Springer.
- Cevolani, G. (2020). Approaching truth in conceptual spaces. *Erkenntnis*, *6*(85), 1485–1500.
- Cevolani, G., & Festa, R. (2020). A partial consequence account of truthlikeness. *Synthese*, *197*, 1627–1646.
- Cevolani, G., Crupi, V., & Festa, R. (2011). Verisimilitude and belief change for conjunctive theories. *Erkenntnis*, *75*(2), 183–202.
- Cevolani, G., Festa, R., & Kuipers, T. A. F. (2013). Verisimilitude and belief change for nomic conjunctive theories. *Synthese*, *190*(16), 3307–3324.
- Crupi, V., & Tentori, K. (2013). Confirmation as partial entailment: A representation theorem in inductive logic. *Journal of Applied Logic*, *11*(4), 364–372.
- Crupi, V., Nelson, J., Meder, B., Cevolani, G., & Tentori, K. (2018). Generalized information theory meets human cognition: Introducing a unified framework to model uncertainty and information search. *Cognitive Science*, *42*, 1410–1456.
- Festa, R. (1993). *Optimum inductive methods*. Netherlands: Springer.
- Festa, R. (2007a). Verisimilitude, qualitative theories, and statistical inferences. In S. Pihlström, P. Raatikainen, & M. Sintonen (Eds.), *Approaching truth: Essays in honour of Ilkka Niiniluoto* (pp. 143–178). London: College Publications.
- Festa, R. (2007b). Verisimilitude, cross classification, and prediction logic. Approaching the statistical truth by falsified qualitative theories. *Mind and Society*, *6*, 37–62.
- Festa, R. (2012). *On the verisimilitude of tendency hypotheses*. In D. Dieks et al. (Eds.), *Probabilities, Laws, and Structures* (pp. 43–55). Netherlands: Springer.
- Fine, K. (2019). Some remarks on Popper's qualitative criterion of verisimilitude. *Erkenntnis*, <https://doi.org/10.1007/s10670-019-00192-5>.
- Hiltebrand, D. K., Laing, J. D., & Rosenthal, H. (1977). *Prediction analysis of cross classifications*. New York: Wiley.
- Kemeny, J. G. (1953). A logical measure function. *The Journal of Symbolic Logic*, *18*(4), 289–308.
- Kuipers, T. A. F. (2000). *From Instrumentalism to Constructive Realism*. Dordrecht: Kluwer Academic Publishers.
- Kuipers, T. A. F. (2019). *Nomic truth approximation revisited*. Berlin: Springer.
- Kullback, S., & Leibler, R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, *22*(1), 79–86.
- Miller, D. (1978). On distance from the truth as a true distance. In J. Hintikka, I. Niiniluoto, & E. Saarinen (Eds.), *Essays on mathematical and philosophical logic* (pp. 415–435). Dordrecht: Kluwer.
- Niiniluoto, I. (1987). *Truthlikeness*. Dordrecht: Reidel.
- Niiniluoto, I. (1998). Verisimilitude: the third period. *The British Journal for the Philosophy of Science*, *49*(1), 1–29.
- Niiniluoto, I. (1999). *Critical scientific realism*. Oxford: Oxford University Press.
- Niiniluoto, I. (2003). Content and likeness definitions of truthlikeness. In J. Hintikka, T. Czarnecki, K. Kijania-Placek, A. Rojszczak, & T. Placek (Eds.), *Philosophy and logic: In search of the polish tradition* (pp. 27–35). Dordrecht: Kluwer Academic Publishers.
- Niiniluoto, I. (2005). Inductive logic, verisimilitude, and machine learning methodology and philosophy of science. In P. Hájek, L. Valdés-Villanueva, & D. Westerståhl (Eds.), *Logic* (pp. 295–314). London: College Publications.
- Oddie, G. (1986). *Likeness to truth*. Dordrecht: Reidel.
- Oddie, G. (2013). The content, consequence and likeness approaches to verisimilitude: compatibility, trivialization, and underdetermination. *Synthese*, *190*(9), 1647–1687.
- Oddie, G. (2016). Truthlikeness. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Winter 2016 edition. <https://plato.stanford.edu/entries/truthlikeness/>
- Oddie, G. (2017). What accuracy could not be. *The British Journal for the Philosophy of Science*, *70*(2), 551–580.
- Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Pettigrew, R. (2019). Epistemic utility arguments for probabilism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Winter 2019 edition. <https://plato.stanford.edu/entries/epistemic-utility/>

- Popper, K. R. (1963). *Conjectures and refutations: The growth of scientific knowledge* (3rd ed.). London: Routledge and Kegan Paul.
- Roche, W. (2018). Is there a place in Bayesian confirmation theory for the reverse Matthew effect? *Synthese*, 195, 1631–1648.
- Rosenkrantz, R. D. (1980). Measuring truthlikeness. *Synthese*, 45(3), 463–487.
- Salmon, W. C. (1969). Partial entailment as a basis for inductive logic. In N. Rescher (Ed.), *Essays in honor of Carl G. Hempel* (pp. 47–82). Dordrecht: Reidel.
- Schoenfield, M. (2019). Accuracy and verisimilitude: The good, the bad, and the ugly. *The British Journal for the Philosophy of Science.*, <https://doi.org/10.1093/bjps/axz032>.
- Schurz, G., & Weingartner, P. (2010). Zwart and Franssen's impossibility theorem holds for possible-world-accounts but not for consequence-accounts to verisimilitude. *Synthese*, 172, 415–436.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.