



Campbell-Moore, C. (2019). Limits in the Revision Theory: More Than Just Definite Verdicts. *Journal of Philosophical Logic*, 48, 11-35.
<https://doi.org/10.1007/s10992-018-9477-y>

Publisher's PDF, also known as Version of record

License (if available):
CC BY

Link to published version (if available):
[10.1007/s10992-018-9477-y](https://doi.org/10.1007/s10992-018-9477-y)

[Link to publication record on the Bristol Research Portal](#)
PDF-document

This is the final published version of the article (version of record). It first appeared online via Springer at <https://link.springer.com/article/10.1007%2Fs10992-018-9477-y> . Please refer to any applicable terms of use of the publisher.

University of Bristol – Bristol Research Portal

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/brp-terms/>



Limits in the Revision Theory

More Than Just Definite Verdicts

Catrin Campbell-Moore¹

Received: 28 July 2016 / Accepted: 11 December 2016 / Published online: 05 September 2018
© The Author(s) 2018

Abstract

We present a new proposal for what to do at limits in the revision theory. The usual criterion for a limit stage is that it should agree with any definite verdicts that have been brought about before that stage. We suggest that one should not only consider definite verdicts that have been brought about but also more general properties; in fact any closed property can be considered. This more general framework is required if we move to considering revision theories for concepts that are concerned with real numbers, but also has consequences for more traditional revision theories such as the revision theory of truth.

Keywords Revision theory · Self-reference · Circular definitions · Taking limits · Probability

1 Introduction

The revision theory is an influential method of working with concepts that are characterised by circular definitions, i.e. definitions that refer to that very notion. The most widely studied instance of the revision theory is the revision theory of truth, which was initially developed by Herzberger and Gupta (independently) to study truth in a way that can account for the liar paradox: a sentence which says of itself that it is not true. They take the equivalence “ φ is true” if and only if φ to be one of definition. To apply the definition one starts with a hypothesis about the interpretation of ‘true’ and then revises this hypothesis in accordance with the definition to determine a new

This work was done partly while at the Munich Center for Mathematical Philosophy, funded by the Alexander von Humboldt Foundation, and partly while at Corpus Christi College, Cambridge, funded by the college.

✉ Catrin Campbell-Moore
catrin@ccampbell-moore.com

¹ Cotham House, University of Bristol, Bristol, BS6 6JL, UK

interpretation. This revised hypothesis can then itself be revised, and so on. This process of taking a hypothesis and ‘revising’ it forms the basis of the revision theory of truth. One thereby obtains a sequence of hypotheses.

The revision sequence that results will typically extend into the transfinite. One needs a criterion for what to do at limit stages as there’s no immediately preceding stage which one can apply the revision step to. The limit stages instead sum up the information in the sequence of hypotheses up to that stage. Exactly how to sum up such information is the topic of this paper: we will propose a limiting policy that imposes more constraints than the one usually considered. Although there are different proposals in the literature for what to do at the limit stage, they all have a common core: a limit stage should sum up any definite verdicts that have been brought about by the sequence up to that stage. We give an alternative proposal for what to do at limits: we suggest one should not only take into account the definite verdicts that have been brought about but also other properties of the hypotheses. We will not specify exactly which properties should be considered but instead give a framework which allows for more general properties. The usual limiting criterion is the weakest of the class of criteria that we consider.

As has been urged by Belnap [1], and further in Gupta and Belnap [6], the revision theory is in fact a quite general account of circular definitions. Whenever a concept is associated with a circular definition a natural revision rule is determined and then the revision theory can be used to account for the operation of this concept. One application of the revision theory is to study concepts that are concerned with real numbers; for example notions of probability or degrees-of-truth. For example it might be used to model an agent rationally reflecting on her credences. If one wants to study such concepts, then the usual limit rule from the revision theory is not sufficient: if one has a revision sequence where the probability of a sentence converges to $1/2$, then one might plausibly want the limit rule to require that the probability of that sentence is equal to $1/2$ at the limit stage. We account for such convergence by considering not just definite verdicts that have been brought about, but also additional properties of the hypotheses such as: the probability of a sentence is at least $1/4$. We will want to restrict our attention to properties that act nicely under limits. For example the property of having the probability of a sentence be strictly less than 1 does not act nicely under limits: a sequence of hypotheses might all have this property and converge to another hypothesis that does not have that property. These notions of convergence and limits can be made precise using topology. We impose the constraint that the only properties that should be considered are those where the set of hypotheses they determine is a closed set in the topological sense. We will then call the associated property a closed property. This restriction to closed properties allows us to show that so long as the underlying space of hypotheses has a certain feature, namely that the topology on it is compact, the limit criterion we impose will be satisfiable.

So the move to considering revision theories where one is interested in concepts like probability motivates a move to considering more general properties of hypotheses that are brought about, as that then allows one to account for convergence. Once we are considering more general properties we might be interested in the strongest such limit criterion: where all closed properties are considered. This proposal also has some interesting consequences for the revision theory of truth: if one adopts this

more restrictive limit criterion, then one obtains the result that limit hypotheses will be required to have additional properties, for example the interpretation of the truth predicate will be maximally consistent. This is a departure from traditional limit criteria but is perhaps desirable because now all the hypotheses in the revision sequence satisfy minimal criteria one might hope for from an interpretation of the truth predicate. This means that we can also obtain interesting theories by determining what ends up holding at all limit stages.¹

The paper proceeds as follows: In Section 2 we will present our proposal for the limit stage criteria in their general form. The proposal is only legitimate if the topology has certain properties, which are outlined in Section 2.1; that section is separated as some readers may wish to skip it. In Section 3 we present the revision theory of truth and discuss how these new limit stage criteria make a difference. In Section 4 we discuss how the limit stage criteria affect the revision theory if the hypothesis space involves reals. Here we see that such a generalisation of the limit criterion is needed to do the intended job of the limit stage. In that section we first generally discuss how the limit criteria should apply (Section 4.1), then present examples of how the strongest criterion, Coherence of all Closed Properties, works (Section 4.2), then we discuss what results if alternative collections of properties are considered (Section 4.3), and finally present a sample revision rule that might be appropriate for a revision theory of probability (Section 4.4). In Section 5 we conclude.

2 Limits in the Revision Theory

In a revision theory one constructs a sequence of hypotheses. Each hypothesis typically provides a hypothetical interpretation of a concept of interest and the sequence is given by revising the hypotheses. In its general form, then, the revision sequence can be seen to be a sequence of hypotheses, h_0, h_1, h_2, \dots , where each hypothesis is drawn from a space of possible hypotheses H . There is an underlying rule, called a revision rule, which determines how one moves from one hypothesis to the next. The revision rule is a function $\rho : H \rightarrow H$, and the revision sequence needs to be such that $h_{n+1} = \rho(h_n)$.

Usually just considering the revision sequence as an ω -length sequence of hypotheses obtained by iterating the revision rule is not enough. For example in the case of truth, if one wants to ensure that $\top \ulcorner \forall n \top^n \ulcorner 0 = 0 \urcorner \urcorner$ comes out as true one might have to move to stages beyond ω .²

¹It should be noted that limit ordinals that are not of the form $\alpha + \omega$ aren't well controlled by this limit policy. A theory which reads off what holds in limits should thus either only focus on ordinals of the form $\alpha + \omega$, or one should weaken the notion of 'brought about' to be that of near stability. See Campbell-Moore [2, Section 5.4] especially Theorem 5.4.2 for theories of this kind, where the near stability variant of Coherence of all Closed Properties is being used.

²Here, $\top^n \ulcorner \varphi \urcorner$ denotes n -applications of the truth predicate: $\overbrace{\top \dots \top}^n \ulcorner \varphi \urcorner$. More carefully, $\top^n x$ is shorthand for $\top f(n, x)$, where f represents the function f defined recursively by $f(0, x) = x$, $f(n + 1, x) = \top f(n, x)$. (See Halbach [7, Theorem 13.9], or in its original form McGege [11].)

A revision sequence is thus a transfinite sequence of hypotheses

$$h_0, h_1, \dots, h_\omega, h_{\omega+1}, \dots$$

These will follow the revision rule for the successor stage, so $h_{\alpha+1} = \rho(h_\alpha)$. We also need to impose certain requirements on what the limit hypotheses should be like in a revision sequence. What requirements can be imposed is the topic of this paper.

In Gupta's first published presentation of his theory he says:

Intuitively what is wanted is a way of summing up the improvements that are brought about by each successive application of τ_M [in our notation: ρ]. That is, we want a way of going from the improvements that are severally brought about by the various applications of τ_M to the improvements that are collectively brought about by those applications. (Gupta [5], p. 39)

That's the informal idea. It is usually turned into a precise criterion as follows:

At a limit stage μ , one "sums up" the effects of earlier revisions: if the revision process up to μ has yielded a definite verdict on an element, d , (i.e., if d has stabilized in the process) then this verdict is reflected in the μ th hypothesis; (Gupta and Belnap [6], p. 168)

In this, Gupta and Belnap are specifically just considering hypothesis spaces that have the form X^D ,³ i.e. where a hypothesis is a choice of a 'value', or member of X , for each member of D . For example a choice of 'truth value', true or not-true, to each sentence.

The criterion that is usually applied at limit stages is thus:

COHERENCE OF DEFINITE VERDICTS. If d is stably x in $\langle h_\alpha \rangle_{\alpha < \mu}$, then h_μ assigns value x to d .

where we say that d is stably x in a sequence of hypotheses if there is a point before μ after which the value of the element d has settled down on x .⁴

This Coherence of Definite Verdicts is then usually used to provide the limit criterion in the definition of a revision sequence: (we add the specifier 'definite verdicts' because we will generalise the definition)

Definition 1 Fix a hypothesis space X^D and a revision rule ρ . A *definite verdicts revision sequence* is a sequence of hypotheses $\langle h_\alpha \rangle$ where:

- For each α , $h_{\alpha+1} = \rho(h_\alpha)$,
- For each limit ordinal μ , Coherence of Definite Verdicts is satisfied.

In this characterisation, the only kinds of properties that are considered in determining whether the limit stage coheres with the previous stages are the definite

³We use X^D to mean either a D -indexed product of X or to denote the collection of functions from D to X , noting that these two spaces have easy translations back and forth. We stick to the product representation because we then use the product topology on this space.

⁴Formally we say d is stably x if $\{h \mid h(d) = x\}$ is stable in $\langle h_\alpha \rangle_{\alpha < \mu}$ in the sense of Definition 2.

verdicts. However in this paper we will argue that one should have a more encompassing characterisation. We take the limit rule to generally be accounted for by something of the form:

If a property of interest of the hypotheses is brought about by the sequence beneath μ then it should be satisfied at the μ th stage.

To obtain specific criteria we need to specify what properties of the hypotheses are of interest and what it means to be brought about.

We will not specify what the correct collection of properties is; we will just impose certain constraints on legitimate collections and consider the range of limit criteria that this can lead to. For this essay we will be taking an extensional perspective on the notion of a property and equating them with sets of hypotheses. We find the talk of properties helpful for presentation but don't take this to impose any formal constraints.⁵

For simplicity of this paper we keep 'being brought about' to be characterised by stability, as is usual in the revision theory. Informally a property is stable in a sequence if the hypotheses in that sequence all end up having that property. Formally, then:

Definition 2 A set of hypotheses, $A \subseteq H$, (informally we are calling A a property) is *stable* in $\langle h_\alpha \rangle_{\alpha < \mu}$ if there is some $\beta < \mu$ such that for all α with $\beta < \alpha < \mu$, $h_\alpha \in A$.

In fact we have leanings towards counting properties that are merely *nearly stable* to also count as being brought about, where near stability weakens stability by requiring that the property is satisfied from some point onwards *except* for finitely many stages following limit ordinals.⁶ This will only make a difference at ordinals, like ω^2 , which are not of the form $\alpha + \omega$, and can lead to some behaviour at such ordinals that might be thought of as desirable, as discussed on Page 23. The near stability variant is adopted in Campbell-Moore [2], but the focus of this paper is the question of which properties should be considered instead of how best to understand the notion of 'brought about', so we will not further consider the near stability variant.⁷

So the general form of the limit criteria is the following, for an appropriate choice of \mathcal{P} (pronounced 'properties of interest') which is a collection of sets of hypotheses: **COHERENCE OF \mathcal{P}** . If $A \in \mathcal{P}$ and A is stable in $\langle h_\alpha \rangle_{\alpha < \mu}$, then $h_\mu \in A$.

One then defines the corresponding revision sequences:

⁵For example we will talk about any set of hypotheses as a property, even those that aren't governed by a defining formula.

⁶More precisely:

$$\exists \beta < \mu \forall \alpha \begin{matrix} < \mu \\ > \beta \end{matrix} \exists N < \omega \forall n \begin{matrix} < \omega \\ > N \end{matrix}, h_{\alpha+n} \in A.$$

Where, e.g., $\forall \alpha \begin{matrix} < \mu \\ > \beta \end{matrix}$ is a restricted quantifier, so is shorthand for $\forall \alpha (\beta < \alpha < \mu \rightarrow \dots)$.

⁷One further consideration for not here focusing on near stability is that we are not certain that it is quite general enough: near stability allows for ignoring finite stages following limit ordinals, but we might think that some more general *arbitrary* stability which allows one to ignore appropriate bounded (but possibly infinite) regions after ordinals. Such considerations would take us too far away from our main topic.

Definition 3 Fix a hypothesis space H and a revision rule ρ . A \mathcal{P} -revision sequence is a sequence of hypotheses $\langle h_\alpha \rangle$ where:

- For each α , $h_{\alpha+1} = \rho(h_\alpha)$,
- For each limit ordinal μ , Coherence of \mathcal{P} is satisfied.

The usual limit criterion is an instance of this criterion where \mathcal{P} is taken to be just the definite verdicts. But we also can allow for other properties to be taken into account, i.e. consider a more encompassing \mathcal{P} .

There are some properties which are inappropriate to consider in this criterion because they aren't the sorts of properties that we would even want to ensure limits cohere with respect to. Consider a property of a real number: being strictly smaller than 1. Even if a sequence of hypotheses brings about this property, we might not want to require that the property is also satisfied at the limit; for example if the real numbers get closer and closer to 1 we might instead want to have that the limit value is equal to 1.

We thus want to restrict the kinds of properties that can be taken into account; and this will be done by using topological considerations: we will assume that our hypothesis space comes with an associated topology, which will determine which sets are closed and which are open (a set is closed iff its complement is open), and then we restrict choices of properties of interest to those that determine a closed set. We can use a result from topology to see why these are the sets that act nicely under limits: A is closed iff for any h , if h can be approximated arbitrarily well by members of A , (i.e. For each U open with $h \in U$, there is some $h_U \in A \cap U$.) then $h \in A$.⁸

We will thus require that members of \mathcal{P} be closed. This assumption will also play an important role in our proof that there will always be some hypothesis which satisfies the limit criterion, at least with an additional assumption on the underlying hypothesis space: that the topology on it is compact.

A further restriction on appropriate choices of \mathcal{P} that we will make is that we will require that the definite verdicts are in \mathcal{P} so that Coherence of \mathcal{P} will then extend Coherence of Definite Verdicts.

Appropriate choices of \mathcal{P} are therefore given by choosing some collection of propositions of interest which at least contain all definite verdicts, and only contain closed properties. The appropriate limit criteria are then bounded by the weakest criterion, which is the usual one, Coherence of Define Verdicts, and the strongest criterion, Coherence of all Closed Properties:

COHERENCE OF ALL CLOSED PROPERTIES. If A is a closed property and A is stable in $\langle h_\alpha \rangle_{\alpha < \mu}$, then $h_\mu \in A$.

We will call the associated revision sequences Closed Properties Revision Sequences. This limit criterion is in fact already familiar in topology: the hypotheses that satisfy

⁸See, e.g., Willard [12, Theorem 4.7b]

Coherence of all Closed Properties are exactly the *cluster points* of the preceding sequence of hypotheses.⁹

There are some restrictions on appropriate topologies in order to ensure that our proposal works as required. The first requirement is that the topology is compact. We need this to ensure that our limit criterion will always be satisfiable. The second requirement is that the definite verdicts are closed. This is required for there to be any choices of \mathcal{P} that satisfy the restrictions we have imposed on legitimate choices.

These considerations and more detailed discussion of the topology will now be given, but because the details now become a bit more abstract and involved we present them in a separate section that can be safely skipped.

2.1 Topological Considerations

This section presents more abstract topological considerations. As just mentioned, the reader who is less interested in such details can skip to Section 3.

We have used the notion of a closed set to set out appropriate choices of \mathcal{P} to be used in the limit criterion. We thus assumed that our hypothesis space comes with an associated topology to determine which sets are closed or not. A topology is given by choosing some of the sets to be the *open* ones, satisfying certain constraints: \emptyset and H are open and the open sets are closed under finite intersections and arbitrary unions. The closed sets are then defined to be the complements of open sets. One can also directly define a topology by choosing an appropriate collection of sets to be the closed sets, this collection will have to contain \emptyset and H and be closed under finite unions and arbitrary intersections.

There may be different topologies on a single hypothesis space, but hypothesis spaces typically come with a standard topology. For example, the standard topology on $\{t, f\}$ has that all sets are both open and closed, and the standard topology on $[0, 1]$ is the topology containing the fewest open sets where each open interval (a, b) is open. We can give a direct characterisation of the closed sets in this topology: they are $[a, b]$ and anything generated from these by finite unions and arbitrary intersections. Often the hypothesis space will be a product of other spaces, e.g. $\{t, f\}^{\text{Sent}_T}$ where a hypothesis is an assignment of truth values to sentences, in which case one will typically use the product topology where the underlying space(s) are given their standard topology:¹⁰

Definition 4 The *product topology* on $\prod X_i$ is the coarsest topology (i.e. contains the fewest open sets) where for any i_0 and U_{i_0} open in X_{i_0} , $\{(x_i) \mid x_{i_0} \in U_{i_0}\}$ is open. Open sets are thus anything that can be generated from such projections by finite intersections and arbitrary unions.

⁹The μ -length sequence would be called a net in topology. See e.g. Willard [12, Definition 11.3].

¹⁰A case where we wouldn't do that is if, for example, the different concepts added have an intimate connection to one another; for example if one added countably many predicates: $P_{\geq r}$ for each rational r , then one would want that the topology on the hypothesis space $(\{t, f\}^{\text{Sent}_T})^{\mathbb{Q}}$ is appropriately connected to the topology $\mathbb{Q}^{\text{Sent}_T}$.

We can equivalently think of it in terms of closed sets: A set is closed if it is the projection of a closed set (i.e. $A = \{(x_i) \mid x_{i_0} \in C_{i_0}\}$ for some C_{i_0} closed in X_{i_0}), or generated from these by finite unions and arbitrary intersections.

We should note an alternative picture one might have: we might start with an idea of what properties we want to consider in this criterion then determine a topology where all those are closed. For this latter approach it will be very important to check that the topology is then compact. See footnote 18 for a case where compactness will form a barrier to a situation where one might want to follow this approach.

The limit criteria we have proposed will only be legitimate if we impose some restrictions on the topologies, and it is these restrictions that we turn to now. The first restriction is that the topology is compact:

Definition 5 A space is *compact* if whenever there is a collection of open sets that cover the whole space (i.e. every point lies in one of the sets) there is some finite sub-collection that also covers the space.

This restriction is required to ensure that our limit criteria are satisfiable, i.e. that there is there always a choice of h_μ that satisfies Coherence of \mathcal{P} , at least whenever \mathcal{P} only contains closed sets:

Proposition 6 *If the topology on H is compact and \mathcal{P} contains just closed sets, then for any sequence $\langle h_\alpha \rangle_{\alpha < \mu}$, there is some h_μ satisfying the Coherence of \mathcal{P} criterion.*

Proof We first note the standard result about compact spaces (which is equivalent to compactness, showing that our result is rather tight):¹¹ If \mathcal{A} is a family of closed sets with the finite intersection property, i.e. for every finite $\mathcal{F} \subseteq \mathcal{A}$, $\bigcap \mathcal{F} \neq \emptyset$, then $\bigcap \mathcal{A} \neq \emptyset$.

Let

$$\mathcal{A} := \{A \in \mathcal{P} \mid A \text{ is stable in } \langle h_\alpha \rangle_{\alpha < \mu}\}.$$

Note that by the assumption on \mathcal{P} this is a family of closed sets. We will show that \mathcal{A} has the finite intersection property and by the above observation we will be able to deduce that $\bigcap \mathcal{A} \neq \emptyset$, i.e. there is some $h_\mu \in \bigcap \mathcal{A}$. Such an h_μ will be as required.

Let \mathcal{F} be finite $\subseteq \mathcal{A}$. We know that for each $A \in \mathcal{F}$ there is some $\beta_A < \mu$ such that for all α with $\beta_A < \alpha < \mu$ we have $h_\alpha \in A$. Therefore for any $\alpha \geq \max\{\beta_A \mid A \in \mathcal{F}\}$ we have $h_\alpha \in A$, and thus $h_\alpha \in \bigcap \mathcal{F}$. So we can see that $\bigcap \mathcal{F} \neq \emptyset$. We have therefore shown that \mathcal{A} has the finite intersection property, as required. \square

So whenever the topology on H is compact, the limit criterion can be safely applied. There are many compact spaces, for example $\{t, f\}$ with the discrete topology,

¹¹The proof of this equivalence is very easy; it follows from the observation that $\bigcap \mathcal{F} = \emptyset$ iff $\{H \setminus B \mid B \in \mathcal{F}\}$ is a cover for H . The result can be found in any introductory topology text book, e.g. Willard [12, Theorem 17.4].

which says that all sets are closed, or $[0, 1]$ with the standard topology.¹² An example of a space that is not compact is $(0, 1)$. One can account for such non-compact spaces by adding extra elements at limits,¹³ but we will not further discuss this.

Furthermore, spaces that are ‘built from’ compact spaces are also compact. This is made precise in Tychonoff’s theorem, which says that the product of compact spaces is compact, if it is endowed with the product topology.

Theorem 1 (Tychonoff’s Theorem) *If K_i is associated with a compact topology for each $i \in I$, then the product topology on $\prod_i K_i$ is compact.*

This allows us to see that, for example, $[0, 1]^{\text{Sent}_{\mathcal{P}}}$ is compact.

Although we now know that in a compact space there will always be limit hypotheses that satisfy the limit criterion, there may be many such hypotheses, and since the proof of existence relied on Tychonoff’s theorem there may not be any way to ‘construct’ a limit stage satisfying the criterion (Tychonoff’s theorem in its general form is equivalent to the axiom of choice). However we don’t think this is a problem: here the limit criterion should be viewed as part of a definition of what it is to be a revision sequence. It is a test of whether a given sequence of hypotheses is a revision sequence or not. For some further comments see Section 3. Sometimes one talks as if the revision theory is in the business of constructing revision sequences, but in fact one typically isn’t interested in what is obtained in a particular revision sequence but is instead interested in what happens in *all* revision sequences for the purposes of reading off categorical information. Even if one wants to show that certain features are *not* brought about by all revision sequences, one does not need to construct a particular revision sequence with that feature, one instead just needs to provide a proof that there are some revision sequences with that feature. This can be done by use of the following result:

Proposition 8 *Suppose A is closed and cofinal, i.e. its complement is not stable in $\langle h_\alpha \rangle_{\alpha < \mu}$. Then there is some hypothesis in A which satisfies Coherence of all Closed Properties.*

The proof is just a slight extension of Proposition 6, see Campbell-Moore [2, Lemma 5.B.1] for further details.

We need to make one further restriction on the hypothesis spaces where this criterion can be applied: We said earlier that appropriate choices of \mathcal{P} were ones that contain only closed sets and at least contain all the definite verdicts. For this to be possible we require that all definite verdicts are closed. This criterion is a little bit subtle because we haven’t explicitly said what a definite verdict is in a general hypothesis space, we only looked at the example of a space X^D . Without further information about the form of the hypothesis space we might just say that a definite verdict is a

¹²More generally any closed and bounded set of reals is compact.

¹³E.g. in this case we might add 0 and 1. Whenever the topology is Tychonoff (which is quite weak; for example all metric spaces are Tychonoff) one can find a ‘compactification’.

singleton of a hypothesis, or in property terminology: the property of being equal to that fixed hypothesis. If we take definite verdicts to be singletons then we can ensure that they are closed by asking that the topology be T_1 . A topology is called T_1 if all singletons are closed. Most natural topologies are T_1 , in particular all the topologies mentioned in this paper are T_1 . If we know more about the particular hypothesis space we might be able to say something more restrictive about what properties are definite verdicts. In the case of the hypothesis space X^D we said a definite verdict was a specific assignment of some member of X to a member of D . The definite verdicts are thus the sets of the form $\{h \in X^D \mid h(d) = x\}$ and we then need these to be closed in order to have legitimate choices of collections of properties of interest. Such sets are closed if the topology on X is T_1 and the topology on X^D is the product topology (or finer than the product topology, i.e. whenever a set is closed in the product topology it is also closed in this topology).

In conclusion, this framework for limit criteria is legitimate whenever the hypothesis space has a compact topology where all definite verdicts are closed.

3 Limits in the Revision Theory of Truth

In this section we'll consider how this new proposal affects the revision theory of truth.

The revision theory of truth can be used to say something about a language with a self-referential truth predicate. In such a language we will have sentences like the liar sentence that says of itself that it is not true. A revision theory will be a transfinite sequence of interpretations of the truth predicate that is governed by the revision rule ρ based on taking the Tarskian T-biconditionals, " φ is true" iff φ , to be a *definitional* equivalence. In such a revision sequence, the liar sentence will continue to change its truth value, flipping in and out of the truth predicate in successive revision stages.

To present this revision theory and the effect of our suggested limit criteria we first present some standard technicalities and terminology required to start the discussion of the revision sequence.

Setup Let \mathcal{L} be the language of Peano Arithmetic,¹⁴ and let \mathcal{L}_T extend this by adding a unary predicate T . We shall use Sent_T to denote the set of sentences of the language \mathcal{L}_T .

The construction works over some fixed background model of \mathcal{L} , which we will fix as the standard model of arithmetic, \mathbb{N} .

We assume some coding of Sent_T to natural numbers which is recursive and one-to-one. For $\varphi \in \text{Sent}_T$, we let $\ulcorner \varphi \urcorner$ denote the numeral corresponding to the code of φ .¹⁵ We will often identify a sentence with its code, but this should not lead to any confusion.

¹⁴Everything that is said could equally be applied to languages that also contain additional vocabulary.

¹⁵The numeral of n is denoted \bar{n} and it corresponds to the expression $\overbrace{S(\dots S(0)\dots)}^n$.

PA is the theory of Peano Arithmetic, and we will always assume that induction axioms are extended to the whole language.

A revision sequence is a sequence of hypotheses, which are proposed extensions of the truth predicate. Our hypotheses are subsets of the domain, or actually of Sent_T as that's all we care about. A hypothesis is thus some $T \subseteq \text{Sent}_T$, and the hypotheses space is $\wp(\text{Sent}_T)$.

A revision sequence will be a sequence of such hypotheses, i.e. we have for each ordinal α some $T_\alpha \subseteq \text{Sent}_T$, which should satisfy a successor clause saying that the successors are obtained by an appropriate revision rule and a limit clause saying that limits satisfy the imposed criterion. The revision rule that is used for truth is based on the Tarskian T-biconditionals and says:

$$\varphi \in \rho(T) \text{ iff } (\mathbb{N}, T) \models \varphi.$$

Where (\mathbb{N}, T) is given by expanding the model \mathbb{N} of \mathcal{L} by using T as the interpretation of the predicate T .

In a sequence following this revision rule the liar sentence, which is a sentence λ where $\text{PA} \vdash \lambda \leftrightarrow \neg T^\ulcorner \lambda \urcorner$, will have its truth value continually switched i.e. $\lambda \in \rho(T)$ iff $\lambda \notin T$.

A limit criterion is given by a choice of \mathcal{P} , which can only contain closed properties. Which properties are closed in this hypothesis space $\wp(\text{Sent}_T)$? I.e. what is the topology on this space? The topology that we will use is the one generated by taking all properties of the form $\varphi \in T$ and $\varphi \notin T$ to be open (and thus also closed). The closed sets are then anything generated from such sets by finite unions and arbitrary intersections. This is the natural topology on the powerset.¹⁶ Here are some examples of closed sets:

Example 9 The following are closed sets:

- $\{T \mid \varphi_0 \in T\}$,
- $\{T \mid \varphi_0 \notin T\}$,
- $\{T \mid \neg\varphi_0 \in T \text{ iff } \varphi_0 \notin T\}$,
- $\{T \mid \text{For all } \varphi, \neg\varphi \in T \text{ iff } \varphi \notin T\}$,
- $\{T \mid T \text{ is a maximally consistent set of sentences.}\}$
- Any property that is describable using propositional connectives and sentences of the form $\varphi \in T$.

The following are *not* closed:

- $\{T \mid \varphi_0 \in T \text{ or } T^\ulcorner \varphi_0 \urcorner \in T \text{ or } T^\ulcorner T^\ulcorner \varphi_0 \urcorner \urcorner \in T \dots \}$,
- $\{T \mid T \text{ is } \omega\text{-consistent}\}$, i.e.:

$$* \quad \text{If } \exists x \in N \varphi(x) \in T \text{ then there is some } n \in \mathbb{N} \text{ with } \varphi(\bar{n}) \in T.$$

¹⁶The standard topology on $\{t, f\}^{\text{Sent}_T}$ is the product topology (see Definition 4), where $\{t, f\}$ is given the discrete topology (where all sets are both open and closed). And since $\wp(\text{Sent}_T)$ is isomorphic to this we can take the corresponding topology on $\wp(\text{Sent}_T)$.

These can be shown by using the following characterisation of closed sets, which is a special case of the general criterion that A is closed iff for any h , if h can be approximated arbitrarily well by members of A (i.e. for each U open with $h \in U$, there is some $h_U \in A \cap U$) then $h \in A$. The special case results from the fact that a basis of open sets in $\wp(\text{Sent}_T)$ is the collection of assignments of definite verdicts to finitely many sentences (i.e. the collection of properties of the form $(\bigwedge_{\varphi \in \Delta_0} \varphi \in T) \wedge (\bigwedge_{\varphi \in \Delta_1} \varphi \notin T)$, for Δ_0 and Δ_1 finite).

Proposition 10 *A property $A \subseteq \wp(\text{Sent}_T)$ is closed iff: for each hypothesis $T \in \wp(\text{Sent}_T)$; if T is such that for every finite set of sentences, Δ , there is some $T_\Delta \in A$ where for each $\varphi \in \Delta$, $\varphi \in T_\Delta$ iff $\varphi \in T$, then $T \in A$.*

This topology is compact by Tychonoff's theorem so we get that if \mathcal{P} contains only closed properties then Coherence of \mathcal{P} is satisfiable.

What results does this lead to? As long as all definite verdicts are in \mathcal{P} this criterion will ensure that definite verdicts that are stable beneath μ should be agreed with by T_μ . But by choosing a more encompassing \mathcal{P} we can impose additional requirements on T_μ . For example if we consider the strongest limit criterion we get the following:

Proposition 11 *Suppose we have a Closed Properties Revision Sequence for Truth. Then each T_α will be maximally consistent (including at the limit ordinals).*

Proof The key component to this is the observation that being maximally consistent is closed, which one can do using the characterisation in Proposition 10.¹⁷

Since the interpretation is maximally consistent at successors because of the choice of revision rule, the result follows by transfinite induction. \square

We have not plugged in the requirement that we wanted the limit stages to be maximally consistent, instead we were able to derive this from a more general requirement. Furthermore, this consequence seems nice because one might think that being maximally consistent is a minimal requirement on an appropriate extension of the truth predicate and now all the hypotheses in a revision sequence will

¹⁷Suppose we have some T and for each finite set of sentences, Δ , some maximally consistent T_Δ where for each $\varphi \in \Delta$, $\varphi \in T_\Delta$ iff $\varphi \in T$. We need to show that T is also maximally consistent. Each finite subset of T is consistent (because it corresponds to some consistent T_Δ) so by the compactness theorem T is also consistent. To show it is maximal consider $\Delta = \{\varphi, \neg\varphi\}$. Then $\neg\varphi \in T_{\{\varphi, \neg\varphi\}}$ iff $\varphi \notin T_{\{\varphi, \neg\varphi\}}$ by maximality of $T_{\{\varphi, \neg\varphi\}}$, so since $T_{\{\varphi, \neg\varphi\}}$ and T agree on these sentences we also have $\neg\varphi \in T$ iff $\varphi \notin T$. T is thus maximally consistent.

satisfy that requirement. An analogous argument to show that the limit hypotheses are ω -consistent will fail because that property is not closed.¹⁸

A consequence of being maximally consistent is that we require that either $\lambda \in T_\omega$ or $\neg\lambda \in T_\omega$, but the limit criterion won't tell us which to put in T_ω . There are choices of T_ω that satisfy Coherence of all Closed Properties where $\lambda \in T_\omega$, and ones where $\neg\lambda \in T_\omega$.¹⁹ This highlights the fact that what we are presenting is not a limit *rule* but instead a limit *criterion*.

Different accounts of what to do at limits have been considered in the literature on the revision theory. There is a substantial difference that is relevant to us here between Herzberger's proposal [8] and Belnap's [1] (which was then adopted in Gupta and Belnap [6]. Gupta's initial proposal [5] falls in line with Herzberger's on this issue). Herzberger says that one should put a sentence in T_μ only if it is stably in the extension of the truth predicate up to μ . He is seeing the limit stage to be defined from the previous stages. The picture is that we start at the beginning with an initial hypothesis (often taken to be $T_0 = \emptyset$) and perform the process of constructing a revision sequence. Gupta's original revision theory of truth was similar with regard to the picture of the process of constructing the revision sequence, but differs from Herzberger's in that it consults the initial hypothesis to decide what to do with unstable sentences at limit ordinals. Belnap [1], however, argues that this is too restrictive: why should the same thing be done to unstable sentences at each limit ordinal. He instead suggests that one is allowed to do different things to unstable sentences at different ordinals, and doesn't specify anything more about how they should be treated. His limit policy provides a criterion rather than a rule.

The feature of our limit policy that it tells you that either $\lambda \in T_\mu$ or $\neg\lambda \in T_\mu$ but won't tell you which, means that our revision theory falls in line with Belnap's on this issue: the revision sequence isn't determined by a choice of initial hypothesis. If one wants to see it as a process of constructing a revision sequence one has to make a choice at each limit ordinal.²⁰ We view what we are doing as providing a criterion for when a sequence of hypotheses is a revision sequence, i.e. a definition of what it is to be a revision sequence.

We thus see our proposal as an extension of Belnap's picture of giving a limit criterion and a definition of what it is to be a revision sequence. In fact Belnap's definition of a revision sequence is the special case of ours where only definite verdicts are considered. We don't see why the restriction to just definite verdicts in Belnap's

¹⁸One might think that being ω -consistent is a property that one would like to take over to limits. To allow for that one might try to alter the topology to also ensure that this is closed; however one will result in a non-compact topology, where the criteria are thus inapplicable. This can be seen by an application of McGee's result.

¹⁹Which can be shown using Proposition 8

²⁰There is one way which ours is worse than Belnap's: for ours the available options for choices at limits will depend on what has happened in the revision sequence up to that stage, whereas for Belnap's one can make the choices before the revision process is underway.

proposal is motivated and therefore we have developed a framework where more general properties can also be considered.²¹

4 A Hypothesis Space with Reals

We will now present a special case of a hypothesis space that involves real numbers. This space is particularly interesting with regard to this generalised limit criteria because the hypothesis space has a natural topology associated with it where convergence plays an important role. In the topology on $\{t, f\}$ what it was to converge was just to stabilise on some particular truth value, whereas convergence in $[0, 1]$ can happen without the sequence ending up always assigning a particular real number. Using the usual limit criterion therefore doesn't work as intended in such a space.

In the truth case we see the general framework of limit criteria as interesting, but they may turn out to be optional alternatives to just considering definite verdicts. Ultimately philosophical motivations should determine the appropriate collection of properties, and for truth we acknowledge that this might end up just being the definite verdicts, resulting in the usual limit criterion. However, for hypothesis spaces where convergence plays an important role, the general framework is essential as the usual choice of \mathcal{P} as the definite verdicts is not a reasonable option.

4.1 The Limit Criteria in this Hypothesis Space

We will here consider the application of these criteria if we have a hypothesis space involving reals. In order to apply the tools of revision theory we will be interested in concepts that attach a real number to all sentences of a language, including sentences that say something about that very concept. The case we'll be talking about here is the notion of probability, but everything that we say in this section would equally apply to other notions such as degrees of truth.

The setup is mostly the same as in the previous section where we considered the revision theory of truth, with a few alterations as follows:

Setup We take the base language here to be the language of real ordered closed fields (with $+$, $-$, \times , 0 , 1 , $<$) and a natural number predicate N . We add to this a probability function symbol P , obtaining \mathcal{L}_P . The sentences are denoted Sent_P . Our

²¹There is also the restriction in the literature in revision theory of truth on being a fully varied revision sequence [4]. This adds to the definition of a revision sequence (as in 'definite verdicts revision sequence') the requirement that if a hypothesis coheres with the whole revision sequence $\langle h_\alpha \rangle_\alpha$ with regard to definite verdicts, then that hypothesis is cofinal in the sequence. We note that a similar restriction could be applied to our revision theory:

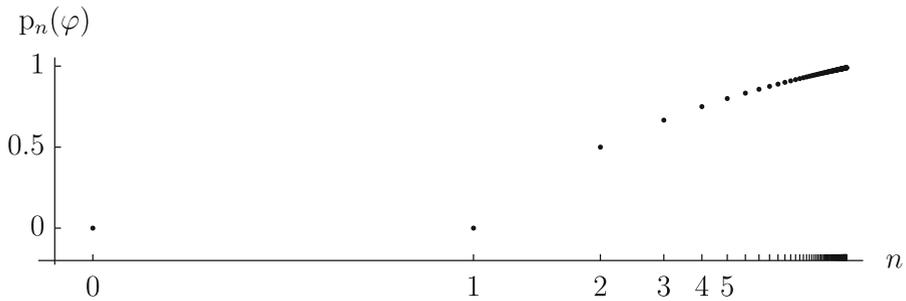
Definition 12 A \mathcal{P} fully varied revision sequence is a \mathcal{P} revision sequence where if a hypothesis h satisfies Coherence of \mathcal{P} (where the ordinal length sequence, $\langle h_\alpha \rangle_\alpha$, is considered), then h is cofinal in the sequence (i.e. $H \setminus \{h\}$ is not stable).

The consequences of such a proposal remain to be investigated.

base model of the language without probability will be \mathbb{R} , the intended interpretation of this language, with a restriction to N recovering \mathbb{N} . Our coding is now taken to code the language \mathcal{L}_P . We thus have a type-free probability notion where there are sentences that talk about their own probability.

A revision sequence for probability is given by a sequence of hypotheses, which are interpretations of P . So a hypothesis is some p , where p assigns to each sentence a number between 0 and 1, i.e. $p : \text{Sent}_P \rightarrow [0, 1]$. The space of all hypotheses is then $[0, 1]^{\text{Sent}_P}$. A revision sequence will be a sequence of hypotheses, p_α .

Coherence of Definite Verdicts is not sufficient in this hypothesis space: Suppose one has a revision sequence up to ω , where $p_n(\varphi) = (n - 1)/n$. This can be graphed as follows:



So if the revision sequence is as described, the probability of φ converges to 1. It thus seems natural to want to require that the ω th hypothesis puts the probability of φ to be equal to 1. However, there is no definite verdict regarding the probability of φ which is brought about by the sequence. Instead the probability just gets closer and closer to 1. This means that using Coherence of Definite Verdicts, the probability value at stage ω could be put to anything at all. This would result in a great loss of valuable information about the probability of φ . In the specific revision rule that we will later consider it will be of great importance to impose such a constraint because otherwise simple sentences like $\top \ulcorner 0 = 0 \urcorner$ might have their probability unconstrained at ω .

This is why we don't just focus on definite verdicts but more general properties. We can get this result that $p_\omega(\varphi) = 1$ if we also consider the properties of the form $r \leq p(\varphi) \leq s$ in our limit criterion. This is because for every real number $\epsilon > 0$, the property $1 - \epsilon \leq p(\varphi) \leq 1$ is stable beneath ω : the values converge to 1, so they get arbitrarily close to 1 and thus however small we make this interval it will nonetheless hold from some point onwards beneath ω . As long as these closed intervals are being taken into account, Coherence of \mathcal{P} will then require $1 - \epsilon \leq p_\omega(\varphi) \leq 1$ for each $\epsilon > 0$, and thus it must be that $p_\omega(\varphi) = 1$.

This example also highlights why we have to impose some restrictions on the class of properties considered: we cannot consider the property $p(\varphi) < 1$ along with all the properties $r \leq p(\varphi) \leq s$ as there is no limit hypothesis which will then satisfy the constraint. The restriction to closed properties is natural because those are the ones that act nicely under convergence. $p(\varphi) < 1$ is thus a property that we will not want to consider in our limit criterion.

To specify which sets are closed we require a topology on the hypothesis space. There is a very natural choice for this hypothesis space, namely the product topology, as defined in Definition 4, where $[0, 1]$ is endowed with the standard topology.

As before, we can describe the closed sets by using the notion of approximation. But since we now have an obvious notion of convergence we state it in that format.²²

Proposition 13 *A property $A \subseteq \wp(\text{Sent}_\top)$ is closed iff:*

*for each hypothesis $[0, 1]^{\text{Sent}_\mathbb{P}}$;
if \mathbb{p} is such that there is a sequence of hypotheses $\langle \mathbb{p}_n \rangle$ (not necessarily following the revision rule) that converges to \mathbb{p} ,*

*i.e. for all $\varphi \in \text{Sent}_\mathbb{P}$ and for all $\epsilon > 0$, there is some $m \in \mathbb{N}$ such that for all $n > m$, $|\mathbb{p}_n(\varphi) - \mathbb{p}(\varphi)| < \epsilon$.
and where each $\mathbb{p}_n \in A$, then $\mathbb{p} \in A$.*

Here are some examples of properties that are closed and not:

Example 14 The following are closed properties:

- $\{\mathbb{p} \mid \mathbb{p}(\varphi) = 1/2\}$,
- $\{\mathbb{p} \mid \mathbb{p}(\varphi) = 1 \text{ or } \mathbb{p}(\varphi) = 0\}$,
- $\{\mathbb{p} \mid \mathbb{p}(\varphi) = 1 - \mathbb{p}(\psi)\}$,
- $\{\mathbb{p} \mid \mathbb{p}(\varphi) \cdot \mathbb{p}(\psi) \geq 1/2\}$,
- $\left\{ \mathbb{p} \mid \text{If } \mathbb{p}(\psi) > 0, \text{ then } \frac{\mathbb{p}(\varphi \wedge \psi)}{\mathbb{p}(\psi)} = 1/2 \right\}$,
- $\{\mathbb{p} \mid \mathbb{p} \text{ is a finitely additive probability function}\}$, i.e.:
 - * $\mathbb{p}(\top) = 1$ for a tautology \top ,
 - * $\mathbb{p}(\varphi) \geq 0$ for all φ ,
 - * $\mathbb{p}(\varphi \vee \psi) = \mathbb{p}(\varphi) + \mathbb{p}(\psi)$ if φ and ψ are logically incompatible.

The following are *not* closed:

- $\{\mathbb{p} \mid \mathbb{p}(\varphi) \in \mathbb{Q}\}$,
- $\{\mathbb{p} \mid 1/2 < \mathbb{p}(\varphi)\}$,
- $\{\mathbb{p} \mid \mathbb{p}(\varphi) \neq 1/2\}$,
- $\{\mathbb{p} \mid \mathbb{p} \text{ is an } \mathbb{N}\text{-additive probability function}\}$, i.e.:
 - * \mathbb{p} is a finitely additive probability function, and
 - * $\mathbb{p}(\exists x \in N \varphi(x)) = \lim_n \mathbb{p}_n(\varphi(\bar{0}) \vee \dots \vee \varphi(\bar{n}))$

We have already shown when the limit rule is satisfiable: if the topology is compact (Proposition 6). The topology we are using on $[0, 1]^{\text{Sent}_\mathbb{P}}$ is in fact compact. Moreover in this topology all definite verdicts, i.e. all $\mathbb{p}(\varphi) = r$, are closed, so there are appropriate choices of \mathcal{P} .

²²This variant way of describing closed sets does require that \mathcal{L} be countable to ensure that sequentially closed and closed coincide.

4.2 Features of Coherence of All Closed Properties

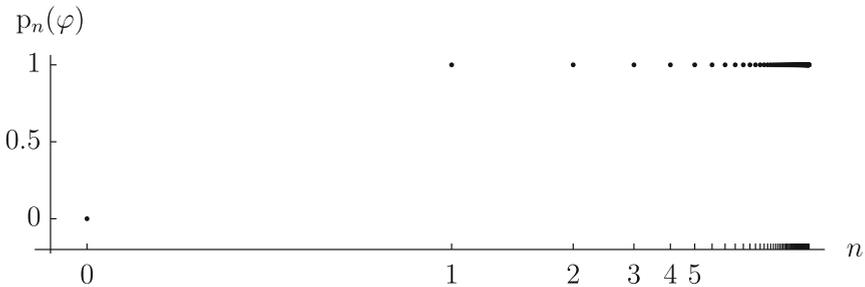
We can now present examples of how this will work. We will initially focus on the strongest limit criterion, Coherence of all Closed Properties, and see what features this leads to as this is the strongest appropriate limit criterion. We will then later discuss which restrictions of \mathcal{P} get us the different features.

Example 15 Here are some examples of how the criterion Coherence of all Closed Properties works with the hypothesis space $[0, 1]^{\text{Sent}_{\mathcal{P}}}$.

FIXES ON A VALUE:

If up to μ the probability of φ ends up fixing on a value r , then $p_{\mu}(\varphi) = r$.

E.g. Suppose $p_n(\varphi)$ are as follows:



The earlier stages “bring about” the property that $p(\varphi) = 1$ because it is stable beneath ω . It is also closed, so the limit hypothesis p_{ω} will also have to satisfy this property, i.e. $p_{\omega}(\varphi) = 1$.

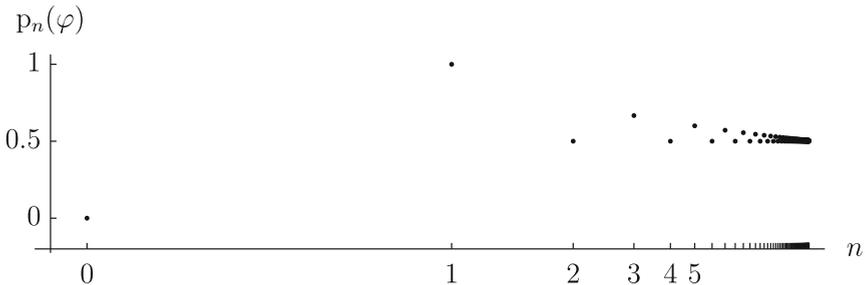
CONVERGES TO A VALUE:

If the probability of φ converges to r beneath μ , i.e.:

$$\text{for all } \epsilon > 0 \text{ there is some } \beta < \mu \text{ such that for all } \alpha \text{ with } \beta < \alpha < \mu, \\ |p_{\alpha}(\varphi) - r| < \epsilon$$

then $p_{\mu}(\varphi) = r$.

E.g. Suppose $p_n(\varphi)$ are as follows:

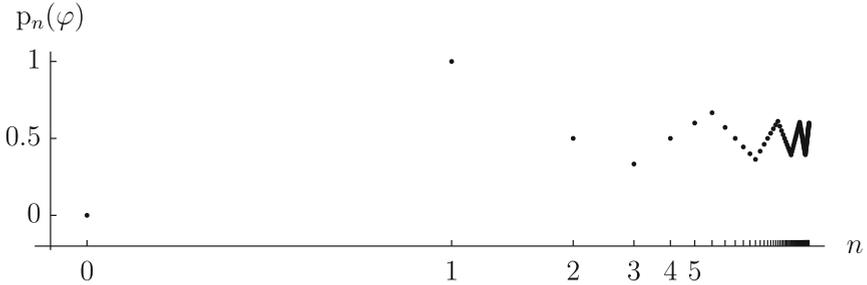


These converge to $1/2$. So by the argument in Section 4.1, $p_{\omega}(\varphi) = 1/2$.

IN AN INTERVAL:

If the probability of φ ends up always (roughly) being in some interval, then the limit probability will also be in that interval.

E.g. Suppose $p_n(\varphi)$ are as follows:

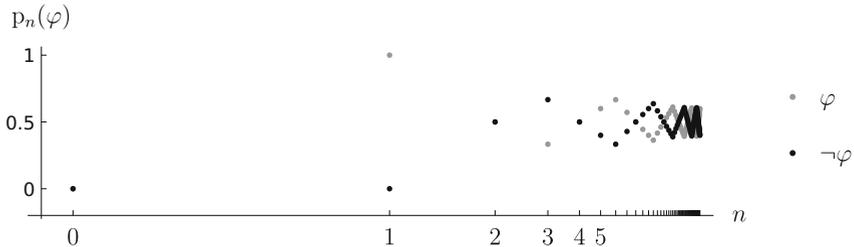


Except for the first few values, the probability of φ beneath ω is always between $0.4 - \epsilon$ and $0.6 + \epsilon$ but does not converge to any value. Our limit constraint requires that $0.4 \leq p_\omega(\varphi) \leq 0.6$ for similar reasoning as to in CONVERGES TO A VALUE.

RELATIONSHIPS:

If the probability of φ ends up always being one minus the probability of ψ beneath μ , then $p_\omega(\varphi) + p_\omega(\psi) = 1$.

E.g. Suppose $p_n(\varphi)$ and $p_n(\neg\varphi)$ are as follows:

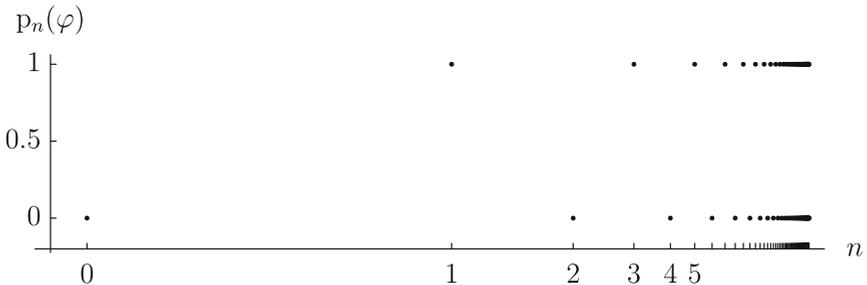


For every stage, except for the first stage, $p_n(\varphi) + p_n(\neg\varphi) = 1$, and $p(\varphi) + p(\neg\varphi) = 1$ is closed as mentioned in Example 14. We will therefore also have that $p_\omega(\neg\varphi) + p_\omega(\varphi) = 1$.

NON-CONVEX:

If the probability of φ ends up always being either a or b beneath μ , then $p_\mu(\varphi)$ is either a or b .

E.g. Suppose $p_n(\varphi)$ are as follows:



For every stage, either $p_n(\varphi) = 0$ or $p_n(\varphi) = 1$. Since $p(\varphi) = 0 \vee p(\varphi) = 1$ is closed, it must also be that $p_\omega(\varphi) = 0$ or $p_\omega(\varphi) = 1$.

We have now seen some examples of how such a limit rule will work. Here is a very important consequence of Coherence of all Closed Properties, which is the probabilistic analogue of our observation that Coherence of all Closed Properties in the truth revision sequence implied that limit interpretations of truth were maximally consistent.

Proposition 16 *Suppose we have a Closed Properties Revision Sequence. If p is a finitely additive probability function at each successor ordinal then p will be a finitely additive probability function at all ordinals.*

Proof As in Proposition 11 using the fact that being a finitely additive probability function is closed. One can observe that this is closed by using Proposition 13 □

An analogue of ω -consistency for probability is:

$$p(\exists x \in N \varphi(x)) = \lim_n p_n(\varphi(\bar{0}) \vee \dots \vee \varphi(\bar{n})),$$

which was called σ -additivity in Leitgeb [10] but which we call \mathbb{N} -additivity. Just as for ω -consistency, this property is not closed so can't be taken over to limits.²³

4.3 Which Choices of \mathcal{P} Give Which Features

In the previous section we considered features of the limit criterion when Coherence of all Closed Properties was used. However one might wish to consider possible weakenings of this criterion by modifying \mathcal{P} . One reason to do this is that one doesn't take think certain behaviour should be imposed on limits, alternatively one may wish to satisfy properties which are inconsistent with such strong limiting behaviour.²⁴ We shall therefore here present some options for how one might weaken the limit rule.

The minimal choice of \mathcal{P} is to just consider all definite verdicts. This would obtain the limiting behaviour in `FIXES ON A VALUE` but not even that in `CONVERGES TO A VALUE`. We therefore suggest that this will usually not be a sufficient constraint on the limit probabilities.

To obtain the limiting behaviour in `CONVERGES TO A VALUE` one needs \mathcal{P} to contain all properties $r \leq p(\varphi) \leq s$. This would not only obtain the behaviour of `FIXES ON A VALUE` and `CONVERGES TO A VALUE` but also that of `IN AN INTERVAL`. This constraint is then equivalent to:

$$\liminf_{\alpha < \mu} p_\alpha(\varphi) \leq p(\varphi) \leq \limsup_{\alpha < \mu} p_\alpha(\varphi)$$

²³In fact for the relative frequency revision sequence one can see that the probabilities at stages $\alpha + \omega$ must fail to be \mathbb{N} -additive by a McGee style argument as presented in Leitgeb [9].

²⁴For an example of that see Campbell-Moore [2, Section 5.3.3].

We believe that this proposal is a significant improvement over the definite verdicts alternative as it is strong enough to obtain desirable features like convergence.²⁵

The behaviour in RELATIONSHIPS is where we see the effect of considering the relationships between different sentences. This would not be obtained if we were only interested in properties of the form $r \leq p(\varphi) \leq s$ but is obtained if we are also interested in properties that say something about how the probability of φ relates to the probability of $\neg\varphi$. It is this feature that allows us to *derive* that the limit interpretations of \mathbf{P} will be probabilistic.

Finally, consider the behaviour in NON-CONVEX. It is debatable whether this behaviour is desirable for the notion of probability. If the probabilities of φ flip between 0 and 1 then we might think that an appropriate way of summing this up would be to take the probability of φ to be $1/2$. We could avoid this behaviour by imposing the constraint that the properties of interest be convex, i.e. if two probability functions satisfy A then so do any weighted averages of them.²⁶ A property like $p(\varphi) = 0 \vee p(\varphi) = 1$ would thus not be convex, and so we could rule this out from being taken over to the limits. Using such a restriction to convex properties we would nonetheless be able to derive that p is a finitely additive probability function as that is a convex property.

4.4 Relative Frequencies Revision Rule

We have so far been talking about how the limit criteria would work if one has the hypothesis space $[0, 1]^{\text{Sent}_{\mathcal{P}}}$. We will here present an example of a revision rule in such a hypothesis space.

We will present here a revision rule based on relative frequencies that originates in Leitgeb [10]. The idea is that the next stage probability should measure how often the sentence was satisfied in the preceding sequence of interpretations. It is thus based on an idea of frequency of truth. For example if one also has a truth predicate in the language one can get that the probability of the liar sentence converges to $1/2$ through the finite stages of this process. The notion of probability that this revision step develops would have to be understood as something like a *semantic probability*, measuring how true a sentence is. It is limited to assigning non-trivial probabilities only to sentences that themselves talk about the notions that are captured in a revision sequence, such as truth or probability. This is because the truth of sentence of the base language, such as whether a coin lands heads, will be fixed in the revision sequence so it's frequency of truth will either be 1 or 0 – this account of probability is thus not one that can assign non-trivial probabilities to such usual cases. The main reason we present this revision step here is a pragmatic one: it is an easy case where we can see nice examples of convergence with a hypothesis space that remains rather simple. If we want to use the revision theory to develop an account of probability as degrees of

²⁵This restriction doesn't get the feature that the limit stages must be finitely additive probability functions. If one thinks that is a desirable constraint one could add it in this framework by also requiring that \mathcal{P} contains the property: p is a finitely additive probability function.

²⁶More carefully: if $p, p' \in A$ then for any $r \in [0, 1]$, then p'' with $p''(\varphi) = r \cdot p(\varphi) + (1 - r) \cdot p'(\varphi)$ is $\in A$.

belief we might have to plug additional information into the revision sequence and make our hypothesis space more detailed, for example by taking the hypotheses to be assignments of values to sentences in each possible world, so the hypothesis space is $([0, 1]^{\text{Sent}_p})^W$, where W is the collection of possible worlds. Closed properties will then include ones involving relationships between interpretations at different worlds, for example λ is in the extension of T at w_0 iff it is in the extension at w_1 . See Campbell-Moore [2, Section 5.3] for further details. To keep things simple, then, we will here focus on the revision rule developed by Leitgeb. For more comments on the interpretation of such a revision rule see Campbell-Moore et al. [3].

Leitgeb uses a sort of revision procedure to prove the consistency of a set of principles including Probabilistic Convention T which says $p(\varphi) = p(T^\top \varphi^\top)$. To construct his model Leitgeb starts by constructing an ω -length revision sequence providing interpretations for truth and probability and then his model is found by taking a ‘limit’ of these interpretations. We shall use his proposal to tell us how to interpret the finite stage probabilities and although our limit stage differs in technical details, it will result in something very close to Leitgeb’s ω -stage definition.

Leitgeb’s finite stage construction says that the n th-stage hypothesis should say that the probability of φ is the relative frequency of φ being satisfied in the sequence of interpretations leading up to that stage, i.e.,²⁷

$$p_n(\varphi) = \frac{\#\{i \in \{0, \dots, n - 1\} \mid (\mathbb{R}, p_i) \models \varphi\}}{n}$$

For a probabilistic liar sentence, π (with $\text{PA} \vdash \pi \leftrightarrow P^\top \pi^\top \leq 1/2$), we then obtain the probabilities:

	0	1	2	3	4	5	6	7	...
$p_-(\pi) = ?$	0	1	1/2	2/3	2/4	3/5	3/6	4/7	...
$(\mathbb{R}, p_-) \models \pi ?$	✓	✗	✓	✗	✓	✗	✓	✗	...

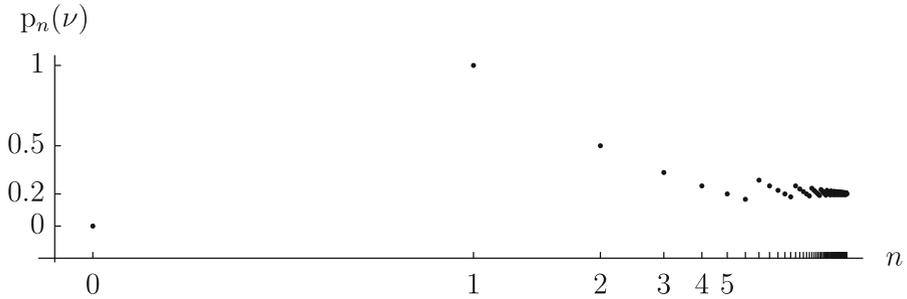
These probability values are represented in the diagram from CONVERGES TO A VALUE on Page 17. In fact a number of the examples we used when presenting the features of the limit criterion earlier were chosen as examples of the relative frequencies revision sequence for different sentences:

- FIXES ON A VALUE: $0 = 0$,
- CONVERGES TO A VALUE: $P^\top 0 = 0^\top = 1$,
- IN AN INTERVAL: One can easily extend what’s been done here to a joint language with both probability and truth and then a sentence whose revision would follow the example in IN AN INTERVAL would be δ with $\text{PA} \vdash \delta \leftrightarrow P^\top \delta^\top < 0.4 \vee (0.4 \leq P^\top \delta^\top \leq 0.6 \wedge T^\top \delta^\top)$,
- RELATIONSHIPS: this was δ as in IN AN INTERVAL along with $\neg\delta$,
- The figure on Page 15: $P^\top 0 = 0^\top = 1$.²⁸

²⁷Here # A counts how many objects there are in A .

²⁸Assuming we start with an initial hypothesis with $p_0(0 = 0) < 1$.

As a last example, a sentence ν with $\text{PA} \vdash \nu \leftrightarrow \neg \text{P}^\Gamma \nu^\neg > r$ will have its probability converge to r .²⁹ E.g. for $\text{PA} \vdash \nu \leftrightarrow \neg \text{P}^\Gamma \nu^\neg > 0.2$:



The transfinite successor stages for truth are usually defined in the same way as the finite successor stages. Following this same idea we would want to define $p_{\alpha+1}(\varphi)$ to be the relative frequency of φ being satisfied in the sequence of models up to $\alpha + 1$. However there are now transfinitely-many such models so how should one define relative frequency in this case? We shall define $p_{\alpha+1}(\varphi)$ to be the relative frequency of φ being satisfied in the finite sequence of models just before $\alpha + 1$. For an alternative approach see Campbell-Moore et al. [3].

For example,

$$p_{\omega+5}(\varphi) = \frac{\#\{\beta \in \{\omega, \omega+1, \dots, \omega+4\} \mid (\mathbb{R}, p_\beta) \models \varphi\}}{5}$$

To give this definition in general one needs to choose the finite sequence of models beneath an ordinal.

Definition 17 For a successor ordinal $\alpha + 1$ we let $\zeta_{\alpha+1}$ denote the greatest limit ordinal beneath $\alpha + 1$, and $k_{\alpha+1}$ be the natural number such that $\alpha + 1 = \zeta_{\alpha+1} + k_{\alpha+1}$.

For example $\zeta_{\omega+5} = \omega$ and $k_{\omega+5} = 5$. We can define:

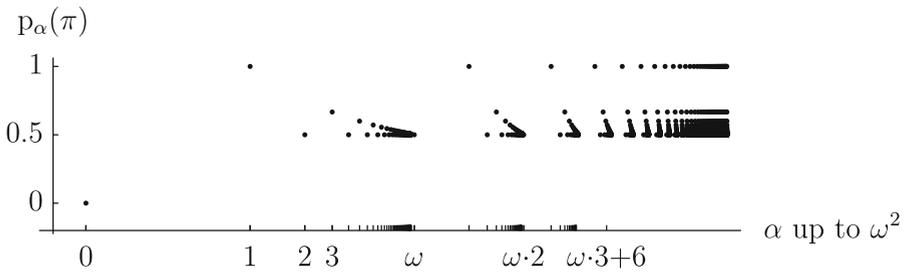
$$p_{\alpha+1}(\varphi) = \frac{\#\{\beta \in \{\zeta_{\alpha+1}, \zeta_{\alpha+1}+1, \dots, \alpha\} \mid (M, p_\beta) \models \varphi\}}{k_{\alpha+1}}$$

For the probabilistic liar sentence (π with $\pi \leftrightarrow \text{P}^\Gamma \pi^\neg \leq 1/2$) one would then obtain the following probabilities.

	ω	$\omega + 1$	$\omega + 2$	$\omega + 3$	$\omega + 4$	$\omega + 5$	$\omega + 6$...
$p_-(\pi) = :$	1/2	1	1/2	2/3	2/4	3/5	3/6	...
$(M, p_-) \models \pi ?$	✓	✗	✓	✗	✓	✗	✓	...

²⁹As will: $\nu \leftrightarrow \neg \text{P}^\Gamma \nu^\neg \geq r, \nu \leftrightarrow \text{P}^\Gamma \nu^\neg < r$ and $\nu \leftrightarrow \text{P}^\Gamma \nu^\neg \leq r$

Which we can picture as:

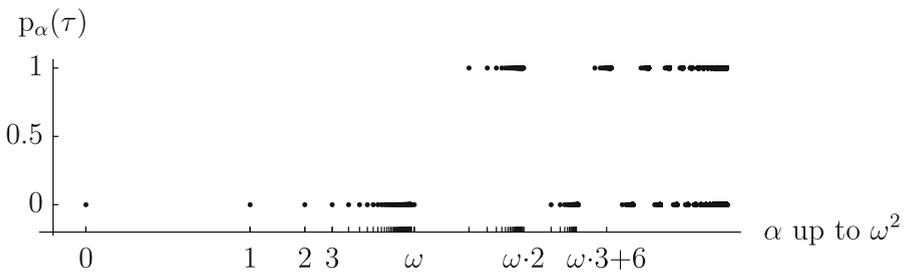


One might think that this Coherence of all Closed Properties will lead to $p_{\omega^2}(\lambda) = 1/2$. It turns out that that is not the case because, for example, $p(\lambda) \geq 1/4$ is not stable beneath ω^2 . As we have already very briefly mentioned, one might extend what counts as being brought about to be characterised by *near stability*³⁰ and the resulting constraint would lead to $p_{\omega^2}(\lambda) = 1/2$.

We can get examples of non-convex properties arising by considering the transfinitary revision sequence. Consider, e.g. τ where:

$$PA \vdash \tau \leftrightarrow (\gamma \wedge \neg T^\Gamma \tau^\neg) \vee (\neg \gamma \wedge T^\Gamma \tau^\neg)$$

where γ is the McGee sentence, i.e. is such that $\gamma \leftrightarrow \neg \forall n T^{n\Gamma} \gamma^\neg$.³¹ The probabilities of this τ beneath ω^2 are:



which has $p(\tau) = 0 \vee p(\tau) = 1$ stable beneath $\omega \cdot \omega$.

5 Conclusions

We have developed a general limit criterion that says that a limit stage should sum up the previous stages by satisfying the properties of interest that are brought about by the preceding sequence. The properties of interest are given by some \mathcal{P} . This new proposal is made precise in the criterion Coherence of \mathcal{P} , for appropriate choices of \mathcal{P} . Restrictions on appropriate choices were imposed: \mathcal{P} should contain all definite verdicts and only closed properties. The weakest such criterion is thus Coherence of

³⁰See Footnote 6.

³¹In particular, $(\mathbb{R}, T_-, p_-) \models \gamma$ at all successor ordinals, but $(\mathbb{R}, T_-, p_-) \models \neg \gamma$ at all ordinals of the form $\alpha + \omega$.

Definite Verdicts, which is exactly the usual constraint on the limit stage. We thus see our proposal as extending the revision theory, as developed in Gupta and Belnap [6], by providing a definition of what it is to be a revision sequence where we lift the requirement of only considering definite verdicts. Note that if one is drawn to a more restrictive limit clause where the revision sequence is *determined* by a choice of an initial stage (by, e.g., following a limit policy such as Herzberger [8] or Gupta [5]) then you might not like the generalisation we have presented, but a reader drawn to the liberal kind of policy should see these generalisations of the limit clause as legitimate proposals that need to be considered.

The restriction on the choice of \mathcal{P} that says it must only contain closed sets ensured that there is always some limit hypothesis satisfying Coherence of \mathcal{P} , at least if the hypothesis space (with its associated topology) is compact. Our proposal is thus restricted to such spaces. We furthermore restrict the hypothesis spaces by requiring that all definite verdicts be closed (otherwise there is no choice of \mathcal{P} that satisfies our requirements).

This general framework highlights the fact that when one presents a revision theory using Coherence of \mathcal{P} , the choice of \mathcal{P} should be well-motivated; and we contend that the choice of \mathcal{P} as just definite verdicts, which results in the traditional definition of what it is to be a revision sequence, has not yet been well-motivated enough. We suggest that the strongest such criterion, Coherence of all Closed Properties, is a particularly interesting criterion and has some motivation for it: a topology on a hypothesis space determines the notion of ‘limit’, or ‘approximation from within’, and then the closed sets are exactly the ones that act nicely under such limits; it thus also makes sense that they are the ones we take into consideration when determining how our limit stage should cohere with the previous stages. Coherence of all Closed Properties is also familiar in topology as it says that the limit hypothesis is a cluster point of the preceding sequence. We suggest that the default is that one considers this strongest limit criterion and that weakenings of it should be motivated, though we do not suggest that there are no such motivated weakenings, we have ourselves suggested that one might weaken the properties considered to just convex ones in the revision theory of probability.

In the revision theory of truth, we have seen that restricting the limit criterion by considering more than just definite verdicts can lead to interesting restrictions; for example Coherence of all Closed Properties ensures that the limit stages will be maximally consistent extensions of the truth predicate. It may turn out that in the truth case the \mathcal{P} that has the best motivation is in fact just definite verdicts, then we will result in the usual revision theory, but we would have achieved additional understanding of why that is the best revision theory of truth.

For some hypothesis spaces, however, the generalisation of the limit criterion is essential. We saw an example of this in Section 4 where the hypothesis space involves real numbers. This is because in such a space convergence plays an essential role and the restriction to just definite verdicts cannot recognise such facts.

Further work should be done to study how the range of different limit criteria affect the revision theory and to present arguments for motivating a particular instance of this limit criterion (via a choice of \mathcal{P}) for the different concepts the revision theory is being used to understand.

Acknowledgements I would like to thank Hannes Leitgeb, Johannes Stern, Seamus Bradley and Leon Horsten and for their helpful discussions and comments on versions of this paper. I am also very grateful to the organisers and participants at a number of conferences, including: Predicate Approaches to Modality, Munich; 45th meeting of the Society for Exact Philosophy, Caltech, Pasadena; Bristol-Leuven conference, Bristol; Logic Seminar, Cambridge; Bristol-München Conference on Truth and Rationality, Bristol.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Belnap, N.D. (1982). Gupta's rule of revision theory of truth. *Journal of Philosophical Logic*, 11(1), 103–116.
2. Campbell-Moore, C. (2016). Self-referential probability. PhD thesis, Ludwig-Maximilians-Universität München.
3. Campbell-Moore, C., Horsten, L., Leitgeb, H. Probability for the revision theory of truth, to appear in this volume.
4. Chapuis, A. (1996). Alternative revision theories of truth. *Journal of Philosophical Logic*, 25(4), 399–423.
5. Gupta, A. (1982). Truth and paradox. *Journal of Philosophical Logic*, 11(1), 1–60.
6. Gupta, A., & Belnap, N.D. (1993). *The revision theory of truth*. Cambridge: MIT Press.
7. Halbach, V. (2011). *Axiomatic theories of truth*. Cambridge: Cambridge University Press.
8. Herzberger, H.G. (1982). Notes on naive semantics. *Journal of Philosophical Logic*, 11(1), 61–102.
9. Leitgeb, H. (2008). On the probabilistic convention T. *The Review of Symbolic Logic*, 1(02), 218–224.
10. Leitgeb, H. (2012). From type-free truth to type-free probability. In G. Restall, & G. Russel (Eds.) *New waves in philosophical Logic edited by Restall and Russell* (pp. 84–94). New York: Palgrave Macmillan.
11. McGee, V. (1985). How truthlike can a predicate be? A negative result. *Journal of Philosophical Logic*, 14(4), 399–410.
12. Willard, S. (1970). *General topology*. New York: Courier Corporation.