Chapter 21
Truthlikeness and the Problem of Measure Sensitivity

Gustavo Cevolani

Abstract The so-called problem of measure sensitivity concerns the use of formal models in philosophical argumentation: as it turns out, the soundness of many arguments is critically sensitive to the choice of the specific models employed. In this paper, I study how this issue affects the theory of truthlikeness (or verisimilitude) and its applications. As an illustration, I focus on the idea of cognitive progress as increasing truthlikeness. In particular, I show that some basic arguments concerning truth approximation through belief change are not invariant across the different truthlikeness measures proposed in the literature.

Keywords Truthlikeness • Verisimilitude • Measure sensitivity • Scientific progress • Belief change • Truth approximation • Formal models

21.1 Introduction

Philosophers routinely use formal models in their analysis of ordinary and scientific reasoning. While the role of model-based argumentation in current philosophical and epistemological analysis is widely acknowledged, a crucial problem with this common practice remains prominent. This is the so-called problem of measure sensitivity: several theoretical arguments, it turns out, are not invariant across different and otherwise plausible models (often measures of some kind, hence the

I would like to thank Vincenzo Crupi, Roberto Festa, Luca Tambolo, and an anonymous reviewer for useful feedback. Financial support from the FIRB project *Structures and Dynamics of Knowledge and Cognition* (Italian Ministry of Scientific Research, Turin unit, D11J12000470001) and from the project *Assessing information models: exploring theories and applications of optimal information search* (University of Turin and Compagnia San Paolo, D16D15000190005) is gratefully acknowledged.

G. Cevolani (✉)
Center for Logic, Language and Cognition and Department of Philosophy and Education, University of Turin, Turin, Italy
e-mail: g.cevolani@gmail.com

M. Massimi et al. (eds.), *EPSA15 Selected Papers*, European Studies in Philosophy of Science 5, DOI 10.1007/978-3-319-53730-6_21

label). More precisely, the soundness of these arguments is critically sensitive to the choice of a particular formal explication.

The issue of measure sensitivity has been recently explored with reference to some central notions in formal epistemology and the philosophy of science. A probably incomplete list includes Bayesian confirmation (Festa 1999; Fitelson 1999; Brössel 2013; Festa and Cevolani 2016), coherence (Schippers 2016), explanatory power (Schupbach and Sprenger 2011; Crupi and Tentori 2012), and accuracy as employed in so called epistemic utility theory (Leitgeb and Pettigrew 2010). The purpose of this paper is to enrich the above list with a new entry: truthlikeness or verisimilitude (I'll use these two terms as synonymous).

Truthlikeness, and related notions like approximate truth, play a crucial role in prominent defenses of scientific realism (e.g., Niiniluoto 2015; Kuipers 2000; Psillos 1999; Chakravartty 2007; Aronson et al. 1995; Weston 1992), in the theory of scientific progress (Niiniluoto 2014; Cevolani and Tambolo 2013), and in the analysis of scientific and ordinary reasoning (e.g., Teller 2001; Cevolani and Crupi 2015; Cevolani 2017; Cevolani and Schurz 2016). It is quite well known that the different measures of truthlikeness proposed in the literature are not even ordinally equivalent (see Oddie 2014, 2013 for recent overviews). This means that, given two propositions h and g, it may well happen that h is assessed as more truthlike than g by some measure, but g is more truthlike than h according to a different measure. However, the implications of such plurality of truthlikeness measures for the theory of truth approximation in general, and for different philosophical arguments making use of this notion, have not been systematically explored so far. In this paper, I offer a first assessment of such implications; due to limits of space, I focus on a small number of measures and on one specific application, leaving a more comprehensive treatment for another occasion.

I proceed as follows. In Sect. 21.2, I present three different measures, which have been proposed within competing approaches to explicate the notion of truthlikeness. I then introduce a toy example which shows how those measures deliver incompatible assessments of truthlikeness already in very simple cases (Sect. 21.3). Having so illustrated the plurality of truthlikeness measures, I explore its consequences for the analysis of cognitive progress as increasing verisimilitude. In Sect. 21.4, I present a couple of arguments concerning the notions of progressive and regressive belief change; I then show, in Sect. 21.5, that they are all crucially sensitive to the specific measure adopted to assess truth approximation. Section 21.6 summarizes the contents of the paper.

21.2 Measuring Truthlikeness

For the sake of simplicity, I'll focus on a finite propositional language $\mathcal{L}_n$ with n atomic propositions (cf., e.g., Zwart 2001). Within $\mathcal{L}_n$, one can express 2^{2^n} logically non-equivalent propositions, including the tautological and the contradictory ones, denoted $\top$ and $\bot$ respectively. Any proposition h of $\mathcal{L}_n$ (with the exception of $\bot$)

is treated as a possible theory or hypothesis about the truth, which is construed as the most informative true statement expressible in the language. The aim of a theory of truthlikeness is to make sense of comparative judgments of the form "hypothesis h is closer to the truth than hypothesis g" (Popper 1963; Niiniluoto 1987; Kuipers 2000; Oddie 2014).

Given two propositions h and g, h is said to be logically stronger than g when h entails g but g doesn't entail h (in symbols: $h \vDash g$ but $g \nvDash h$); thus, $\perp$ is the logically strongest proposition, and $\top$ the weakest one. The strongest contingent propositions are the 2^n so-called constituents or state descriptions of $\mathcal{L}_n$; these are consistent conjunctions of n "basic" propositions, i.e., atomic propositions or their negations. 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 of $\mathcal{L}_n$. Intuitively, a constituent completely describes a possible state of affairs of the relevant domain (a "possible world"); thus, t can be construed as "the (whole) truth" in $\mathcal{L}_n$, i.e., as the complete true description of the actual world. Given an arbitrary proposition h, its range is the set $R(h)$ of constituents entailing h (or, equivalently, the class of possible worlds in which h is true). Of course, h is true if and only if (iff) t is in its range, and false otherwise. The "complete falsehood" is represented by the "worst" constituent f of $\mathcal{L}_n$, which is the conjunction of the negations of all true basic propositions, i.e., of all the conjuncts of t. The following, standard toy example will be useful to fix ideas (cf. Oddie 2014).

Example 1 Suppose that the weather in some given place is described by the statements of a simple language $\mathcal{L}_3$ with only three atomic propositions: that it is cold (c), rainy (r), or windy (w). Thus, there are only 8 possible state of affairs (constituents) the world can be in: the weather may be either cold, rainy, and windy, or cold, rainy and still, and so on. One of these will be the actual one, as described by the true constituent t. In the following, I'll assume that the weather is actually cold, rainy and windy, i.e., that the truth is $t \equiv c \wedge r \wedge w$. All propositions of $\mathcal{L}_3$ are then construed as competing hypotheses about the actual meteorological conditions, and their truthlikeness is assessed in terms of their closeness to t. To mention but two examples, both the complete falsehood $f \equiv \neg c \wedge \neg r \wedge \neg w$ and the negation of the truth $\neg t \equiv \neg c \vee \neg r \vee \neg w$ appear to be quite far from the truth. In fact, the former is true exactly in the world which is the most dissimilar relative to the actual one; and the latter is true in all worlds except the actual one. Whether f or $\neg t$ is closer to the truth, however, will depend on the specific measure adopted to assess their truthlikeness.

Interestingly, the minimalist conceptual framework of Example 1 is already sufficient to compare some of the main accounts of verisimilitude currently on the market. Such accounts are traditionally classified as belonging to two competing approaches to explicating truthlikeness: the content approach and the similarity (or likeness) approach (Zwart 2001; Oddie 2013); I present them in turn below.

21.2.1 The Content Approach

Intuitively, h is close to the truth when h tells many things about the world, and many of those things are (approximately) true. It follows that a true and highly informative proposition will be quite close to the truth. Following this idea, the content approach defines truthlikeness in terms of two factors: the truth value of h, and its information content, defined essentially as the logical strength of h. Thus, *ceteris paribus*, a true proposition will be more verisimilar than a false one, and the greater the amount of information provided by h, the greater its verisimilitude. A truthlikeness measure which satisfies the above desiderata can be defined as follows (Miller 1994). Assuming that $p(h)$ is the probability of h, let $cont(h) = 1 - p(h)$ be a measure of its content. Moreover, let $cont(h \vee g)$ measure the common content of h and g, which is plausible since $h \vee g$ is the strongest consequence of both h and g (cf. Hempel and Oppenheim 1948,171). Finally, let

$$q(h|g) \stackrel{df}{=} \frac{cont(h \vee g)}{cont(g)} = p(\neg g|\neg h)$$

be a measure of the "deductive dependency" of h on g, i.e., of the proportion of the content of h that is entailed by g. According to Miller (1994,pp. 214 ff.), $q(t|h)$ provides a "rough" measure of the truthlikeness of h, i.e., of how much of the whole truth is entailed by h:

$$vs_q(h) \stackrel{df}{=} q(t|h) = p(\neg h|\neg t) \tag{21.1}$$

As one can check, this measure varies between $vs_q(t) = 1$ and $vs_q(\neg t) = 0$. Thus, the whole truth t is (of course) the most verisimilar proposition, while its negation is the least verisimilar one. This is in line with the intuitions behind the content approach, since $\neg t$ is the *weakest* falsehood of $\mathcal{L}_n$, i.e., the only factual proposition which is both false and maximally uninformative. Note that a tautology is also minimally truthlike according to this account: although it is true, it doesn't provide any information at all about the world, hence $vs_q(\top) = vs_q(\neg t) = 0$. On the contrary, the complete falsehood f is fairly informative, and hence its truthlikeness is not minimal; indeed, it may be quite high (for instance, assuming an uniform probability distribution on the possible worlds in Example 1, one can check that $vs_q(f) = \frac{6}{7} \simeq 0.86$). This highlights a general feature of any adequate account of truthlikeness: a false but informative proposition (in the present case, f) can be closer to the truth than a true but uninformative one (like, for instance, $\top$).

21.2.2 The Similarity Approach

Virtually all theorists agree with the basic intuition behind the content approach, that truthlikeness must be positively correlated with both truth and information. According to the similarity or likeness approach, however, any adequate definition of truthlikeness should "take the *likeness* in *truthlikeness* seriously", as Oddie (2014) puts it. The general idea is that h is verisimilar when it is true in those possible worlds which are close to the actual one and false in those which are far from it. Thus, similarity to the truth is literally a matter of closeness between what the theory says and what the truth is.

The starting point of the similarity approach is defining a measure $\lambda(w,t)$ of the likeness or closeness of an arbitrary possible world to the actual world (Oddie 2013,sect. 5). In the literature, it is customary to define $\lambda(w,t)$ as $1-\delta(w,t)$, where δ is a normalized measure of the distance of a constituent w from the true constituent t. In our propositional framework, one can simply define $\delta(w,t)$ as the number of atomic propositions on which w and t disagree, divided by n (this amounts to the so-called Hamming distance between w and t). 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 $\lambda(t,f) = 0$.

The closeness $\lambda(h,t)$ of theory h to the truth is then defined as a function of $\lambda(w,t)$ for all constituents w in the range of h (which describes possible state of affairs compatible with h). For instance, if one only takes into account the constituent of h which is closest to t, one obtain the following "min" measure of the "approximate truth" of h:

$$at(h) \overset{df}{=} \lambda_{min}(h,t) \overset{df}{=} 1 - \min_{w \vDash h} \delta(w,t) \tag{21.2}$$

While measure $at(h)$ underlies, more or less explicitly, many discussions of realism and truth approximation (e.g. Psillos 1999; Aronson et al. 1995; Weston 1992; Teller 2001), it should not be confused with a proper measure of truthlikeness, since the latter denotes closeness to the whole truth, while $at(h)$ only says how close h is to *being true* (i.e., to including the actual world t). The crucial difference is that truthlikeness assessments are sensitive also to the amount of information provided by h about t, while approximate truth is not: as it is easy to check, for all true h, $at(h) = 1$, independently of the informativeness of h. Thus, $at(h)$ provides at best one "ingredient" of an adequate measure of the verisimilitude of h.

One way to define such a measure is to consider the average closeness of all worlds in h to the actual world; this leads to the so called Tichý-Oddie "average" measure of truthlikeness (Oddie 2014):

$$vs_{av}(h) \overset{df}{=} \frac{\sum_{w \vDash h} \lambda(w,t)}{|R(h)|} \tag{21.3}$$

where $|R(h)|$ is the number of constituents entailing h. This measure varies between $vs_{av}(t) = 1$ and $vs_{av}(f) = 0$, with $vs_{av}(\top) = \frac{1}{2}$ providing a sort of natural middle point. Note that, contrary to what happens with the content approach, the weakest falsehood $\neg t$ is less verisimilar than a tautology, but far better than the complete falsehood f. For instance, with reference to Example 1, one can check that $vs_{av}(\neg t) = \frac{3}{7} \simeq 0.43$.

Niiniluoto's favored "min-sum" measure provides another, more sophisticated definition of truthlikeness, explicitly couched as a weighted sum of a truth-factor and an information-factor. The former is provided by the approximate truth of h; the latter is the normalized closeness of all the worlds in h from t, defined as:

$$\lambda_{sum}(h,t) \overset{df}{=} 1 - \frac{\sum_{w \vDash h} \delta(w,t)}{\sum_{w} \delta(w,t)} = \frac{\sum_{w \vDash \neg h} \delta(w,t)}{\sum_{w} \delta(w,t)}. \tag{21.4}$$

The final measure is the following (cf. Niiniluoto 1987,242):

$$\begin{aligned} vs_{ms}(h) &\overset{df}{=} \gamma\lambda_{min}(h,t) + \gamma'\lambda_{sum}(h,t) \\ &= \gamma at(h) + \gamma'\lambda_{sum}(h,t) \end{aligned} \tag{21.5}$$

with $0 < \gamma, \gamma' \leq 1$. Again, vs_{ms} ranges from $vs_{ms}(t) = 1$ to 0. The tautology has maximum approximate truth but null information factor, hence its truthlikeness is $vs_{ms}(\top) = \gamma$. Concerning the proposition with the least degree of truthlikeness, contrary to the accounts considered so far, it can vary depending on the choice of (the ratio of) the two weights λ and λ' for the truth- and the information-factor (see below for a relevant example).

21.3 The Plurality of Truthlikeness Measures

The measures of truthlikeness presented in the foregoing section are just three examples of those discussed in the literature. However, they have been chosen to be fairly representative of the plurality of accounts to truthlikeness currently on the market.[1] Indeed, as the following example shows, those three measures already disagree quite wildly in assessing the truthlikeness of very simple propositions.

Example 2 Suppose that $\mathcal{L}_3$ is the toy weather language from Example 1. Assuming again that the truth is $t \equiv c \wedge r \wedge w$, let us consider two hypotheses about the current weather: h, which says that it is hot and rainy, and g, which agrees on what h says but adds that it is also still; in symbols:

[1]A more comprehensive treatment should at least address the third, so called "consequence" approach to truthlikeness (Schurz 2011b; Oddie 2013; Cevolani et al. 2011), that I can't consider here due to space limits.

$$h \equiv \neg c \wedge r$$
$$g \equiv \neg c \wedge r \wedge \neg w$$

Note that both h and g are false, but g makes one "mistake" more than h; thus, it seems that their truthlikeness should be different. Indeed, this is what all three accounts considered here deliver; interestingly, however, they disagree on the relevant truthlikeness ordering (the proof will follow in a moment):

$$vs_q(h) < vs_q(g)$$
$$vs_{av}(h) > vs_{av}(g)$$
$$vs_{ms}(h) > vs_{ms}(g) \text{ iff } \gamma > \tfrac{1}{4}\gamma'$$

In short, the three measures deliver (partially) diverging truthlikeness orderings for the two hypotheses considered above.

At first sight, the task of comparing two hypotheses like "hot and rainy" and "hot, rainy, and still" in terms of their relative closeness to the truth "cold, rainy, and windy" may seem uncontroversial. However, a problem arises concerning how to balance the two ingredients of truthlikeness, i.e., truth and informative content. In fact, hypothesis g has greater content than hypothesis h (since $g \vDash h$), which is good; but g is more informative than h just because it adds to it a falsehood ($\neg w$), which is bad. Or, equivalently: although h makes one mistake less than g, it eschews that mistake just being "silent", or suspending the judgment, about whether it is windy or not: i.e., h avoids error only at the price of a smaller information content. It follows that the assessment of the relative truthlikeness of h and g will depend on how different measures balance the desire for truth and that for information (cf. Levi 1967).

To illustrate, let us see in details how the three measures considered here deal with the two hypotheses of Example 2. As for the content-based measure vs_q, note that g is strictly stronger than h and hence has greater content. It follows that $p(g) < p(h)$ and that $p(\neg g|\neg t) > p(\neg h|\neg t)$. By Eq. 21.1, this means that $q(t|g) > q(t|h)$, i.e., that g has greater truthlikeness than h. For instance, assuming a uniform probability distribution on the constituents of $\mathcal{L}_3$, vs_q delivers the following assessment: $vs_q(h) = \frac{5}{7} < \frac{6}{7} = vs_q(g)$. Thus, g must be closer to the truth than h according to vs_q and, more generally, according to the content approach to truthlikeness.

Things are different for the two similarity-based measures vs_{av} and vs_{ms} (cf. Fig. 21.1). The former computes truthlikeness as the average closeness to the truth of the worlds in h and g. Note that g is itself a constituent, so that $vs_{av}(g) = \lambda(g, t) = \frac{1}{3}$, since g describes correctly only one aspect of the truth out of three. As for h, it is compatible with two situations: one is g itself, while the other is the world in which it is hot, rainy, and windy, which is fairly close to the truth, since it is mistaken only on the fact that is it cold. In sum, the average closeness to the truth of h is $vs_{av}(h) = \frac{1}{2}(\frac{1}{3} + \frac{2}{3}) = \frac{1}{2}$. It follows that h is more verisimilar than g.

Fig. 21.1 The two hypotheses $h \equiv \neg c \wedge r$ and $g \equiv \neg c \wedge r \wedge \neg w$ from Example 2 in the weather space of Example 1. *Dots* represent possible worlds as described in $\mathcal{L}_3$ (a *bar over a proposition* denotes negation) and *circles* denote increasing distance from the truth $t \equiv c \wedge r \wedge w$

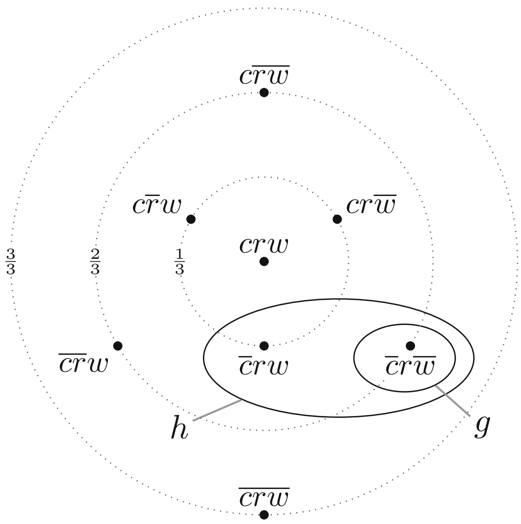

Finally, measure vs_{ms} delivers a more nuanced assessment. As far as the truth-factor is concerned, h has greater approximate truth than g, since h has a world which is closer to t than the only world in g; in fact, $at(h) = \frac{2}{3} > \frac{1}{3} = at(g)$. The situation is however reversed for the information-factor, which rewards g over h for its greater content, as few calculations show: $\lambda_{sum}(h, t) = \frac{3}{4} < \frac{5}{6} = \lambda_{sum}(g, t)$. In sum, one obtains that $vs_{ms}(h) = \frac{2}{3}\gamma + \frac{3}{4}\gamma'$ and $vs_{ms}(g) = \frac{1}{3}\gamma + \frac{5}{6}\gamma'$. Since h has greater approximate truth and less content than g, and g has greater content but less approximate truth, which one is the more verisimilar crucially depends on the ratio between the two weights γ and γ'. It is easy to compute that h is closer to the truth than g just in case $\gamma > \frac{1}{4}\gamma'$, i.e., unless the truth-factor counts much less than the information-factor.

21.4 Progressive and Regressive Belief Change

As one may expect, the plurality of truthlikeness measures illustrated in the last section can impact in various ways on philosophical arguments making use of this notion. This is the problem of measure sensitivity: the soundness of different arguments surrounding truthlikeness may crucially depend on the specific measure adopted to estimate closeness to the truth. As an example of such problem, we shall focus on some recent results concerning belief change and cognitive progress construed as increasing truthlikeness (e.g. Niiniluoto 2011; Cevolani et al. 2011; Schurz 2011b; Kuipers 2011).

The so called AGM theory of belief revision studies (in purely logical, i.e., non-probabilistic, terms) how the beliefs or theories of a rational agent should change in front of incoming evidence or information. In our framework, both the agent's theory h and input information e can be expressed as propositions of $\mathcal{L}_n$. The most basic operations of belief change are known as expansion and contraction, respectively. The former amounts to adding e to h, assuming that e doesn't contradict h; thus, the expansion $h{+}e$ of h by e can be plainly defined as their conjunction:

$$h{+}e \stackrel{df}{=} h \wedge e \tag{21.6}$$

As for contraction, it amounts to removing e from (the consequences of) h in a conservative way, in agreement with a principle of "minimal change" or "informational economy" (Gärdenfors 1988). When $h \vDash e$, this requires that h is weakened and made compatible with $\neg e$, to the effect that the contraction $h{-}e$ of h by e doesn't entail e anymore. A way to achieve this is considering the set $D_h(\neg e)$ of those worlds in the range of $\neg e$ which are closest to h (Niiniluoto 2011):

$$D_h(\neg e) = \{w_i \in R(\neg e) : \lambda_{min}(w_i, h) \geq \lambda_{min}(w_j, h) \text{ for all } w_j \in R(\neg e)\}.$$

One can then define the contraction of h by e as follows:

$$(h{-}e) = \bigvee D_h(\neg e) \vee h \tag{21.7}$$

i.e., as the proposition the range of which is that of h enlarged by $D_h(\neg e)$.[2]

Assuming that h represents a scientific theory and e a piece of empirical evidence, the AGM theory provides a model of theory change, albeit a very abstract and simplified one. To assess its relevance, philosophers of science have recently investigated whether such a model can be applied as a model of cognitive and scientific progress (see, e.g., Niiniluoto 2011). Assuming that the aim of inquiry is approaching the truth about the target domain, one can ask under what conditions belief change leads an agent's theory closer to the truth. As an example, one may expect that the addition of true (or approximately true) inputs e to theory h should result in a new theory $h{+}e$ which is closer to the truth than h was. Or, that removing from h a false consequence e should make the contraction $h{-}e$ more verisimilar than h. Other examples concern *regressive* belief change: adding falsehoods and removing truths from h may be expected to decrease the verisimilitude of the final theory. As we shall see in a moment, none of the above expectations turn out to be fulfilled in general.

Let us take a closer look at possible principles governing progressive (i.e., truthlikeness increasing) and regressive (i.e., truthlikeness decreasing) belief change. The following ones have been considered in the literature as, at least *prima facie*, plausible conditions:

[2]More interesting changes can be constructed as a combination of the two basic operations just introduced: for instance, the revision of h by e may be defined as the contraction of h by $\neg e$ followed by the expansion of the resulting theory by e, i.e., as $(h - \neg e){+}e$.

$$\text{If } e \text{ is true, then } h{+}e \text{ is at least as close to the truth as } h \tag{21.8}$$

$$\text{If } e \text{ is false, then } h{-}e \text{ is at least as close to the truth as } h \tag{21.9}$$

$$\text{If } e \text{ is false, then } h \text{ is at least as close to the truth as } h{+}e \tag{21.10}$$

$$\text{If } e \text{ is true, then } h \text{ is at least as close to the truth as } h{-}e \tag{21.11}$$

The former two conditions say that gaining true information and giving up false information cannot lead one farther from the truth. The latter two say, conversely, that adding falsehoods and removing truths cannot improve the truthlikeness of one's beliefs. As plausible as such principles may appear, arguably none of them can be defended as a general condition of progressive or regressive belief change (cf. Niiniluoto 2011; Cevolani et al. 2011). For instance, as far as condition 21.8 is concerned, the following example shows that sometimes the addition of truths to a false theory may well lead one farther from the truth (cf. Schurz 2011a).

Example 3 Again with reference to the weather Example 1, suppose that an agent accepts the false hypothesis h that when it is it cold, it is both dry and still: $h \equiv c \to \neg r \wedge \neg w$. When the agent receives the true information that it is hot or rainy, the resulting expansion of h by $e \equiv \neg c \vee r$ is the bad theory that at least two of the three basic propositions that it's cold, rainy and windy are false: $h{+}e = (\neg c \wedge \neg r) \vee (\neg c \wedge \neg w) \vee (\neg r \wedge \neg w)$. According to most accounts of truthlikeness, even if e is true, $h{+}e$ may well be farther from the truth than h, thus violating condition 21.8. Indeed, one obtains the following results for the three measures considered here:

$$\begin{aligned} vs_q(h{+}e) &= \tfrac{3}{7} > \tfrac{2}{7} = vs_q(h) \\ vs_{av}(h{+}e) &= \tfrac{1}{4} < \tfrac{1}{3} = vs_{av}(h) \\ vs_{ms}(h{+}e) &= \tfrac{1}{3}\gamma + \tfrac{1}{4}\gamma' < \tfrac{2}{3}\gamma + \tfrac{1}{6}\gamma' = vs_{ms}(h) \text{ iff } \gamma > \tfrac{1}{4}\gamma' \end{aligned}$$

Thus, two out of three measures agree that in the present case expanding by a true input leads to a less verisimilar theory.

Similar counterexamples can be raised against each of the other principles 21.9, 21.10, and 21.11. Accordingly, there is widespread agreement that none of the four principles mentioned earlier provides an acceptable condition on accounts of progressive or regressive belief change.

21.5 Measure Sensitivity and Truthlikeness

How do the three measures considered here fare with respect to the principles concerning progressive and regressive belief change outlined in the previous section? It is not difficult to check that both similarity-based measures vs_{av} and vs_{ms} violates each one of condition 21.8, 21.9, 21.10, and 21.11; this supports the claim that none of them can be taken as a generally sound principle (Niiniluoto 2011). However, the content-based measure vs_q meets both conditions 21.8 and 21.11. The reason is that this measure satisfies the following principle, which makes truthlikeness covary with logical strength[3]:

The Value of Content If $h \vDash g$, then h is at least as close to the truth as g.

Now, the definitions of expansion and contraction (see Eqs. (21.6) and (21.7) immediately imply that $h+e \vDash h$ and that $h \vDash h-e$. It follows that *any* expansion of h will be at least as close to the truth as h itself, and *no* contraction of h can be closer to the truth than it; thus, conditions 21.8 and 21.11 are automatically satisfied.

Searching for more adequate conditions on progressive and regressive belief change, Niiniluoto (2011) has defended the following restricted formulations of conditions 21.8 and 21.11, respectively:

$$\text{If } h \text{ and } e \text{ are both true, then } h+e \text{ is at least as close to the truth as } h \tag{21.12}$$

$$\text{If } h \text{ and } e \text{ are both true, then } h \text{ is at least as close to the truth as } h-e \tag{21.13}$$

The intuition is that if an agent is sure that some source of information is fully reliable, i.e., it provides only true information, then the more information the better. In particular, condition 21.12 vindicates a very basic instance of cognitive progress, according to which the accumulation of truths—e.g., collecting results from a reliable experimental apparatus—improves truth approximation (Niiniluoto 2015). Both vs_q and Niiniluoto's favored measure vs_{ms} guarantee that conditions 21.12 and 21.13 are satisfied. More generally—in view of the observation above that $h+e \vDash h$ and that $h \vDash h-e$—these conditions will be met by all truthlikeness measures satisfying the following principle:

The Value of Content for Truths If h and g are both true, and $h \vDash g$, then h is at least as close to the truth as g.

(Note that this principle follows from, but of course is not implied by, the Value of Content.) Both vs_q and vs_{ms} meet the Value of Content for Truths[4]; this is not the case, however, for vs_{av}, as the following counterexample shows.

Example 4 Suppose that an agent accepts the true hypothesis h that it is cold, or rainy, or windy; thus, the agent correctly excludes the possibility f which is farthest from the truth $t \equiv c \wedge r \wedge w$ (since $h \equiv c \vee r \vee w \equiv \neg f$). Now suppose that the agent receives the true piece of information e, excluding two further incorrect

---

[3]Proof: this follows from the fact that if $h \vDash g$, then $p(h) \leq p(g)$ and hence $p(\neg h|\neg t) \geq p(\neg g|\neg t)$; given definition (21.1) this implies in turn $vs_q(h) \geq vs_q(g)$.

[4]Proof. Since the Value of Content implies the Value of Content for Truths, and vs_q meets the former, it also meets the latter. As for vs_{ms}, if h and g are true, their approximate truth is the same (i.e., 1). Moreover, since $h \vDash g$, $R(h) \subseteq R(g)$; it follows that $\sum_{w \vDash h} \delta(w,t) \leq \sum_{w \vDash g} \delta(w,t)$ and hence that $\lambda_{sum}(h,t) \geq \lambda_{sum}(g,t)$ according to definition 21.4. Then, by definition 21.5, $vs_{ms}(h) \geq vs_{ms}(g)$.

Table 21.1 Conditions on progressive and regressive belief change (as numbered in the text) and measure sensitivity. A checkmark (✓) indicates that the corresponding truthlikeness measure meets the condition; a cross (**x**) that the condition is violated; a double cross (**x**†) that the measure violates the condition by meeting the reverse one (as explained in the text)

	Progressive belief change					
	h	e		vs_q	vs_{av}	vs_{ms}
(21.12)	true	true	$vs(h+e) \geq vs(h)$	✓	**x**	✓
—	true	false	$[h-e \equiv h]$	—	—	—
(21.8)	false	true	$vs(h+e) \geq vs(h)$	✓	**x**	**x**
(21.9)	false	false	$vs(h-e) \geq vs(h)$	**x**†	**x**	**x**
	Regressive belief change					
	h	e		vs_q	vs_{av}	vs_{ms}
(21.13)	true	true	$vs(h-e) \leq vs(h)$	✓	**x**	✓
(21.10)	true	false	$vs(h+e) \leq vs(h)$	**x**†	**x**	**x**
(21.11)	false	true	$vs(h-e) \leq vs(h)$	✓	**x**	**x**
(21.10)	false	false	$vs(h+e) \leq vs(h)$	**x**†	**x**	**x**

possibilities: that it is cold, dry and windy, and that it is cold, rainy, and still (hence $e \equiv \neg c \vee (r \leftrightarrow w)$). Although e is true, it excludes possibilities which are close to the truth; as a consequence, expanding h by e may lower the average truthlikeness of the agent's beliefs. Indeed, one can check that

$$vs_{av}(h+e) = \frac{8}{15} \simeq 0.53 < 0.57 \simeq \frac{4}{7} = vs_{av}(h).$$

Thus, although both h and e are true, $h+e$ is less verisimilar than h. A similar example can be given to show that vs_{av} violates condition 21.13.

Table 21.1 displays the conditions considered so far within a truth table for h and e; the upper part refers to the case of progressive, or truthlikeness increasing, belief change and the lower part to the case of regressive, or truthlikeness decreasing, belief change.[5]

A couple of points are worth noting. On the one hand, there is no condition satisfied by all accounts of truthlikeness; on the other hand, the average measure vs_{av} violates all considered conditions. Moreover, with the exception of two principles, each condition is satisfied by at least one truthlikeness measure, and violated by the others. In particular, even the two special conditions 21.12 and 21.13, dealing with the revision of true theories, are met by both the content-based measure vs_q and by the similarity-based measure vs_{ms} but violated by the average measure vs_{av}.

[5]The second row of the first part of the table, corresponding to the contraction of a true h by a false e, is excluded for logical reasons: since true theories can't have false consequences, the contraction is empty and $h-e$ is the same as h. Note also that condition 21.10, dealing with the expansion of either true or false theories by false inputs, corresponds to two distinct rows in the table.

The only point of agreement seems to concern condition 21.9—for which removing falsehoods from a false theory increases truthlikeness—and condition 21.10—for which adding falsehoods to a theory decreases truthlikeness, which are violated by all the three measures. This agreement, however, is more apparent than real. In fact, even if vs_q violates conditions 21.9–21.10, it does so only by satisfying the following, highly implausible principles, which are, respectively, the reverse versions of those conditions:

$$\text{If } e \text{ is false, then } h \text{ is at least as close to the truth as } h{-}e \tag{21.14}$$

$$\text{If } e \text{ is false, then } h{+}e \text{ is at least as close to the truth as } h \tag{21.15}$$

According to such principles, which immediately follow from the Value of Content, it would be plainly impossible to approach the truth by removing falsehoods from one's current beliefs, while adding falsehoods to a theory would always count as a step toward the truth. Not surprisingly, none of conditions 21.14 and 21.15 have been defended in the literature, nor are satisfied by other truthlikeness measures. Thus, one can rely on vs_q to argue against conditions 21.9–21.10 only at the price of accepting their counterparts 21.14–21.15.

Since it seems intuitively clear that, e.g., adding falsehoods to a theory may well result in a less truthlike theory, the above considerations cast serious doubts on both measure vs_q and on the Value of Content, which is violated by both vs_{av} and vs_{ms}. In fact, the latter principle entails the following condition, according to which truthlikeness increases with logical strength even among falsehoods:

The Value of Content for Falsehoods If h and g are both false, and $h \vDash g$, then h is at least as close to the truth as g

The above principle is rejected by virtually all truthlikeness theorists, the only notable exception being Miller (1994), who accepts it as a sound consequence of his content account of truthlikeness. Indeed, the negation of the Value of Content for Falsehoods is arguably a basic adequacy condition for any account of truth approximation (cf., e.g., Oddie 2013; Niiniluoto 1987; Cevolani 2016).

21.6 Conclusion

In this paper, I explored the issue of measure sensitivity in the context of the theory of truthlikeness. As for Bayesian confirmation and other standard notions in the philosophical debate, it turns out that arguments employing the idea of truth approximation are critically sensitive to the specific measure adopted to formalize it. To illustrate the general problem, I focused on three such measures and on a particular application, i.e., assessing truth approximation via belief change. More specifically, I showed that none of these measures support all arguments which have been proposed for or against some basic principles governing progressive and regressive theory change (cf. Table 21.1). As I argued, this measure sensitivity problem points at a more fundamental one: that different explications of truthlikeness disagree on a couple of highly debated principles concerning the relative truthlikeness of propositions differing in logical strength. Due to space limitations, I could not enter in this debate, nor could I explore in full details the implications of the above results for the ongoing discussion on the notion of truthlikeness (e.g. Oddie 2013). The present discussion should already suffice, however, as an effective illustration of the problem of measure sensitivity as it raises in the context of truthlikeness theories and their applications.

References

Aronson, J., R. Harré, and E.C. Way. 1995. *Realism rescued. How scientific progress is possible*. London: Open Court.

Brössel, P. 2013. The problem of measure sensitivity redux. *Philosophy of Science* 80(3): 378–397.

Cevolani, G. 2016. Carnapian truthlikeness. *Logic Journal of the IGPL* 24(4): 542–556.

Cevolani, G. 2017. Fallibilism, Verisimilitude, and the preface Paradox. *Erkenntnis* 82(1): 169–183.

Cevolani, G., and V. Crupi 2015. Subtleties of naïve reasoning. Probability, confirmation, and verisimilitude in the Linda paradox. In *Epistemology of ordinary knowledge*, ed. M. Bianca and P. Piccari, 211–230. Newcastle upon Tyne: Cambridge Scholars.

Cevolani, G., V. Crupi, and R. Festa. 2011. Verisimilitude and belief change for conjunctive theories. *Erkenntnis* 75(2): 183–202.

Cevolani, G., and G. Schurz. 2016, forthcoming. Probability, approximate truth, and truthlikeness: More ways out of the Preface Paradox. *Australasian Journal of Philosophy*. doi:10.108/00048402.2016.1224265.

Cevolani, G., and L. Tambolo. 2013. Progress as approximation to the truth: A defence of the verisimilitudinarian approach. *Erkenntnis* 78(4): 921–935.

Chakravartty, A. 2007. *A metaphysics for scientific realism: Knowing the unobservable*. Cambridge/New York: Cambridge University Press.

Crupi, V., and K. Tentori 2012. A second look at the logic of explanatory power (with two novel representation theorems). *Philosophy of Science* 79(3): 365–385.

Festa, R., and G. Cevolani. 2016. Unfolding the grammar of Bayesian confirmation: Likelihood and antilikelihood principles. *Philosophy of Science* 84(1): 56–81.

Festa, R. 1999. Bayesian confirmation. In *Experience, reality, and scientific explanation*, ed. M.C. Galavotti and A. Pagnini, 55–87. Dordrecht: Springer.

Fitelson, B. 1999. The plurality of Bayesian measures of confirmation and the problem of measure sensitivity. *Philosophy of Science* 66(3): 378.

Gärdenfors, P. 1988. *Knowledge in flux: Modeling the dynamics of epistemic states*. Cambridge, MA: MIT Press.

Hempel, C.G., and P. Oppenheim. 1948. Studies in the logic of explanation. *Philosophy of Science* 15(2): 135–175.

Kuipers, T.A.F. 2000. *From instrumentalism to constructive realism*. Dordrecht: Kluwer Academic Publishers.

Kuipers, T.A.F. 2011. Basic and refined nomic truth approximation by evidence-guided belief revision in AGM-terms. *Erkenntnis* 75(2): 223–236.

Leitgeb, H., and R. Pettigrew. 2010. An objective justification of Bayesianism I: Measuring inaccuracy. *Philosophy of Science* 77(2): 201–235.

Levi, I. 1967. *Gambling with truth*. Cambridge, MA: MIT Press.

Miller, D. 1994. *Critical rationalism: A restatement and defence*. Chicago: Open Court.
Niiniluoto, I. 1987. *Truthlikeness*. Dordrecht: Reidel.
Niiniluoto, I. 2011. Revising beliefs towards the truth. *Erkenntnis* 75(2): 165–181.
Niiniluoto, I. 2014. Scientific progress as increasing verisimilitude. *Studies in History and Philosophy of Science Part A* 46: 73–77.
Niiniluoto, I. 2015. Optimistic realism about scientific progress. *Synthese*. doi:10.1007/s11229-015-0974-z.
Oddie, G. 2013. The content, consequence and likeness approaches to verisimilitude: Compatibility, trivialization, and underdetermination. *Synthese* 190(9): 1647–1687.
Oddie, G. 2014. Truthlikeness. In *The Stanford encyclopedia of philosophy*, ed. E.N. Zalta. Summer 2014. https://plato.stanford.edu/archives/spr2014/entries/truthlikeness/.
Popper, K.R. 1963. *Conjectures and refutations: The growth of scientific knowledge*, 3rd ed. London: Routledge and Kegan Paul.
Psillos, S. 1999. *Scientific realism: How science tracks truth*. London/New York: Routledge.
Schippers, M. 2016. Bayesian coherentism and the problem of measure sensitivity. *Logic Journal of the IGPL* 24(4): 584–599.
Schupbach, J.N., and J. Sprenger. 2011. The logic of explanatory power. *Philosophy of Science* 78(1): 105–127.
Schurz, G. 2011a. Abductive belief revision in science. In *Belief revision meets philosophy of science*, ed. E.J. Olsson and S. Enqvist, 77–104. Dordrecht: Springer.
Schurz, G. 2011b. Verisimilitude and belief revision. With a focus on the relevant element account. *Erkenntnis* 75(2): 203–221.
Teller, P. 2001. Twilight of the perfect model model. *Erkenntnis* 55(3): 393–415.
Weston, T. 1992. Approximate truth and scientific realism. *Philosophy of Science* 59(1): 53–74.
Zwart, S.D. 2001. *Refined verisimilitude*. Dordrecht: Kluwer Academic Publishers.