Philosophical Studies
https://doi.org/10.1007/s11098-023-02071-w

Vindicating the verifiability criterion

Hannes Leitgeb[1]

Accepted: 27 October 2023

Abstract
The aim of this paper is to argue for a revised and precisified version of the infamous Verifiability Criterion for the meaningfulness of declarative sentences. The argument is based on independently plausible premises concerning probabilistic confirmation and meaning as context-change potential, it is shown to be logically valid, and its ramifications for potential applications of the criterion are being discussed. Although the paper is not historical but systematic, the criterion thus vindicated will resemble the original one(s) in some important ways. At the same time, it will also be more modest insofar as meaningfulness will turn out to be relativized linguistically and probabilistically, and different choices of the linguistic and probabilistic parameters may lead to different verdicts on meaningfulness.

Keywords Verifiability · Confirmation · Meaning · Context-change potential · Pragmatism · Probability

It is hard to think of any other philosophical statement that has received so much criticism over the years as the logical empiricists' infamous Verifiability Criterion of meaning(fulness) for declarative sentences.[1] The formulations of the Criterion changed over the years, but one variant of its original formulation may be summarized as:

> The connection between meaning and confirmation has sometimes been formulated by the thesis that a sentence is meaningful if and only if it is verifi-

[1] Here are some of the most relevant references: Hempel (1950, 1951), Passmore (1967), Soames (2003), and Lycan (2019) are highly critical of the Verifiability Criterion, while Uebel (2019) and Creath (2021) hint at its partial salvageability, and Lutz (2012, 2017), Justus (2014), and Glock (2021) seek to effect partial salvage, though not in the way the present paper will do.

✉ Hannes Leitgeb
hannes.leitgeb@lmu.de

[1] Faculty of Philosophy, Philosophy of Science and Religious Studies, Ludwig-Maximilians-University, Geschwister-Scholl-Platz 1, D-80539 Munich, Germany

Published online: 18 December 2023

able, and that its meaning is the method of its verification (Carnap, 1936a, p. 421).[2]

And yet, in what follows, I will show that there is a strong argument for a reasonably revised and precisified version of the Verifiability Criterion. The underlying aims will not be exegetical but systematic: far from being just a flawed proposal by philosophers of the past, there is a version of the Verifiability Criterion that follows from plausible up-to-date assumptions about confirmation and meaning and which deserves a place in contemporary semantics, pragmatics, epistemology, and philosophy of science.

Reasoning towards that conclusion will proceed in four stages corresponding to Sects. 1–4. Section 1 will take the familiar step of improving the original formulation of the Verifiability Criterion by replacing verifiability by (dis-)confirmability, Sect. 2 will be devoted to the explication of (dis-)confirmability, and Sects. 3 and 4 to that of meaning and meaningfulness. By then, I will have all premises available for stating the argument in Sect. 5, discussing its ramifications in Sect. 6, and summarizing what was achieved and where one might go from there in Sect. 7. Along the way, I will build upon existing work on probabilistic confirmation, meaning as context-change potential, and related probabilistic variants of the Criterion (Skyrms, 1984, 1985; Sober, 1990, 2008). Although the whole enterprise is only inspired by classical sources on the Verifiability Criterion without being committed to them, the criterion thus vindicated will end up resembling the original one(s) in important ways; at the same time, it will be more modest and more likely to be applied for other purposes than originally expected.

1 From verification to (dis-)confirmation

The first step towards a defense of the Verifiability Criterion we do not actually need to take ourselves, as (some of) the logical empiricists already took it early on: to replace verifiability by confirmability. Indeed,

[2] Carnap (1936a) continues by dismissing this "Older Requirement of Verifiability" in favor of a substantial modification. I will mostly rely on Carnap references in this paper. This said, the Verifiability Criterion did not actually play a major role in Carnap's work in semantics, in which he reconstructed meaning in terms of truth conditions, or in his work on probability, in which he gave probabilistic explications of confirmation but not of meaningfulness; see Supplement E, Section 1, of Leitgeb and Carus (2021) for further discussion. However, Carnap's work in philosophy of science—such as, in his most mature formulation, Carnap (1956)—did offer important deductive criteria for the empirical significance of theoretical terms and sentences (see Creath, 1976 for a defense of Carnap, 1956 against potential counterexamples, and see Lutz, 2012; Justus, 2014, and Lutz 2017 for further critical assessments and developments). The basic idea of Carnap (1956) was to determine the empirical significance of sentences from that of terms, and to determine the empirical significance of a term by the existence of a sentence including that term as its only descriptive term, such that the sentence makes a difference for the prediction of an observable state of affairs. The present paper will deal solely with a *probabilistic* explication of meaningfulness of *sentences*, where meaningfulness will consist in *making communicative difference* that shows up probabilistically.

> *no complete verification* is possible but only a process of gradually increasing *confirmation* (Carnap, 1936a, 1936b, p. 425, his emphases).

E.g., it would be practically, nomically, and logically impossible to use a finite conjunction of observation sentences to strictly verify a universal law hypothesis that quantifies over all physical bodies at all places and times (unless one does not regard laws as sentences but as rules of inference, in which case reasoning with negations or disjunctions of laws would become problematic). Hence, strict verifiability seems to be too restrictive as a meaningfulness criterion. In contrast, a gradually increasing confirmation of universal sentences is feasible (see also Carnap, 1963 and Passmore, 1967 on this point). While Carnap (1936a, 1937) still aimed to explicate confirmation deductively, Carnap (1945) already turned to probabilistic explications in which increasing absolute confirmation would ultimately be reconstructed as resulting from iterated incremental confirmation (which will the topic of next section).

Furthermore, mere confirmability should really be extended to the disjunctive *confirmability or disconfirmability*, as a declarative sentence is meaningful just in case its negation is (see e.g. Hempel, 1950, pp. 52–53) and the disconfirmation of a declarative sentence amounts to a confirmation of its negation. So we get the following first improvement of the original criterion, which is in line with early work done by the logical empiricists themselves:

VC A (declarative) sentence is meaningful iff it is confirmable or disconfirmable.

2 The explication of (dis-)confirmation

Clearly, VC just by itself is neither clear nor precise enough. How exactly should we understand '(dis-)confirmable'? (Dis-)confirmability by what? And how should one account for background information that might influence (dis-)confirmation?

Fortunately, we can still follow some of the logical empiricists' footprints here. The next step is to choose a language $\mathcal{L}$ of declarative sentences, such that $\mathcal{L}$ is either a fragment of natural language or a logical reconstruction of a fragment of natural language. We are going to focus on the (dis-)confirmability of members of $\mathcal{L}$ by members of $\mathcal{L}$, where one may want to restrict the (dis-)confirming sentences to some subset $\mathcal{S}_1$ of $\mathcal{L}$ and the (dis-)confirmed sentences to a subset $\mathcal{S}_2$ of $\mathcal{L}$. (Perhaps one is interested just in the confirmability of certain sentences by certain other sentences). Additionally, I will assume the language $\mathcal{L}$ to be closed under negation, conjunction, and disjunction, while leaving open all other questions of logical structure. But, of course, in typical applications of the framework, the language $\mathcal{L}$ will also include quantified sentences, sentences with modal expressions, and more. I will not deal with questions, imperatives or other kinds of non-declarative sentences for which confirmation would seem ill-defined.

This being in place, we can exploit Carnap's (1962, Preface) own successful explication of incremental (dis-)confirmation as probability increase (decrease) upon updating by conditionalization, which has become the standard for almost all of the modern literature on confirmation (see Crupi & Tentori, 2016 for a survey and

discussion): consider the set $Prob(\mathcal{L})$ of epistemic probability measures[3] P on $\mathcal{L}$, that is, rational degree-of-belief functions on $\mathcal{L}$, and define for all A in $\mathcal{S}_1$, B in $\mathcal{S}_2$, P in $Prob(\mathcal{L})$, and all (non-empty) subsets $\mathcal{P}$ of $Prob(\mathcal{L})$,

CON1 A confirms (disconfirms) B relative to P iff $P(B|A) > P(B)$ ($P(B|A) < P(B)$),
CON2 B is confirmable (disconfirmable) by members of $\mathcal{S}_1$ rel. to $\mathcal{P}$ iff there are A in $\mathcal{S}_1$ and P in $\mathcal{P}$, such that A confirms (disconfirms) B relative to P.

Using this, we can specify VC from above in the clearer, more precise, and more explicit form:

VC($\mathcal{L}, \mathcal{S}_1, \mathcal{S}_2, \mathcal{P}$) A sentence in $\mathcal{S}_2$ is meaningful rel. to $\mathcal{S}_1$ and $\mathcal{P}$ iff it is confirmable by members of $\mathcal{S}_1$ rel. to $\mathcal{P}$ or disconfirmable by members of $\mathcal{S}_1$ rel. to $\mathcal{P}$.

Thus, the '(dis-)confirmation' of sentences in $\mathcal{S}_2$ is understood probabilistically in line with CON1, the '-able' in '(dis-)confirmable' is explained by existential quantification over $\mathcal{S}_1$ and $\mathcal{P}$ as stated in CON2, and any relevant background information is supposed to be captured by what the epistemic probability measures in $\mathcal{P}$ have in common. E.g., in a scientific context, auxiliary theoretical and experimental assumptions might well restrict the class of probability measures relative to which (dis-)confirmation is determined to a proper subset $\mathcal{P}$ of $Prob(\mathcal{L})$.[4] Any concrete instance of the resulting Verifiability Criterion "scheme" VC($\mathcal{L}$, $\mathcal{S}_1$, $\mathcal{S}_2$, $\mathcal{P}$) will then result from concrete choices of $\mathcal{L}$, $\mathcal{S}_1$, $\mathcal{S}_2$, $\mathcal{P}$.

I hasten to emphasize that this is not at all the first attempt at a probabilistic reconstruction of criteria that are at least "in the ballpark" of the Verifiability Criterion (see Lutz, 2012, Chapter 8, and Lutz 2017, p. 226f, for surveys, see Watanabe, 1969 for an early reference on this): in particular, in the modern literature, Skyrms (1984, 1985) reconsidered a variant of the Verifiability Criterion in Bayesian terms except for replacing the semantic notion of *meaningfulness* on the left-hand side of the criterion by the purely epistemic notion of (*a posteriori*) *knowability*.[5] Similarly, Sober (in several publications from Sober, 1990 to Sober, 2008) proposed

[3] P is a member of $Prob(\mathcal{L})$ just in case it maps the sentences of $\mathcal{L}$ into the interval [0, 1], such that for all logical truths A it holds that $P(A) = 1$, and if A and B are logically inconsistent with each other, then $P(A \vee B) = P(A) + P(B)$. (So only finite additivity of P is assumed). The conditional probability $P(B|A)$ is given by $\frac{P(B \wedge A)}{P(A)}$ when $P(A) > 0$.

[4] This also goes some way towards accommodating the claim by Quine (1951) that confirmation is holistic; see Supplement E of Leitgeb and Carus (2021) for a discussion.

[5] "The positivists were onto something fundamental with the verification principle, but they misidentified it. Knowability was conflated with meaningfulness, epistemology with semantics. When transposed to its proper pragmatic setting, the verification principle can be seen as a truism elucidating the concept of knowledge" (Skyrms, 1984, p. 19). And: "The status of the [verification] principle of empiricism was obscured by calling it a principle of *meaningfulness* rather than a principle of knowability. Pragmatic, epistemological status was conflated with semantic status. When reformulated in the proper pragmatic framework, the principle of empiricism is a cornerstone of the theory of knowledge" (Skyrms, 1985, p. 28, his emphasis).

and defended a probabilistic criterion of *contrastive empirical testability*, another epistemic concept which, like Skyrms, he distinguishes from the semantic concept of meaningfulness that figures in the Verifiability Criterion.[6] While I will build on some of Skyrms' and Sober's important insights in Sect. 6, there is actually no need to deviate so drastically from the original formulation of the criterion as they did: it is still possible to defend a Verifiability Criterion of meaningfulness proper, as I will argue in the next two sections. The key move will be not to identify meaningfulness with *truth-conditional* meaningfulness, as Skyrms and Sober had presupposed, but to tie meaningfulness in the verifiability criterion to meaningfulness in the *pragmatic-epistemic sense of modern dynamic semantics*.[7]

3 Enter dynamic meaning

The third step of my argumentation incorporates a highly successful conception of meaning that has emerged in semantics in recent decades: meaning as *context change potential*, as developed in dynamic semantics (or update semantics; see Lewis, 2017 for a survey). In the words of Veltman (1996):

> The slogan 'You know the meaning of a sentence if you know the conditions under which it is true' is replaced by this one: 'You know the meaning of a sentence if you know the change it brings about in the information state of anyone who accepts the news conveyed by it'. Thus, meaning becomes a dynamic notion: the meaning of a sentence is an operation on information states (Veltman, 1996, p. 221)

The idea is that the (dynamic) meaning of a sentence is not identified with the truth conditions of the sentence but rather with the updates to conversational contexts it(s assertion) brings about. Meanings are thus "context change potentials (CCPs), which are functions from context to context" where "There is no one concept of context that is inherent to the dynamic semantic framework, though contexts are generally thought to be representations of the current state of the conversation, or the information states of the conversational participants" (Lewis, 2017, p. 1). Whilst the original formulations of dynamic semantics used qualitative formalizations of contexts and beliefs (as sets of worlds or of propositions) as well as of updates (involving intersections with propositions), it has become more and more common in the more recent literature to turn to probabilistic formalizations thereof:

[6] "It seems clear that meaningfulness and testability are different. I suppose that the sentence 'undetectable angels exist' is untestable, but the sentence is not meaningless gibberish. We know what it says, what logical relations it bears to other statements, and we discuss whether it is knowable; none of this would be possible if the string of words literally made no sense" (Sober, 2008, pp. 149–150).

[7] I should also add that Sober is skeptical of Bayesian assignments of prior probabilities to theories or evidence; see Sober (2008, pp. 24–32). For a criticism of Sober's explication of testability, see Justus (2011) and Lutz (2012, Chapter 8).

> Following a productive tradition in semantics [...] we view the basic function of language understanding as belief update: moving from a prior belief distribution over worlds (or situations) to a posterior belief distribution given the *literal meaning* of a sentence. Probabilistic conditioning [...] is a very general way to describe updating of degrees of belief [...] Note that conditioning in this way is the natural analogue of the conception of belief update as intersection familiar from dynamic semantics" (Goodman & Lassiter, 2015, p. 664, their emphasis)[8]

In this probabilistic version of dynamic semantics, it is hence a probability measure or a set of such measures that captures the context, and updates involve probabilistic conditionalization. Following the usual tenets of Bayesian decision theory, the probability measures that result from such updates will then serve as the cognitive basis of one's rational decision-making, which is why the probabilistic update that the assertion of a sentence brings about also captures something like the "pragmatic meaning" of that sentence, approximating a more traditional pragmatist understanding of meaning.[9]

My aim is now to express this conception of meaning more formally again and to use it in my argument for the Verifiability Criterion. When doing so, I will deviate from much of the literature on dynamic semantics in the following manner: although compositionality of dynamic meaning is not strictly a must in dynamic semantics[10], many accounts of dynamic semantics restrict update by intersection (in the qualitative case) or by conditionalization (in the probabilistic case) just to *atomic* sentences, that is, the base cases, and determine updates on *complex* sentences compositionally from updates on their proper sentential parts. E.g., the context change potential of $A \vee B$ in Veltman (1996) update semantics depends functionally on the context change potentials of A and B. In contrast, I will not determine dynamic meaning compositionally but rather define the context change potential of every sentence of $\mathcal{L}$ directly through update by conditionalization. Accordingly, the dynamic meaning assignment to sentences below will neither satisfy the requirement of 'compositional dynamicness' nor that of 'conversation systems dynamicness' that Rothschild and Yalcin (2016) have highlighted as constituting two salient notions of dynamicness in semantics.[11] My reason for doing so is that update in dynamic semantics is supposed to correspond to learning evidence (conveyed by an asserted declarative sentence), and the default Bayesian reconstruction of learning is conditionalization on

[8] See Goodman and Lassiter (2015) for a list of probabilistic works in dynamic semantics. For a recent account of probabilistic semantics and pragmatics more generally, see e.g. Moss (2015).

[9] In the words of Ramsey (1927, p. 170): "The essence of pragmatism I take to be this, that the meaning of a sentence is to be defined by reference to the actions to which asserting it would lead".

[10] E.g., Discourse Representation Theory (see Kamp, 1995), a famous instance of dynamic semantics, is non-compositional (see Geurts et al., 2020, Section 7).

[11] Determining the context change potential of $A \vee B$ compositionally from the context change potentials of A and B is not the same as determining first the truth-conditional meaning of $A \vee B$ compositionally from the truth-conditional meanings of A and B in order to then determine the context change potential of $A \vee B$ by update on the truth-conditional meaning of $A \vee B$; this latter "static" form of compositionality is perfectly compatible with how I will be proceeding.

the evidence (or generalizations thereof). If updating on, say, $A \vee B$, did not proceed by conditionalization on $A \vee B$, a dynamic Dutch book argument (see Lewis, 1999a) could be made against the receiver of the asserted $A \vee B$, and the receiver could be shown not to maximize the expected accuracy of her degrees of belief (see Greaves & Wallace, 2006). In other words: for pragmatically and epistemically rational agents, update by conditionalization should be the default option. Although there are special cases, such as sentences with indefinite or epistemic expressions, on which rational update may well differ from conditionalization, I will just put them aside here. (I will return briefly to questions of compositionality in Sect. 6.3. and to the joint "dynamization" of matters of meaning and confirmation as a future research topic in Sect. 7). This said, even though I am not going to exploit semantic rules that are properly dynamic, I will still invoke a dynamic conception of meaning as context change potential in what follows.

So let us consider again a language $\mathcal{L}$ as before and a subset $\mathcal{P}$ of the set $Prob(\mathcal{L})$ of all epistemic probability measures on $\mathcal{L}$, where I will also assume that $\mathcal{P}$ is closed under conditionalization on sentences with positive probability. For each member P of $\mathcal{P}$, let P_A be the result of updating P by conditionalization on A (when $P(A) > 0$), that is, for all B in $\mathcal{L}$: $P_A(B) = P(B|A)$.[12] Then we can define for all $\mathcal{P}$, and for all A for which there is at least one P in $\mathcal{P}$ such that $P(A) > 0$ (hence excluding logically false sentences):

DM1 The (dynamic) meaning of A with respect to $\mathcal{P}$ is the context change potential function ccp_A from $\mathcal{P}$ to $\mathcal{P}$ that is given by:

$$ccp_A\colon P \mapsto P_A$$

(When $P(A) = 0$, we take $ccp_A(P)$ to be undefined. So ccp_A is really a partial function).

The set $\mathcal{P}$ is meant to represent the range of rational degree-of-belief functions available to some community of competent speakers of $\mathcal{L}$ in a context. While any single member P of $\mathcal{P}$ might only capture a single person's beliefs in that context, the set $\mathcal{P}$ as a whole is supposed to be large enough to "wash out" individual beliefs in favor of social meaning. If there is significant variance in meaning across individuals in a linguistic community, then this will show up in a larger range $\mathcal{P}$ of available epistemic probability measures. This is analogous to intensional truth-conditional semantics in which the interpretation mapping applied at a single possible world would merely correspond to truth and extension at that world, whereas the interpretation mapping over the whole set of possible worlds is meant to capture the totality of meaning postulates that an interpretation of the language in question needs to satisfy.

How (static) meaning in the sense of truth conditions relates to (dynamic) meaning in the sense of context change potentials is a matter of ongoing debate; let me just emphasize that the two conceptions of meaning are not necessarily in conflict

[12] I will bracket all questions concerning whether there is always a uniquely determined rational response to any particular body of evidence; but see Kopec and Titelbaum (2016) for more on this.

with each other, nor is it necessarily the case that one of them is prior to the other. Instead, one may regard them as capturing two important aspects of meaning—say, "semantic vs. pragmatic meaning"—that mutually complement each other. So far as $\mathcal{L}$ and $Prob(\mathcal{L})$ are concerned, one may think of any probability measure P in $Prob(\mathcal{L})$ as corresponding to a probability measure P^* on the propositions expressed by the members of $\mathcal{L}$, such that $P(A) = P^*([A])$, where $[A]$ is the proposition expressed by A as determined by truth-conditional semantics. Consequently, for every context change potential ccp_A in the sense of DM1, there is a corresponding context change potential $ccp_{[A]}$ that is given by conditionalization on propositions that are determined truth-conditionally. But even if so, this would still not mean that truth-conditional meaning would have to be *metasemantically* prior to dynamic meaning: e.g., it might be that the truth-conditional assignment of propositions to sentences is constrained by the *prior assignment of dynamic meanings to these sentences* to the effect that for every ccp_A there is a $ccp_{[A]}$ as described before.

In any case, for my present purposes it is sufficient to distinguish probabilistic dynamic meaning from truth-conditional meaning and to put on record that probabilistic dynamic meaning constitutes a well-established and theoretically fruitful notion of meaning from contemporary semantics. This will be important in so far as I am going to suggest to understand meaningfulness in the Verifiability Criterion in dynamic terms. I will not take a stand on how confirmability relates to truth-conditional meaning, and wherever some logical empiricists might have suggested confirmability to be a criterion for truth-conditional meaningfulness, I will not follow their lead.[13] Accordingly, where logical empiricists wanted to claim not only that unverifiable/unconfirmable claims lack meaning but that they are also therefore not true, I will not follow their lead either. This constitutes yet another way in which the position argued for in this paper is weaker than the traditional logical empiricist doctrines.[14]

The final step I need to take before I can turn to my argument for the Criterion is to derive a concept of (dynamic) meaning*fulness* from DM1 above.

4 A dynamic and relativized notion of meaningfulness

In one important sense of the term, one should regard every sentence A of a language $\mathcal{L}$ as used in previous sections as meaningful, since any such A must conform to the grammatical rules of $\mathcal{L}$ in order to count as a sentence of $\mathcal{L}$ in the first place.

[13] An example would be Carnap (1963, p. 874) who states that non-confirmability & non-disconfirmability entails lack of cognitive meaning, and lack of cognitive meaning entails that a sentence does not express a proposition. From the viewpoint of the present paper, a charitable interpretation would be that Carnap aims to formulate a bridge principle between dynamic meaning and truth-conditional meaning that is meant to hold for a particular class of empiristically preferable languages. A less charitable interpretation would be that he confused dynamic meaning with truth-conditional meaning. (However, Carnap, 1936b is perfectly clear about the difference between truth and confirmation). Either way, the present paper will not take a stand on this.

[14] I am grateful to an anonymous reviewer for bringing this up.

Accordingly, every member A of $\mathcal{L}$ has been assigned a dynamic meaning in DM1 in last section.[15] However, sometimes we may want to require more of a meaningful sentence than just grammatical well-formedness: asserting a sentence that is meaningful (in a context) should have the potential of making a communicative difference, that is, of having communicative impact. One way of tracking communicative impact is to determine whether the assertion of the sentence would have any effect on the epistemic probabilities of other sentences that have already been determined to be meaningful (in that context) or which are presupposed as meaningful (in that context). In terms of a causal analogy: in order for an event to have causal power or impact on some other events, its occurrence should make a causal difference to the occurrence of these other events. Analogously, in what follows, I am going suggest a "power" or "impact" notion of meaningfulness of a sentence with respect to others that is based on communicative difference-making.

Let me start with an example: 'At the Miami Open 2022 Alcaraz won 35 out of 50 drop shots' is a well-formed English sentence but would nevertheless not mean much to anyone who is not sufficiently familiar with tennis terminology. They might have heard about Alcaraz and guess that 'drop shot' expresses a type of tennis stroke, but they might be uncertain of the actual meaning of the term. One way of making that (lack of) understanding manifest is to determine what difference updating on that sentence would make for someone's attitudes towards those sentences that they understand already and which do not involve tennis terminology. Indeed, for anyone who (rightly) understands 'drop shot' to mean an attempt at hitting the ball so lightly that it drops immediately behind the net, being told

(A) At the Miami Open 2022 Alcaraz won 35 out of 50 drop shots

should increase the probability of

(B) At the Miami Open 2022 Alcaraz won 35 out of 50 attempts at hitting the ball so lightly that it drops immediately behind the net.

On the other hand, being told the same sentence A would not (or not much) affect the probability of B for anyone who is uncertain of the meaning of the term 'drop shot': e.g., initially, such a person might assign a certain degree of belief to A (say, $\frac{1}{2}$ or $\frac{1}{5}$ or the like, not being sure about the meaning of 'drop shot' or Alcaraz' performance) and also a certain degree of belief to B (not being sure about Alcaraz' performance), and then retain the original degree of belief in B when being told A. At least so far as B is concerned, A is not meaningful for any such person, in the sense that an assertion of A would not have any communicative effect on their degree of

[15] Since it is precisely the members of $\mathcal{L}$ that were assigned $\mathcal{P}$-probabilities in DM1, we could also say that a meaningful sentence in this first sense of the term is one that is assigned a ($\mathcal{P}$-)probability. This seems to have been Reichenbach's (1938, p. 54) view: "*First principle of the probability theory of meaning: a proposition has meaning if it is possible to determine a weight, i.e., a degree of probability, for the proposition*" (his emphasis).

belief in B. This does not mean that A would be semantically or epistemically problematic *per se*: A has perfectly fine truth conditions, and betting on it or against it, guided by an epistemic probability of $\frac{1}{2}$ or $\frac{1}{5}$ or the like, may well lead to a win (so there is no Dutch book against the agent). It is just that communicating A does not make any difference to the agent's degrees of belief in sentences such as B.

Let us capture this idea more formally again: let $\mathcal{S}_O$ be a non-empty set of "old" sentences in $\mathcal{L}$ that the group of subjects in question already understands (sufficiently for some purpose) or which they take to be understood and in terms of which the meaningfulness of another non-empty set $\mathcal{S}_N$ of "new" sentences is to be appraised. (In the previous example, A would have been a member of $\mathcal{S}_N$ and B of $\mathcal{S}_O$). In our intended applications, $\mathcal{S}_O$ will thus be disjoint from $\mathcal{S}_N$, and we exclude again all sentences A from $\mathcal{S}_N$ for which ccp_A would be undefined on all P. Then we can define for all A in $\mathcal{S}_N$:

DM2 A is (dynamically) meaningful with respect to $\mathcal{S}_O$ and $\mathcal{P}$ iff restricting the (dynamic) meaning ccp_A of A with respect to $\mathcal{P}$ to the members of $\mathcal{S}_O$ yields a function that differs from the identity mapping. That is: there are B in $\mathcal{S}_O$ and P in $\mathcal{P}$, such that $ccp_A(P)(B) \neq P(B)$.

Thus, A is meaningful with respect to $\mathcal{S}_O$ and $\mathcal{P}$ just in case updating on A makes a difference to at least some of the $\mathcal{P}$-probabilities on $\mathcal{S}_O$. The resulting notion of meaningfulness is dynamic by building on DM1 from the last section, it is relativized by depending on the choices of $\mathcal{S}_O$ and $\mathcal{P}$, and it may be called 'pragmatic' in the following sense: a sentence A is meaningless (not meaningful) with respect to $\mathcal{S}_O$ and $\mathcal{P}$ just in case asserting A does not have any communicative impact on the part of the context that is given by the $\mathcal{P}$-probabilities of $\mathcal{S}_O$-sentences, and hence asserting it does not have impact on any of the decision-theoretic dispositions that are captured by that part of the context either. Where the dynamic meaning of a sentence concerns what kind of impact an assertion of the sentence has, the dynamic meaningfulness of sentence concerns whether an assertion of the sentence has any impact at all (on $\mathcal{P}$-probabilities of $\mathcal{S}_O$-sentences). The idea of throwing light on the meaning of certain sentences (members of $\mathcal{S}_N$) based on some "O-sentences" (members of $\mathcal{S}_O$) I am borrowing from Lewis (1970) who used a similar setting when proposing his method of defining "new" theoretical terms on the basis of some given O-terms, where 'O' does not necessarily mean 'observational' but rather *old* or *understood* (see Lewis, 1970, p. 428).

I will deal with concrete applications of this concept of meaningfulness in Sect. 6. For now, let me illustrate its negation—meaninglessness—with the help of a toy model: assume that $\mathcal{L}$ is the set of formulas of propositional logic in the propositional variables $p_1, \ldots, p_n, q$, and $\mathcal{L}^-$ is the subset of $\mathcal{L}$ that consists just of formulas in which q does not occur. Consider any non-empty set $\mathcal{P}^-$ of probability measures on $\mathcal{L}^-$: it is easy to see that for each P^- in $\mathcal{P}^-$ there is a unique extension of P^- to a probability measure P on $\mathcal{L}$ such that for all A in $\mathcal{L}^-$: $P(q \wedge A) = tP^-(A)$ and $P(\neg q \wedge A) = (1-t)P^-(A)$ (where t is a number in the interval [0, 1] that has been fixed antecedently). Let $\mathcal{P}$ be the set of uniquely determined extensions of the

members of $\mathcal{P}^-$ to probability measures on $\mathcal{L}$ in the sense just explained. Then it follows that q is meaningless with respect to $\mathcal{L}^-$ and $\mathcal{P}$ according to DM2 from above: learning q would not make any difference whatsoever with respect to $\mathcal{P}$-probabilities of formulas based solely on $p_1, \ldots, p_n$. At the very least, toy models like this demonstrate that the notion of meaningfulness from above is not affected by any immediate triviality problems.

Meaningfulness in the sense of DM2 is highly idealized, as even miniscule changes of probabilities of members of $\mathcal{S}_O$ suffice for meaningfulness. It would not be difficult to determine a more robust notion by either taking into account the extent to which updates change the $\mathcal{P}$-probabilities on $\mathcal{S}_O$ or by introducing a quantitative concept of meaningfulness from the start. But since the underlying rationale would remain the same, it will be more useful for my present purposes to stick to the simpler definition DM2 when I will finally turn to the argument for the Verifiability Criterion.

5 Putting things together: the argument

We are now ready to formulate the promised argument for our version of the Verifiability Criterion from Sect. 2: assume $\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$ to be determined as explained in previous sections (subject to the same presuppositions as stated there). The premises of the argument are: DM1, DM2, the axioms of probability, the definitions of P_A and conditional probability, CON1, CON2 (and of course bits of real number theory and logic). The discussion from the previous sections should have made clear that these premises are plausible. Finally, the conclusion of the argument is VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$). It remains to be seen that reasoning from these premises to that conclusion constitutes a logically valid argument. That is indeed so, as for all members A of $\mathcal{S}_N$:

(1) A is meaningful with respect to $\mathcal{S}_O$ and $\mathcal{P}$ iff, by DM2,
(2) There are B in $\mathcal{S}_O$, P in $\mathcal{P}$, s.t. $ccp_A(P)(B) \neq P(B)$ iff, by DM1,
(3) There are B in $\mathcal{S}_O$, P in $\mathcal{P}$, s.t. $P_A(B) \neq P(B)$ iff, by def. of P_A,
(4) There are B in $\mathcal{S}_O$, P in $\mathcal{P}$, s.t. $P(B|A) \neq P(B)$ iff, by real calculus,
(5) There are B in $\mathcal{S}_O$, P in $\mathcal{P}$, s.t. $P(B|A) > P(B)$ or $P(B|A) < P(B)$ iff,
by the axioms of prob., the def. of cond. prob., and Bayes' Theorem,
(6) There are B in $\mathcal{S}_O$, P in $\mathcal{P}$, s.t. $P(A|B) > P(A)$ or $P(A|B) < P(A)$ iff,
by CON1 and logic,
(7) There are B in $\mathcal{S}_O$, P in $\mathcal{P}$, s.t. B confirms A rel. to P or there are B in $\mathcal{S}_O$, P in $\mathcal{P}$, s.t. B disconfirms A rel. to P iff,
by CON2 (and logic),
(8) A is confirmable by members of $\mathcal{S}_O$ rel. to $\mathcal{P}$ or A is disconfirmable by members of $\mathcal{S}_O$ rel. to $\mathcal{P}$.

The resulting equivalence of (1)–(8) is just the Verifiability Criterion VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$). The parameters $\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$ relative to which meaningfulness is explained in (1) need to coincide to those for confirmability in (8) in order for this to go through.

The key step of the argument is the transition from (5) to (6) in which Bayes' Theorem allows one to reverse the roles of A and B and thus to turn the meaningfulness of A with respect to Bs into the confirmability of A by Bs. Hence, the Verifiability Criterion is vindicated.

6 Discussion

It is time to reap the fruit of our improved Verifiability Criterion, to determine some of its consequences, and to assess its methodological standing in view of our argument for it from the last section.

6.1 Applying the criterion

Let me start by returning to the original motivation for the criterion. Following the logical empiricists' intended application, our Verifiability Criterion "scheme" VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) could be instantiated as follows:

- Let $\mathcal{L}$ be some (logically regimented) fragment of ordinary, scientific, and philosophical language.

The subset $\mathcal{S}_N$ of $\mathcal{L}$ would consist of whatever possibly problematic "metaphysical sentences" (or "theological sentences" or the like) that ought to be judged for their meaningfulness. For instance, using an example from (Ayer, 1936, p. 7):

- Let $\mathcal{S}_N$ be the singleton set of Bradley's (slightly changed by Ayer) (A) 'The Absolute enters into, but is itself incapable of, evolution and progress'.

The goal would thus be to assess the meaningfulness of A by studying its potential impact on a set $\mathcal{S}_O$ of "understood" sentences of $\mathcal{L}$ relative to a set $\mathcal{P}$ of rational degree-of-belief functions on $\mathcal{L}$. Following empiricist inclinations, we might e.g. have:

- Let $\mathcal{S}_O$ be a set of observation sentences (determined in some manner acceptable to logical empiricists), including, say, (B) 'The temperature of physical body b at place s and time t is d degrees centigrade', and the like.[16]

Alternatively, a perhaps more useful choice still at least in the scientific spirit of the logical empiricists would be some set $\mathcal{S}_O$ of representative members of languages of empirically successful scientific theories, including sentences that express universal

[16] "such a metaphysical pseudo-proposition as 'the Absolute enters into, but is itself incapable of, evolution and progress,' is not even in principle verifiable. For one cannot conceive of an observation which would enable one to determine whether the Absolute did, or did not, enter into evolution and progress" (Ayer, 1936, p. 7).

laws, statistical regularities, observation data, physical and psychological existence claims, and the like.

Finally, $\mathcal{P}$ might be identified as follows:

- Let $\mathcal{P}$ be the set of rational degree-of-belief assignments available to the community of logical empiricists at the time.

Presumably,

- *A* is neither confirmable nor disconfirmable by members of $\mathcal{S}_O$ relative to $\mathcal{P}$.

That is because, e.g., updating any member *P* of $\mathcal{P}$ on the temperature sentence *B* from above is going to leave the *P*-probability of the Bradleyan sentence *A* about the Absolute invariant.

Hence, it follows from VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) that

- *A* is not meaningful with respect to $\mathcal{S}_O$ and $\mathcal{P}$,

as claimed by the logical empiricists. In this way, we can apply VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) as originally intended.

In the special case in which $\mathcal{S}_O$ is the set of all observation sentences in $\mathcal{L}$, and $\mathcal{P}$ is a singleton set of just one probability measure *P*, it also follows that *A being meaningless with respect to* $\mathcal{S}_O$ *and* $\mathcal{P}$ is in fact equivalent to the Bayesian reconstruction of *A being such that it cannot be observationally tested* (*relative to P*) that served as the starting point of Sober's more sophisticated probabilistic theory of *contrastive empirical testability* (see Sober, 2008, pp. 150–152, for details) that had been mentioned in Sect. 2. As Sober makes clear, empirical test*ability* is an important methodological concept that is just as sensible as that of empirical *testing*: "Testing is to testability as dissolving [of salt in water] is to solubility" (Sober, 2008, p. 149). It is a virtue of the present approach that empirical testability may still be viewed as a special instance of meaningfulness in the sense developed above relative to parameters set appropriately. Sober (2008) has been criticized by Justus (2011) and Lutz (2012): in particular, Justus (2011, p. 429f) criticizes Sober (2008, p. 151) contrastive testability criterion as suffering from "imprecision" concerning the terms 'auxiliary assumption', 'now', 'justified in believing', and 'justification depending on believing' it invokes. This criticism does not apply to VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) as formulated above, since all relevant parameters have been made at least formally precise and explicit.

But that is not quite the end of the story, precisely because of the parameters that have been made explicit. For meaningfulness as featured in the new "enlightened" Verifiability Criterion is doubly relativized, and the parameters may also be set differently. First, keeping $\mathcal{S}_O$ fixed for the moment, there is the relativity to $\mathcal{P}$, which might e.g. have the consequence that *A* from above is indeed meaningless relative to the set $\mathcal{P}$ of epistemic probability measures available to the logical empirists while at the same time meaning*ful* relative to the set $\mathcal{P}'$ of epistemic probability measures *available to British idealists at Oxford in the period until Bradley's death in 1924*.

E.g., for these Oxford idealists, A might well be confirmed by instances of biological or cultural evolution as perhaps described by some member B of $\mathcal{S}_O$, since updating some of their degree-of-belief functions on B might indeed increase the probability of "The Absolute enters into… evolution…". (E.g., this might be because some of these "Bradleyan" P might assign a high prior probability to 'If A then B', for whatever reasons.) This kind of relativity has been highlighted by Skyrms (1984, 1985) who rightly says about Bayesian accounts of the Verifiability Criterion:

> The foregoing formulation [of a credibility theory of meaning] is relative to a probability function, and we want to take seriously the idea that whatever the constraints of rationality may be, they leave open a rich variety of eligible probability functions. Empirical meaningfulness is then empirical meaningfulness for epistemic agent X at time t. *One man's metaphysics may be another man's empirically meaningful proposition*" (Skyrms, 1984, p. 15, his emphasis).

Secondly, there is the relativity to $\mathcal{S}_O$: A might also be meaningful relative to some alternative set $\mathcal{S}'_O$ of "old" or "understood" sentences and a suitable set of probability measures. E.g., A might be meaningful relative to the set $\mathcal{S}'_O$ of sentences appearing in Bradley's *Appearance and Reality* prior to the first occurrence of A, and relative to the set of probability measures available to philosophy students at Oxford at the time for whom reading some of these sentences in *Appearance and Reality* may well have supported A.

Clearly, meaningfulness in these last two "Bradleyan senses" no longer corresponds to empirical testability. Of course, the logical empiricists would have protested that, while meaningfulness relative to an observational/scientific $\mathcal{S}_O$ would capture that updating on A has intersubjectively accessible linguistic impact, meaningfulness relative to a "Bradleyan" $\mathcal{S}'_O$ would merely express A's relevance to the internal affairs of Bradley's world picture. And Carnap (1962) would have recommended that the probability measures relative to which meaningfulness is to be determined ought to derive from conditionalizing some distinguished "initial" probability measures on incoming streams of empirical evidence, where these distinguished "initial" probability measures would have to satisfy some logicality, structurality, or learnability constraints that would go beyond the axioms of probability, by which the overall inductive system would ultimately rule out all posterior "Bradleyan probability measures" from being rationally available.

Be that as it may, for my present purposes it is more important to observe that different values for '$\mathcal{S}_O$' and '$\mathcal{P}$' allow for different kinds of applications of the Verifiability Criterion and for different criteria of meaningfulness that result from them. For instance: are set-theoretic statements that are undecidable in ZFC meaningful with respect to the language of arithmetic and the set of probability measures available to ordinary professional mathematicians? Are statements of string theory meaningful with respect to the language of relativity theory/quantum mechanics and the set of probability measures available to present-day physicists? Are prescriptive (normative or evaluative) sentences meaningful with respect to a set of descriptive sentences and the set of probability measures available to competent ordinary speakers or professional moral philosophers? (Assuming that probabilities have been

assigned to prescriptive sentences, too). At least in principle, any such question can be addressed using appropriate instances of VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$), and there are many more potential applications that may not have concerned the logical empiricists but which may well be relevant today. The revised and precisified Verifiability Criterion may be of use whether or not its parameters are set in ways envisioned by the logical empiricists.[17]

This said, for much the same reason, the explication of the verifiability criterion in terms of VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) does not leave the criterion with the normative force the logical empiricists had hoped for originally: for the meaningfulness of, e.g., metaphysical discourse is going to depend on how the relevant parameters are set, e.g., what the probabilistic profile of the relevant community is like. In fact, taken by itself, the revised verifiability criterion is not normative at all but closer to a probabilistic criterion of linguistic aboutness or dependency along the lines of (the non-probabilistic) Lewis (1988) and Yablo (2014) on the semantic explication of aboutness and subject matter. However, VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) may still be *combined* with normative claims concerning how its parameters ought to be chosen for certain purposes, e.g., to guarantee mutual intersubjective understanding in the empirical sciences. The more positive take on the new VC's lack of normative power is thus that it enforces a clear separation of the descriptive and the normative components of the original verificationist package.

6.2 Logic, synonymy, compositionality

I will turn to some more general semantic features of this combined account of meaning and confirmation now.

The logical empiricists restricted the Verifiability Criterion to logically *contingent* declarative sentences as they (correctly) thought that logical truths and falsehoods could neither be verified/confirmed properly nor were in need of verification/confirmation.[18] Accordingly, even if a logical truth, e.g., a sentence of the form $A \vee \neg A$, were a member of $\mathcal{S}_N$, it would simply follow that the sentence would be (dynamically) meaningless with respect to $\mathcal{S}_O$ and $\mathcal{P}$, which would reflect probabilistically that logical truths are devoid of content. (Logical falsities were excluded from $\mathcal{S}_N$ from the start: recall Sect. 4).

The standard conception of logical consequence from dynamic semantics is perfectly compatible with the present probabilistic framework, too: for in dynamic semantics one typically first defines acceptance of a sentence in a context as invariance of the context under update with the sentence (an idea going back to Heim, 1983) and then defines logical consequence by acceptance, e.g., amongst other

[17] The resulting package of meaning and confirmation also bears some affinity with recent "use conceptions" of meaning underlying distributional semantics in natural language processing (see Clark, 2015).

[18] E.g., when introducing his replacement of the traditional Verifiability Criterion by proposals for the construction of empiricistically acceptable languages, Carnap (1937) restricts his requirements to synthetic sentences (thus excluding logical ones). E.g.: "*Requirement of Confirmability:* 'Every synthetic sentence must be confirmable"' Carnap (1937, p. 34).

options, as preservation of acceptance in all contexts (see the definition of $valid_3$ in Veltman, 1996, p. 224). In the present setting: for all $A_1, \ldots, A_n, B$ in $\mathcal{L}$, for all P in $Prob(\mathcal{L})$,

- A is (fully) accepted in P if and only if $ccp_A(P) = P$,

by which acceptance of A in P is just equivalent to $P(A) = 1$, and

- $A_1, \ldots, A_n \models B$ if and only if for all P in $Prob(\mathcal{L})$: if each of $A_1, \ldots, A_n$ is accepted in P, then B is accepted in P,

which is easily seen to yield the same logical consequence relation of classical logic that had been "hardwired" into the definition of probability measures from the start (recall Footnote 3). And, of course, every logical truth is indeed accepted in every P in $Prob(\mathcal{L})$.

Moving away from logical truths and returning to the tennis example from Sect. 4, if the set $\mathcal{P}$ available to a community of tennis aficionados satisfies

(Ass.) for all P in $\mathcal{P}$: $P(\forall x(Dropshot(x) \leftrightarrow A[x])) = 1$

where $A[x]$ is short for x being an attempt at hitting the ball so lightly that it drops immediately behind the net, then also the sentence

$$\forall x(Dropshot(x) \leftrightarrow A[x])$$

is meaningless with respect to $\mathcal{P}$ (and whatever $\mathcal{S}_O$). Whilst this may sound curious and may remind one of the famous *Old Evidence Problem* of Bayesian confirmation theory (see Glymour, 1980), it does not mean that the sentence would be dynamically meaningless for all probability measures in the total set $Prob(\mathcal{L})$ (in contrast to logical truths which are meaningless in that sense), and it does not mean either that the sentence would not have any truth-conditional meaning. It only reflects that telling any of our tennis aficionados 'a dropshot is an attempt at hitting the ball so lightly that it drops immediately behind the net' would not have any communicative effect on them that could be measured by epistemic probability,[19] since they already accept the sentence no matter what. Accordingly, for them—much like a logical truth—the sentence could not be (dis-)confirmed nor would be in need of (dis-)confirmation.

More interestingly, (Ass.) from above entails with the axioms of probability that the sentence $Dropshot(a)$ ('a is a dropshot') is synonymous to $A[a]$ (with the $A[x]$ from above) in the sense given by:

[19] There are also communicative effects beyond what is measurable by epistemic probability: e.g., two probabilistically indistinguishable sentences may be such that one is funny and makes another person from the same community laugh while the other one does not. I am excluding such additional communicative effects from the discussion here. (I am grateful to an anonymous reviewer for having highlighted this in a comment).

- A is (dynamically) synonymous to B with respect to $\mathcal{P}$ if and only if the dynamic meaning of A with respect to $\mathcal{P}$ is identical to the dynamic meaning of B with respect to $\mathcal{P}$, that is: $ccp_A = ccp_B$.[20]

Similarly, it follows that *Dropshot*(a) is meaningful with respect to $\{A[a]\}$ and $\mathcal{P}$, and $A[a]$ is meaningful with respect to $\{$*Dropshot*$(a)\}$ and $\mathcal{P}$, both in the sense of DM2. But, of course, A being (dynamically) synonymous to B with respect to $\mathcal{P}$ is generally a much stronger requirement than A being meaningful with respect to $\{B\}$ and $\mathcal{P}$. Even there being just a meaning *resemblance* between A and B with respect to $\mathcal{P}$—some $\mathcal{P}$-probabilities conditional on A being more or less correlated with $\mathcal{P}$ -probabilities conditional on B—would generally be way stronger than the two sentences being meaningful with respect to each other.

Finally, as the formal toy example from Sect. 4 makes clear, meaningfulness in the sense of DM2 is just as non-compositional as dynamic meaning in the sense of DM1 (recall Sect. 4) and as probabilistic confirmation (see Fitelson, 2002; Schippers & Schurz, 2020 on the closely related problem of "irrelevant conjunction"). E.g., while q was found to be meaningless with respect to $\mathcal{L}^-$ and $\mathcal{P}$ in Sect. 4, the conjunction $q \wedge p_1$ is clearly meaning*ful* with respect to $\mathcal{L}^-$ and $\mathcal{P}$ so long as there is at least one P^- in $\mathcal{P}^-$ for which $0 < P^-(p_1) < 1$ is the case; for then the probability measure P that extends such P^- has the property that $P(q \wedge p_1 | p_1) = t > t \cdot P^-(p_1) = P(q \wedge p_1)$. That is: a complex meaningful sentence can have meaningless proper sentential parts (as observed already by Skyrms, 1984, p. 16). For analogous reasons, a complex meaningless sentence can be solely composed of meaningful proper sentential parts: the logical truth $p_1 \vee \neg p_1$ serves as an example.

While these features clearly drive a wedge between dynamic and truth-conditional meaningfulness, they cease to be worrisome again if cast in terms of communicative impact: a complex sentence can have communicative impact even when some of its proper sentential parts do not (for other parts of it may have impact), and a complex sentence may not have any communicative impact even when all of its proper parts do (for their individual impacts may cancel each other out). None of this undermines the dynamic notion of meaningfulness that figures in our reconstruction of the Verifiability Criterion nor does it affect its relationship to confirmability.

6.3 The status of the criterion

Many of the traditional objections to the Verifiability Criterion consisted in the construction of counterexamples to specific variants of the Criterion that are no longer relevant in the present context. E.g., the criterion Ayer famously suggested in the second edition of his *Language, Truth and Logic* (Ayer, 1946) involved the idea of

[20] With the help of VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) it is easy to see that this definition of dynamic synonymy coincides with Reichenbach's (1938, p. 54f) "Second principle of the probability theory of meaning: two sentences have the same meaning if they obtain the same weight, or degree of probability, by every possible observation", except that Reichenbach's 'by every possible observation' needs to be changed into 'conditional on every sentence in $\mathcal{L}$ (for which conditional probability is defined)'.

a meaningful sentence making a difference to the derivability of directly verifiable sentences if combined with suitable auxiliary premises. Ayer's respective existential quantification over auxiliary premises, which led to famous triviality objections, could be thought of as corresponding to the existential quantification over epistemic probability measures from a given set $\mathcal{P}$ in DM2, but since that set simply figures as one of the parameters to which meaningfulness is relativized (as discussed in Sect. 6.1), no special objections emerge from this.[21] Other objections, which concerned issues of non-compositionality, are still relevant but were considered already at the end of last section.

But there are still two more general worries left to be discussed: worries concerning the methodological status of the Criterion (see e.g. Passmore, 1967)—e.g. on what grounds it might be justified—and its potentially self-undermining character (see e.g. Lycan, 2019, p. 108). There is a long-standing debate on whether these more general objections actually apply to Carnap's practical take on verificationism as a proposal for choosing an empiristically acceptable language: see e.g. Ricketts (1994) for a defense of Carnap's position against Putnam's criticism thereof. I will conclude my discussion by solely addressing the two worries to the extent to which they concern VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) from above.

In order to do so, it is useful to distinguish between two distinct, yet closely related, semantic projects (see Partee, 2011): on the one hand, the scientific-linguistic (or perhaps naturalized-philosophical) project of describing, predicting, and explaining semantic phenomena concerning natural language(s), and, on the other hand, the normative-philosophical project of rationally reconstructing and improving semantic phenomena concerning natural language(s) by introducing formal semantics for constructed languages, in particular, formal semantics for logical reconstructions of fragments of natural language. Within the former linguistic project, each instance of VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) is but an empirical thesis that gains support from the premises on which the argument from Sect. 5 was based, including dynamic semantics (as a branch of linguistic semantics) and Bayesian cognitive psychology. Within the latter more traditionally philosophical project, each instance of VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) constitutes a joint clarification, sharpening, and systematization of meaningfulness and confirmability that gains support from the premises on which the argument from Sect. 5 was based, including a dynamic explication of meaning (as a branch of philosophical semantics) and a Bayesian explication of confirmation. Any instance of VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) as an empirical thesis would thus be undermined by empirical evidence to the contrary—evidence that might ultimately prove that instance to be false. And an instance of VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) as a philosophical explication would be undermined by arguing that it was not similar enough to what was to be explicated, or that it was not exact, fruitful, or simple enough: arguments that would ultimately show the explication was not successful enough given its intended purposes. (Compare Carnap's criteria of successful explications in Chapter 1 of Carnap, 1962). Either way, the methodological status of VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) seems

[21] I am grateful to an anonymous reviewer for urging me to comment on the relationship with Ayer's criterion.

sufficiently transparent, and the previous sections should have shown that in fact a lot speaks in favor of its truth *qua* empirical thesis or its success *qua* explication.

Finally, "self-applications" of VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) do not seem particularly problematic either, at least once the notion of "self-application" has been disambiguated appropriately: VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) is a criterion for the object language $\mathcal{L}$ but is itself a sentence in a suitable metalanguage $\mathcal{L}'$ of $\mathcal{L}$; so in order to apply VC to "itself", one would actually have to turn to another version VC($\mathcal{L}'$, $\mathcal{S}'_O$, $\mathcal{S}'_N$, $\mathcal{P}'$) of the Criterion that concerns meaningfulness and confirmability of sentences in $\mathcal{L}'$, including (the relevant instance) VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$).[22] Suitable versions of CON1, CON2, DM1, and DM2 for $\mathcal{L}'$ would then also apply to VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$): subjects would believe the Criterion to certain degrees (in line with the subjects' epistemic probability measures in $\mathcal{P}'$), the Criterion would be confirmable or disconfirmable by other sentences of $\mathcal{L}'$ relative to members of $\mathcal{P}'$, and if VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) belongs to the set $\mathcal{S}'_N$ of sentences of special interest, then VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) would turn out to be meaningful with respect to $\mathcal{S}'_O$ and $\mathcal{P}'$ just in case it would be confirmable by members of $\mathcal{S}_O$ relative to $\mathcal{P}'$ or disconfirmable by members of $\mathcal{S}_O$ relative to $\mathcal{P}'$. In the extreme case in which the arguments from the present paper had been completely convincing to the community corresponding to $\mathcal{P}'$, the Criterion VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) would be fully accepted in every P in $\mathcal{P}'$ and hence be meaning*less* with respect to $\mathcal{S}'_O$ and $\mathcal{P}'$, for the same reason for which (Ass.) had been found to be meaningless for the community of tennis aficionados in Subsection 6.2. Accordingly, in that case, the Criterion could not be (dis-)confirmed by members of $\mathcal{S}_O$ relative to $\mathcal{P}'$ nor would it be in need of confirmation. But nothing bad would follow from this—asserting the criterion in front of members of that community would merely have ceased to have communicative effect again.

7 Conclusions and outlook

This paper developed an argument for a refined version of the Verifiability Criterion of meaningfulness: the argument was found to be logically valid, it was based on independently plausible premises concerning confirmation and meaning, and its outcome was an interesting and non-trivial schematic bridge principle between semantics and epistemology that relates the communicative meaningfulness of declarative sentences to matters of epistemic probabilistic relevance. The Criterion thus vindicated concerns a genuine and theoretically useful notion of meaning, however, one that is dynamical-pragmatic in nature rather than truth-conditional. According to it, e.g., a sentence such as 'The Absolute enters into, but is itself incapable of, evolution and progress' does turn out to be dynamically-pragmatically meaningless with respect to a given set of observation sentences and the logical empiricists' rational

[22] Alternatively, $\mathcal{L}$ would have to be type-free in the sense that VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) would both talk about and be included in $\mathcal{L}$. But that would lead to familiar worries concerning self-referentiality, ungroundedness, and paradoxicality well-known from formal theories of type-free truth, that is, worries orthogonal to debates about the Verifiability Criterion.

degree-of-belief functions. While the new Criterion stayed reasonably close to variants of the Criterion proposed by some of the logical empiricists themselves, it also differed from these variants in crucial ways: in particular, both meaningfulness and confirmability are relativized to probabilistic contexts and sets of sentences on which the communicative impact of the assertion of a sentence is to be determined; and the theory itself does not come with reasons for preferring certain probabilistic contexts and sets of sentences over others. That is also why it remains doubtful whether the new Criterion by itself could be used to demarcate the class of "scientifically respectable" sentences from those that are not, as the logical empiricists had hoped. However, due to its schematic and relativized form, the Criterion may allow for completely new applications and may trigger completely new questions, and it can be combined with separate normative assumptions concerning how its parameters ought to be set for certain purposes.

Let me conclude by sketching two of these potentially interesting new applications or questions (the details of which need to be left to future work). First, the *topological study of meaningfulness* (or, by the Criterion, of confirmability/disconfirmability): for fixed $\mathcal{L}$ and $\mathcal{P}$, and for arbitrary $\mathcal{S}, \mathcal{S}' \subseteq \mathcal{L}$, let us define: the set $\mathcal{S}$ of sentences is ($\mathcal{P}$-)meaningful-with-respect-to the set $\mathcal{S}'$ of sentences just in case there is an A in $\mathcal{S}$ that is meaningful with respect to $\mathcal{S}'$ and $\mathcal{P}$. (Alternatively: $\mathcal{S}$ is ($\mathcal{P}$-)meaningful-with-respect-to $\mathcal{S}'$ just in case there are sufficiently many such A in $\mathcal{S}$ that are meaningful with respect to $\mathcal{S}'$ and $\mathcal{P}$). Now consider a set $\{\mathcal{S}_1, \ldots, \mathcal{S}_n\}$ of non-empty subsets of $\mathcal{L}$. It is easy to see that 'meaningful-with-respect-to' expresses a symmetric relation on $\{\mathcal{S}_1, \ldots, \mathcal{S}_n\}$. So long as for every i there is an A_i in $\mathcal{S}_i$, such that for some P in $\mathcal{P}$ it holds that $P(A_i) < 1$, the relation is also reflexive, but it is not necessarily transitive: it determines a meaningfulness graph or similarity relation rather than an equivalence relation. With this in place, one can start investigating topological properties of meaningfulness, such as: which of the sets $\mathcal{S}_1, \ldots, \mathcal{S}_n$ of sentences is connected to which other via a sequence of edges in the graph? Which is isolated from the others? Which is central in being connected to all others in one step? The members of which "screen off" the meaningfulness of the members of a certain set with respect to the members of another (in the probabilistic sense of 'screening off' that is well-known from probabilistic causal models)? And so forth. In this way, the theory developed in this paper might be used to determine how particular linguistic domains relate to each other epistemically, and to assess (perhaps negatively) those linguistic domains that fail to sufficiently epistemically relate to others. How would, e.g., sets of sentences specific to certain scientific, philosophical, political,… communities do in that respect? And what would this tell us about the similarities or dissimilarities between the conceptual frameworks underlying these linguistic domains (e.g. in the probabilistic-geometric sense studied in Sznajder, 2016)?

Secondly, returning to Sects. 3 and 4, the joint "dynamization" of meaningfulness and confirmability: the method of update or context change on which DM1 in Sect. 3 is based is probabilistic conditionalization, independently of what the logical form of the given sentence A of $\mathcal{L}$ is like on which the update takes place. Accordingly, the corresponding notion of meaningfulness that is given by DM2 in Sect. 4 identifies the communicative impact of asserting A with probabilistic change

induced by conditionalizing on A. But, as explained in Sect. 3, the method of update or context change by which a properly dynamic semantics interprets certain logically complex sentences A, such as those including indefinite or epistemic expressions, may well differ from plain conditionalization. If so, and if we understand DM1 and DM2 accordingly, would it still be possible to defend VC($\mathcal{L}$, $\mathcal{S}_O$, $\mathcal{S}_N$, $\mathcal{P}$) for languages $\mathcal{L}$ that include such sentences A? If the answer is positive, then for such sentences A neither side of the Verifiability Criterion—neither dynamic meaningfulness nor (dis-)confirmability—would be given by the mere existence of sentences that are positively or negatively probabilistically relevant to A. By using the Verifiability Criterion in this way to translate dynamic semantics into the theory of confirmability, the machinery developed in this paper might lead to a fruitful marriage of ideas between the subject matters of semantics/pragmatics and those of formal epistemology/general philosophy of science.

Even if logical empiricism itself is dead, its Verifiability Criterion may well deserve a happy afterlife.

Acknowledgements I am very grateful to Ivano Ciardelli, Benedikt Leitgeb, Thomas Uebel, Paul Weingartner, and two anonymous reviewers for their invaluable feedback on earlier versions of this article.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest The author has no relevant financial and non-financial interests to disclose

References

Ayer, A. J. (1936). *Language, truth and logic*. Victor Gollancz.
Ayer, A. J. (1946). *Language, truth and logic* (2nd ed.). Victor Gollancz.
Carnap, R. (1936). Testability and meaning (Part 1). *Philosophy of Science, 3*(4), 419–471.
Carnap, R. (1936b). Wahrheit und Bewährung. In *Actes du Congrès international de philosophie scientifique, Sorbonne, Paris 1935, fasc. 4. Unité de la science, Hermann, 18–23. A modified version is translated as "Truth and Confirmation", in Feigl and Sellars (1949)* (pp. 119–127).
Carnap, R. (1937). Testability and meaning (Part 2). *Philosophy of Science, 4*(1), 1–40.
Carnap, R. (1945). On inductive logic. *Philosophy of Science, 12*(2), 72–97.
Carnap, R. (1956). The methodological character of theoretical concepts. In H. Feigl & M. Scriven (Eds.), *The foundations of science and the concepts of psychology and psychoanalysis* (pp. 38–76). University of Minnesota Press.
Carnap, R. (1962). *Logical foundations of probability* (2nd ed.). University of Chicago Press.
Carnap, R. (1963). Reply to Henle. In Schilpp (1963) (pp. 874–877).

Clark, S. (2015). Vector space models of lexical meaning. In S. Lappin & C. Fox (Eds.), *Handbook of contemporary semantic theory* (pp. 493–522). Wiley-Blackwell.
Creath, R. (1976). On Kaplan on Carnap on significance. *Philosophical Studies, 30*(6), 393–400.
Creath, R. (2021). Logical empiricism. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Winter 2021 Edition). https://plato.stanford.edu/archives/win2021/entries/logical-empiricism/
Crupi, V., & Tentori, K. (2016). Confirmation theory. In A. Hájek & C. Hitchcock (Eds.), *Oxford handbook of philosophy and probability* (pp. 650–665). Oxford University Press.
Feigl, H., & Sellars, W. (Eds.). (1949). *Readings in philosophical analysis*. Appleton-Century-Crofts.
Fitelson, B. (2002). Putting the irrelevance back into the problem of irrelevant conjunction. *Philosophy of Science, 69*, 611–622.
Geurts, B., Beaver, D. I., & Maier, E. (2020). Discourse representation theory. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2020 Edition). https://plato.stanford.edu/archives/spr2020/entries/discourse-representation-theory/
Glock, H.-J. (2021). Ayer's verificationism: Dead as a dodo? In A. T. Tuboly (Ed.), *The historical and philosophical significance of Ayer's language, truth and logic* (pp. 251–278). Palgrave Macmillan.
Glymour, C. (1980). *Theory and evidence*. Princeton University Press.
Goodman, N. D., & Lassiter, D. (2015). Probabilistic semantics and pragmatics: Uncertainty in language. In S. Lappin & C. Fox (Eds.), *The handbook of contemporary semantic theory* (pp. 655–686). Wiley.
Greaves, H., & Wallace, D. (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind, 115*, 607–632.
Heim, I. (1983) "On the Projection Problem for Presuppositions. In M. Barlow et al. (Ed.), *WCCFL 2: Second Annual West Coast Conference on Formal Linguistics* (pp. 114–125).
Hempel, C. G. (1950). Problems and changes in the empiricist criterion of meaning. *Revue Internationale de Philosophie, 4*(11), 41–63.
Hempel, C. G. (1951). The concepts of cognitive significance: A reconsideration. *Proceedings of the American Academy of Arts and Sciences, 80*(1), 61–77.
Justus, J. (2011). Evidentiary inference in evolutionary biology, review of Sober (2008). *Biology and Philosophy, 26*, 419–437.
Justus, J. (2014). Carnap's forgotten criterion of empirical significance. *Mind, 123*, 415–436.
Kamp, H., et al. (1995). Discourse representation theory. In J. Verschueren (Ed.), *Handbook of pragmatics* (pp. 253–257). John Benjamins.
Kopec, M., & Titelbaum, M. G. (2016). The uniqueness thesis. *Philosophy Compass, 11*(4), 189–200.
Leitgeb, H., & Carus, A. (2021). Rudolf Carnap. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2021 Edition). https://plato.stanford.edu/archives/sum2021/entries/carnap/
Lewis, D. (1970). How to define theoretical terms. *Journal of Philosophy, 67*(13), 427–446.
Lewis, D. (1988). Statements partly about observation. *Philosophical Papers, 17*(1), 1–31.
Lewis, D. (1999a). Why conditionalize? In Lewis (1999b) (pp. 403–407).
Lewis, D. (1999b). *Papers in metaphysics and epistemology*. Cambridge University Press.
Lewis, K. (2017). Dynamic semantics. *Oxford handbook online*. Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199935314.013.14
Lutz, S. (2012). *Criteria of empirical significance. foundations, relations, applications*. Dissertation, Utrecht University.
Lutz, S. (2017). Carnap on empirical significance. *Synthese, 194*, 217–252.
Lycan, W. G. (2019). *Philosophy of language. A contemporary introduction*. Routledge.
Moss, S. (2015). On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics, 8*(5), 1–81.
Moore, G. E., & Ramsey, F. P. (1927). Facts and propositions. *Proceedings of the Aristotelian Society, Supplementary Volumes, 7*, 153–170.
Partee, B. H. (2011). Formal semantics: Origins, issues, early impact. *The Baltic International Yearbook of Cognition, Logic and Communication, 6*, 1–52.
Passmore, J. (1967). Logical positivism. In P. Edwards (Ed.), *The encyclopedia of philosophy* (Vol. 5, pp. 52–57). Macmillan.
Quine, W. V. O. (1951). Two dogmas of empiricism. *The Philosophical Review, 60*(1), 20–43.
Reichenbach, H. (1938). *Experience and prediction. An analysis of the foundations and the structure of knowledge*. The University of Chicago Press.
Ricketts, T. (1994). Carnap's principle of tolerance, empiricism, and comventionalism. In P. Clark & B. Hale (Eds.), *Reading Putnam* (pp. 176–200). Blackwell.

Rothschild, D., & Yalcin, S. (2016). Three notions of dynamicness in language. *Linguistics and Philosophy, 39*, 333–355.
Schilpp, P. A. (Ed.). (1963). *The philosophy of Rudolf Carnap*. Open Court.
Schippers, M., & Schurz, G. (2020). Genuine confirmation and tacking by conjunction. *British Journal of the Philosophy of Science, 71*, 321–352.
Skyrms, B. (1984). *Pragmatics and empiricism*. Yale University Press.
Skyrms, B. (1985). Pragmatics and the principle of empiricism. In N. Rescher (Ed.), *The heritage of logical positivism* (pp. 23–28). University Press of America.
Soames, S. (2003). *Philosophical analysis in the twentieth century* (Vol. 1). Princeton University Press.
Sober, E. (1990). Contrastive empiricism. In C. Savage (Ed.), *Scientific theories* (pp. 392–410). University of Minnesota Press.
Sober, E. (2008). *Evidence and evolution: The logic behind the science*. Cambridge University Press.
Sznajder, M. (2016). What conceptual spaces can do for Carnap's late inductive logic. *Studies in History and Philosophy of Science Part A, 56*, 62–71.
Uebel, T. (2019). Verificationism and (some of) its discontents. *Journal for the History of Analytical Philosophy, 7*(4), 1–31.
Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic, 25*, 221–261.
Watanabe, S. (1969). *Knowing and guessing. A quantitative study of inference and information*. Wiley.
Yablo, S. (2014). *Aboutness*. Princeton University Press.