



Probability for the Revision Theory of Truth

Catrin Campbell-Moore¹ · Leon Horsten¹ · Hannes Leitgeb²

Received: 13 September 2016 / Accepted: 24 May 2017 / Published online: 13 September 2018

© The Author(s) 2018

Abstract

We investigate how to assign probabilities to sentences that contain a type-free truth predicate. These probability values track how often a sentence is satisfied in transfinite revision sequences, following Gupta and Belnap’s revision theory of truth. This answers an open problem by Leitgeb which asks how one might describe transfinite stages of the revision sequence using such probability functions. We offer a general construction, and explore additional constraints that lead to desirable properties of the resulting probability function. One such property is Leitgeb’s Probabilistic Convention T, which says that the probability of φ equals the probability that φ is true.

Keywords Liar paradox · Semantic paradox · Revision theory of truth · Probabilistic convention T

1 Introduction

The revision theory of truth is an influential way to account for a type-free truth predicate. This theory constructs a series of hypotheses, or extensions, of the truth predicate. In this process one ‘revises’, or moves from one hypothesis to the next, by an application of Tarski biconditionals. After some sufficiently long initial sequence one can look back at the revision sequence and draw certain conclusions about the semantic status of different sentences. The sentence $T^{\top}0 = 0^{\top}$, for example, settles down on being true, whereas the liar sentence (which says of itself that it is not true) will have its truth value continuing to switch throughout the revision sequence. One then concludes that $T^{\top}0 = 0^{\top}$ is *true*, and that the liar sentence is *paradoxical*.

✉ Leon Horsten
Leon.Horsten@bristol.ac.uk

¹ Department of Philosophy, University of Bristol, Bristol, UK

² Munich Center for Mathematical Philosophy, Ludwig Maximilians Universität München, München, Germany

There are many ways of not settling down on a truth value. For instance, a sentence might fall into a pattern such as

$$t, \dots, t, f, t, \dots, t, f, t, \dots$$

where the strings of t 's ('true') coming between two successive f 's ('false') can be of any finite length, and the classification of this sentence will be the same as that of the liar sentence: they are both paradoxical. However, we suggest that a finer-grained classification system is needed because their *probability* of being true is different. In this paper we construct such probability functions that measure how often a sentence is true in a revision sequence of hypotheses.

This idea was pioneered in Leitgeb [9], where the finite stages of the revision theory of truth are used to define probabilities for sentences that may include a type-free truth predicate. We suggest that Leitgeb's construction yields a satisfactory account of the semantic probability of the sentences whose semantic behaviour is captured by the sequence of *finite* revision stages. For instance, the liar sentence is given semantic probability value $1/2$ because it always changes truth value in the transition from stage n to stage $n + 1$.

However, the semantic status of some sentences is not appropriately captured by the finite revision stages: one needs to proceed into the transfinite to see the true colours of such sentences. In this article, we explore ways that Leitgeb's construction can be extended to determine probability values dependent on the transfinite stages of the revision sequence, thus answering a question in Leitgeb [10].

Moreover, we will be interested in such probability functions that also satisfy certain notable global properties. One of these is Leitgeb's *Probabilistic Convention T*, which says that the probability of a proposition φ equals the probability of the proposition that φ is true. Another global property is regularity, which says that sentences have probability 0 only if they are always false in the revision sequence.

This article is structured as follows. We start by commenting on the kind of probability that is being assigned and how that is to be interpreted. We then, in Section 4.1, briefly review the revision theory of truth on which our probability assignments are based. In Section 4.2 we present the theory of Leitgeb [9], on which our proposal builds. Our proposal extends his by assigning probabilities at a transfinite ordinal stage, measuring how often the sentence is true in the revision sequence up to that ordinal. In Section 5 we then present the idea of our new proposal and some initial workings towards it. In fact, we formulate two distinct proposals which are different ways of extending our idea; however the two proposals are closely related in that the one is the 'standard part' of the other. These proposals are presented in Sections 5.2 and 5.3. It will emerge that which probability values get assigned and how the probability values work depends on a choice of an ultrafilter. In Section 6 we discuss different properties of the ultrafilters one might wish to impose, and consider their consequences for the probability functions that they determine. In Section 7 we conclude.

2 Some Preliminaries

Consider a language \mathcal{L}_T which adds to some background language \mathcal{L} (which at least contains Peano Arithmetic) a predicate T , called the *truth predicate*. We assume a

standard Gödel coding of the sentences of \mathcal{L}_T into the natural numbers, where the code of φ is denoted $\# \varphi$, and is referred to in the object language by $\ulcorner \varphi \urcorner$. In fact, in the interest of readability, and without real harm, we will at times be sloppy about coding matters.

In the revision theory of truth we are only interested in the extension of the truth predicate. So we fix some background model M of \mathcal{L} , which is a standard model of Peano Arithmetic for the arithmetic vocabulary. A *model* of \mathcal{L}_T is given by adding to M an interpretation of the truth predicate (a ‘hypothesis’), T , thus obtaining a model $\mathcal{M} = (M, T)$ of the language \mathcal{L}_T .

We will fix some limit ordinal, κ , and define a probability function Pr which measures how often φ is true in the revision sequence up to the stage κ . There are different ways of choosing this ordinal κ ; we mention a few here. One might do everything with κ as the class of ordinals On , though then one will have to take ultrafilters on proper class-sized objects. Alternatively one may fix κ as ‘long enough’. One proposal for such long-enough κ would be to set $\kappa = \omega_1$. The motivation for this would be that the semantic behaviour of every sentence of \mathcal{L}_T is conclusively exhibited well before ω_1 in the revision sequence. Finally, one might instead think of κ as variable limit ordinal and that what we are doing is assigning probabilities at each stage in the revision sequence; one might then read off facts about the probability by determining what is brought about by all Pr with κ large enough. For the purposes of this article any of these three interpretations are admitted.

Our aim is then to find some Pr that measures how often a sentence is true in a revision sequence up to the ordinal κ . Since the probability we are going to assign to a sentence depends only on the sequence of truth and falsity values, we will directly define such probability values as functions of such *truth profiles*, using 0 for ‘false’ and 1 for ‘true’:

Definition 1 A truth profile is a function $f : \kappa \rightarrow \{0, 1\}$.

Many, but of course not all, truth profiles are expressed by sentences of \mathcal{L}_T :

Definition 2 The truth profile of a sentence φ is the function

$$\text{TruthProfile}_\varphi(\alpha) := \begin{cases} 1 & \text{if } (M, T_\alpha) \models \varphi \\ 0 & \text{otherwise} \end{cases}$$

We want to define some $\text{Pr}(f)$ that measures how many 1s there are in a truth profile. We then abuse notation and set

$$\text{Pr}(\varphi) := \text{Pr}(\text{TruthProfile}_\varphi).$$

This definition of probability of a sentence in terms of the probability of its truth profile will henceforth be taken for granted.

3 Semantic Probability

In this article we are constructing a probability function that assigns probability values of sentences according to how often they are true in the revision theory of truth. But what kind of notion of *probability* is it?

The truth or falsehood of *grounded* sentences of \mathcal{L}_T are determined by the background model (the standard natural number structure). So there seems no immediate room for non-extremal probability values there. So the usual applications of probability that say interesting things about, e.g., tosses of coins will not be covered by this notion. But the probability notion we develop will allow for non-extremal probabilities in cases of paradoxical sentences, where the background model does not determine their truth value. Perhaps the best we can say about their truth status is based on the *frequency* of them being true in the revision sequence.

Since these frequencies are not grounded in empirical facts, ours is not a standard variant of the frequentist conception of probability. Instead, the truth frequencies are determined by revision sequences. The revision sequences are in turn driven by the Tarski-biconditionals, which are the basic *semantical rules* governing the truth predicate. So we are using a *semantic notion of probability*: the probabilities that we assign to paradoxical sentences might, in the spirit of Roeper and Leblanc [13, p. xi; p. 111–113], be taken to be their *semantic values*.

Semantic probabilities are not degrees of belief, and they are not chances as usually conceived. But there may nonetheless be connections between the semantic probabilities we assign and those other notions of probability. For example, perhaps one's degrees of belief should be weighted averages of semantic probabilities.¹ This proposal could be motivated by modifications of the traditional arguments for probabilism. One such argument for probabilism is based on *accuracy*, or epistemic utility, considerations (see Joyce [8]). This argument relies on the idea that beliefs should be as close to the truth, or 'accurate', as possible. One might think that the semantic probabilities we obtain in this paper should be treated as the semantic values of sentences in such arguments, and one's degrees of belief should be as close to such semantic values as possible. Along the lines of Williams [15] one will be able to show that this results in an agent's degrees of belief being required to be weighted averages of the semantic probabilities. A Dutch-book style argument will similarly be able to be used to this same conclusion (again, see Williams [15]). In the special case where the agent knows what the semantic probabilities are, this will also mean that her degrees of belief should then equal those semantic probabilities.

4 Background

4.1 The Revision Theory of Truth

In the revision theory of truth (Gupta and Belnap [6]), one constructs revision sequences, which are sequences of interpretations of a truth predicate. These can be used to determine appropriate interpretations of T, or at least facts about how appropriate interpretations will treat different sentences. Starting with an initial hypothesised interpretation, T_0 , one can revise this hypothesis to determine a new hypothesis, which will itself in turn be revised, etc. This process of revision develops

¹Many thanks to an anonymous referee for this suggested connection to the accuracy arguments.

a sequence of hypotheses. The revision of a hypothesis works by following the Tarski jump:

$$T_{\alpha+1} = \{\#\varphi \mid (M, T_\alpha) \models \varphi\}$$

The paradoxicality of the liar sentence then finds its expression in the fact that from any stage n to the next revision stage $n + 1$, the truth value of liar is reversed.

Typically, however, one wishes to extend the revision sequence into the transfinite in order to account for certain intuitive characteristics of truth. Consider the sentence:

$$\forall n T^n \ulcorner 0 = 0 \urcorner = 0^\ulcorner,$$

with T^n denoting a string of n occurrences of T . This sentence is not paradoxical: we expect any $T^n \ulcorner 0 = 0 \urcorner = 0^\ulcorner$ to be true simpliciter, and thus $\forall n T^n \ulcorner 0 = 0 \urcorner = 0^\ulcorner$ to be true. But this is not guaranteed to be obtained in any finite stage of the revision sequence: each finite stage may only recognise a finite number of iterations of the truth predicate applied to to $0 = 0$, and thus may take $\forall n T^n \ulcorner 0 = 0 \urcorner = 0^\ulcorner$ to be false. To guarantee the truth of $\forall n T^n \ulcorner 0 = 0 \urcorner = 0^\ulcorner$ one would then need to proceed to an ω^{th} stage. This is done by ensuring that if the revision sequence has settled down on some particular truth value then that truth value is also assigned at the limit stage. There are different rules for what to do at the limits, and we do not assume any particular such rule, but one standard one is the *Herzberger limit rule* where T_λ contains only the *stable* truths at any limit ordinal λ :

$$T_\lambda = \{\#\varphi \mid \exists \beta < \omega \forall \alpha \begin{matrix} \geq \beta \\ < \lambda \end{matrix} \varphi \in T_\alpha\}$$

We in fact need to go *beyond* T_ω and into the transfinite: even though for every n , the sentence $T^n \ulcorner 0 = 0 \urcorner = 0^\ulcorner$ must be in T_ω , the sentence $\forall n T^n \ulcorner 0 = 0 \urcorner = 0^\ulcorner$ may not yet be in T_ω . We thus have to proceed one step further into the transfinite, following the Tarski revision step for successor stages. Proceeding in this manner, we see that $\forall n T^n \ulcorner 0 = 0 \urcorner = 0^\ulcorner$ must enter the extension of the truth predicate by stage $\omega + 1$, and stays there forever after. It is easy to construct non-paradoxical sentences that first enters the extension of truth predicate (and stays in forever after) at a stage that is much larger than $\omega + 1$, at least when we start with the empty initial hypothesis.²

4.2 Leitgeb

In Leitgeb [9], a proposal is made for how to define probabilities based on relative frequencies in ω -length revision sequences. The present article can in part be seen as an answer to a question in Leitgeb [10], which asks how one might define such relative frequency based probabilities for revision sequences of ordinal length significantly greater than ω . Our paper provides an answer to this question. Leitgeb specifically asks whether one can take *Banach limits* of such longer revision sequences, which we can provide an answer for if we restrict our ultrafilter to be ‘Banach’ (see Section 6.4).

²These matters are discussed in Welch [14] and Horsten et al. [7, section 3].

Tarski’s theorem on the undefinability of truth shows that no classical model makes Tarski’s convention T true in an unrestricted way, i.e., there is no model where the Tarski-biconditionals

$$\varphi \leftrightarrow T^\top \varphi^\top$$

are true for all $\varphi \in \mathcal{L}_T$. In Leitgeb [9, p. 219] Leitgeb introduces his *Probabilistic Convention T*, which can be seen as an approximation of the unrestricted Tarski-biconditionals. Probabilistic Convention T says that for every sentence φ :

$$\Pr(\varphi) = \Pr(T^\top \varphi^\top).$$

This principle immediately has certain consequences: for example, Leitgeb notes that every probability function that satisfies Probabilistic Convention T must assign exactly probability 1/2 to the liar sentence.

Leitgeb shows that Probabilistic Convention T is consistent with the axioms of finitely additive probability by using the finite stages of the revision theory and providing a summary probability function capturing the ideas of relative frequencies in these finite stages of the revision theory.

We now present his proposal, which we subsequently generalise to apply also to transfinite revision sequences.

Leitgeb’s basic idea for obtaining probabilities based on relative frequencies in ω -length revision sequences is to take longer and longer finite samples, where relative frequencies can easily be defined, and take limits to obtain the final probability value. More carefully, then, if f is an ω -length truth profile—often it will be the truth values of some sentence in the ω -length revision sequence—we can directly define relative frequencies in each initial segment $\{0, 1, \dots, n - 1\}$:³

$$\text{RelFreq}_n(f) := \frac{\sum_{i=0}^{n-1} f(i)}{n} = \frac{|\{i \in \{0, \dots, n - 1\} : f(i) = 1\}|}{n}.$$

Leitgeb then uses these RelFreq_n to approximate the sought-after \Pr_ω , defining \Pr_ω to be a *limit* of the finite approximations.

The relative frequency idea, together with the classical (Cauchy-Weierstraß) conception of limit of a sequence, supports the view that if $\Pr_n(f)$ converges to r as n goes to ω , then $\Pr_\omega(f)$ should be taken to be identical with r .

However, not all sequences have convergent probability, so this does not determine \Pr_ω . As an example, consider a truth profile which switches value at each stage of the form 3^n . For such a profile, when the value is switching the number of stages in the final sequence of values is much larger than all the preceding stages; in fact twice as large. This means that its probability value then oscillates between being $\geq 2/3$ after each sequence of 1s, and being $\leq 1/3$ after each sequence of 0s. There is a sentence that has this non-convergent profile as its truth profile:⁴

$$\forall n \left(T^n \top 0 = 0^\top \wedge \neg T^{n-1} \top 0 = 0^\top \rightarrow \forall p ((3^p \leq n \wedge 3^{p+1} > n) \rightarrow \text{Even}(p)) \right)$$

³Where $|X|$ denotes the cardinality of X .

⁴At least if we take $T_0 = \emptyset$, as then n is the first stage at which $T^n \top 0 = 0^\top$ is satisfied. So the antecedent of this formula identifies n as the stage of the revision theory we are at. $\text{Even}(p)$ is shorthand for $\exists x (2 \cdot x = p)$.

So the familiar notion of limit does not give us a probability function that is defined on all profiles, or even all sentences of \mathcal{L}_T . To determine a probability value for *all* profiles, one can take a *generalised limit* at stage ω of the finite relative frequencies. There are of course many ways of doing this, many of which would result in a finitely additive probability function defined on \mathcal{L}_T . But Leitgeb proposes that we take a *Banach limit* at stage ω , and define the sought after probability function Pr to be such a Banach limit. A Banach limit is a generalised limit that is *shift-invariant*:

Definition 3 For an truth profile $f : \kappa \rightarrow \{0, 1\}$, the shift Sf of f is the function such that

$$Sf(\alpha) = f(\alpha + 1).$$

Definition 4 Pr is shift-invariant iff for all f , $\text{Pr}(f) = \text{Pr}(Sf)$.

Leitgeb then shows how being shift-invariant guarantees that Pr indeed meets the adequacy condition of satisfying Probabilistic Convention T (Leitgeb [9], Theorem 1, p. 219–220).

In the previous section we saw that the sentence $\forall n \top^n \top 0 = 0^\top$ is not necessarily classified as true at stage ω . Since $\text{Pr}_n(\forall n \top^n \top 0 = 0^\top)$ may be 0 for all n , it follows that $\text{Pr}_\omega(\forall n \top^n \top 0 = 0^\top)$ may be 0. But just as it seems intuitively clear that $\forall n \top^n \top 0 = 0^\top$ is true, it seems that it *should* get a probability (approximating) 1. It seems that the overwhelming likelihood of $\forall n \top^n \top 0 = 0^\top$ being true becomes visible only in the transfinite. There are also paradoxical sentences whose ‘true colours’ only become visible after stage ω . It is not hard to construct sentences that start oscillating in a liar-like fashion only after some transfinite stage α (but do something different before stage α).⁵ Thus we have also probabilistic reasons for extending Leitgeb’s construction into the transfinite.⁶

5 Two Proposals

In this section, we try to improve on the revision theories of probability that we have discussed above, and formulate two new proposals. In the first proposal, a class of real-valued probability measures is defined. In the second proposal, a class of hyper-real probability functions is constructed.

5.1 The Idea Behind the Proposals

We want Pr to be something like the relative frequency of φ being true, but we only have a clear idea of how to define the required relative frequency if we have finitely

⁵See footnote 2.

⁶Campbell-Moore [5] presents a proposal which might also be seen as a method of extending Leitgeb’s construction to the transfinite; but Campbell-Moore points out a number of undesirable features of her proposal. This paper can be seen as an improvement on her proposal. There are also some technical similarities between PrApprox that we construct and her limit stage requirements: If C is a closed property (see Campbell-Moore [5]) and $\{X \in [\kappa]^{<\omega} \mid \text{RelFreq}_X(f) \in C\} \in \mathcal{U}$ then $\text{PrApprox}(f) \in C$.

many stages. What we will do, then, is first take finite samples and work out relative frequencies for these, and then sum up the findings from this finite sampling into a probability value. In order to sum up the findings, we will be using an ultrafilter.

Leitgeb’s construction only looked at up to $n - 1: \{0, 1, \dots, n - 1\}$, but we here define finite sampling more generally. For a finite, non-empty set of ordinals X we define

$$\text{RelFreq}_X(f) := \frac{\sum_{\alpha \in X} f(\alpha)}{|X|} = \frac{|\{\alpha \in X : f(\alpha) = 1\}|}{|X|}.$$

So, in particular, if X is the interval $\{0, 1, \dots, n - 1\}$, we have RelFreq_X is the same as RelFreq_n as defined earlier.

We now want to choose Pr in such a way that it sums up the results from the finite sampling from ordinals $< \kappa$. Using standard notation, we call the collection of all such finite, non-empty samples

$$[\kappa]^{<\omega} := \{X \subseteq \kappa \mid X \neq \emptyset \text{ and } X \text{ is finite}\}.$$

So we define Pr using the $\text{RelFreq}_X(f)$ for $X \in [\kappa]^{<\omega}$. The probability value we will assign will consider what happens in ‘enough’ samples. If *enough* samples have $\text{RelFreq}_X(f) = r$ then we will assign $\text{Pr}(f) = r$. We will capture this notion of ‘enough’ by an *ultrafilter*, \mathcal{U} .⁷ This imposes rather strong constraints on the notion of ‘enough’, for example either a collection of samples counts as enough or its complement does. But these are required to ensure that our proposals end up assigning every sentence a probability value.

We then want:

$$\text{Pr}(f) = r \in \mathbb{R} \text{ if } \{X \in [\kappa]^{<\omega} \mid \text{RelFreq}_X(f) = r\} \in \mathcal{U},$$

where \mathcal{U} is an ultrafilter on $[\kappa]^{<\omega}$.

This does not yet fix the probability values of all sentences: it might be that for no real number r do enough samples have $\text{RelFreq}_X(\varphi) = r$.⁸

To be able to assign values to all sentences we make two different proposals. The first is to assign a sentence a real number as a probability value if that number is *approximated* in enough samples. This proposal is developed in Section 5.2, and is called PrApprox . The second proposal we make does require that the value the sentence is assigned is exactly achieved in enough samples, it just allows that the probability values be non-standard real numbers, or *hyperreals*. We develop this proposal in Section 5.3, and the probability notion developed is labelled PrHyp . These

⁷ $\mathcal{U} \subseteq \wp([\kappa]^{<\omega})$ is an *ultrafilter* iff:

- $[\kappa]^{<\omega} \in \mathcal{U}$,
- $\emptyset \notin \mathcal{U}$
- If $A \in \mathcal{U}$ and $B \supseteq A$, then $B \in \mathcal{U}$,
- If $A \in \mathcal{U}$ and $B \in \mathcal{U}$, then $A \cap B \in \mathcal{U}$.
- Either $A \in \mathcal{U}$ or $[\kappa]^{<\omega} \setminus A \in \mathcal{U}$.

It is rather common to use the metaphor of ‘enough’ to describe such ultrafilters.

⁸For example, consider $\top \uparrow 0 = 0 \uparrow$ which has the truth profile f, t, t, t, \dots if we start with $T_0 = \emptyset$. For any fine, non-principal ultrafilter this will not have its relative frequency equal to any fixed value in enough samples. See Section 6 for definitions of these kinds of ultrafilters.

two proposals that we present are very closely connected: PrApprox is just the standard part of PrHyp, as is made precise in Theorem 9, but we think it is valuable to present PrApprox directly both because we think that helps understanding PrHyp without needing to first know anything about the hyperreals, and because we think it is a valuable proposal in itself.

Our definitions are always relative to an ultrafilter \mathcal{U} . The probability values that are assigned by these are thus dependent on the ultrafilter chosen.⁹ In Section 6 we consider some further global properties that can be imposed on ultrafilters and discuss the consequences they have for the probability values assigned by the proposals.

5.2 Real-Valued Probabilities

The proposal developed in this section extends the previous suggestion, which took a real number as the value of a sentence if that number was the relative frequency in enough samples, by looking at *approximation* instead of equality. This will allow that every sentence is assigned a particular real number as its probability value, of course relative to a chosen ultrafilter.

We use $\text{RelFreq}_X(f) \approx_{1/n} r$ as shorthand for $|\text{RelFreq}_X(f) - r| < 1/n$.

Definition 5 $\text{PrApprox}(f) = r \in \mathbb{R}$ iff for all $n \in \omega$,

$$\{X \in [\kappa]^{<\omega} \mid \text{RelFreq}_X(f) \approx_{1/n} r\} \in \mathcal{U}.$$

It turns out that this does in fact ensure that each sentence obtains a probability value, and even ensures that the probability values they receive satisfy the axioms of finitely additive probability.

Proposition 6 PrApprox assigns a value to each sentence and this is a finitely additive probability function.

Proof This can be proved directly, but it can be seen more concisely as a corollary of Proposition 9 and the fact that PrHyp is a perfectly additive probability function. \square

We take this to be a potential suggestion for our sought-after Pr. However we will also give another alternative in the next section.

5.3 Hyperrational Probabilities

The probability functions that were discussed in the previous section are everywhere defined, real valued and finitely additive. But they are not σ -additive, i.e., there may be cases where

$$\Pr(\exists x \varphi(x)) \neq \lim_{n \rightarrow \omega} \Pr(\varphi(\bar{0}) \vee \dots \vee \varphi(\bar{n}))$$

(Leitgeb [10], Theorem 1, p. 219–220). Indeed, they do not satisfy any natural infinite additivity rule.

⁹See Benci et al. [3] for further discussion of the dependence on the ultrafilter.

We will now define a class of probability functions that do satisfy a natural and strong infinite additivity principle. This principle is called *perfect additivity*. It says, roughly, that for *any* family (countable or uncountable) of pairwise disjoint events, the probability of its union is the sum of the probabilities of the members of the family.¹⁰ The price that we will have to pay is that these probability functions take their value in a non-archimedean extension of \mathbb{Q} (*hyperrational spaces*). This construction uses the machinery developed in Benci et al. [2].

This construction allows us to keep the equivalence

$$\text{PrHyp}(f) = r \text{ iff } \{X \in [\kappa]^{<\omega} \mid \text{RelFreq}_X(f) = r\} \in \mathcal{U},$$

instead of resorting to using approximation in the right hand side of this equivalence. However it will do so by assigning *non-standard* probability values to sentences where no real satisfies the right hand side of this equivalence.

As before, we consider a revision sequence for \mathcal{L}_T of length κ , and we take an ultrafilter \mathcal{U}_κ on $[\kappa]^{<\omega}$.

We will construct the hyperreal space in the usual way, by using equivalence classes based on the ultrafilter chosen.

Definition 7 For $F, G : [\kappa]^{<\omega} \rightarrow \mathbb{Q}$, define: $F \sim_\kappa G$ iff

$$\{X \in [\kappa]^{<\omega} \mid F(X) = G(X)\} \in \mathcal{U}.$$

Here F and G can be thought of as giving relative frequencies of truth on finite sets (‘samples’) of hypotheses.

Now we use the \sim_κ equivalence classes to be the objects in our hyperrational space: $H_\kappa = \{[F]_{\sim_\kappa}\}$. On the basis of elementary arguments from non-standard analysis, this can be seen as a non-Archimedean extension of the rational $[0,1]$ interval by associating a rational number r with $[F_r]_{\sim_\kappa}$, where F_r is the constant function assigning r to every X .

The field operations can be similarly used on this hyperrational space by defining them point-wise; e.g. $[F]_{\sim_\kappa} + [G]_{\sim_\kappa} = [F + G]_{\sim_\kappa}$ —one can check that these are well-defined. One can also define $[F]_{\sim_\kappa} < [G]_{\sim_\kappa}$ iff $\{X \mid F(X) < G(X)\} \in \mathcal{U}$. One can then define notions like ‘infinitesimally close to’: $x \approx x'$ iff for all $n \in \omega$, $|x - x'| < 1/n$.

Then for any function $F : [\kappa]^{<\omega} \rightarrow [0, 1]^\mathbb{Q}$ we can define the (hyperrational) *probability of F* to be $[F]_{\sim_\kappa}$. We obtain our desired probability values for truth profiles by using RelFreq:

Definition 8 For f a truth profile, define $F^f : [\kappa]^{<\omega} \rightarrow \mathbb{Q}$ by $F^f(X) := \text{RelFreq}_X(f)$. Then

$$\text{PrHyp}(f) := [F^f]_{\sim_\kappa}.$$

This definition extends our original attempted definition in that we have for $r \in \mathbb{Q}$, $\text{PrHyp}(\varphi) = r$ iff $\{X \in [\kappa]^{<\omega} \mid \text{RelFreq}_X(\varphi) = r\} \in \mathcal{U}$.

¹⁰This is discussed in detail in Brickhill and Horsten [4, section 2.5]. (The exact perfect additivity principle that we get is given in Proposition 8 of that paper.)

PrHyp is always an everywhere defined finitely additive probability function (Benci et al. [2], section 3.4). Moreover, PrHyp is *perfectly additive* (Benci et al. [2], Proposition 8, p. 132–133).¹¹

In the mathematical literature, probability theory is *almost* regarded as applied mathematical analysis. In other words, in mathematical probability theory, many feel that if an account does not contain a rule that relates infinite events to sub-events that are infinitely small in comparison to it —of course σ -additivity, which makes use of the classical notion of *limit*, is the standard way of doing this— then it cannot really be said to be a *theory* of probability. In the philosophical literature on probability, in contrast, finite additivity is often regarded as sufficient.

For many purposes in philosophy we can confine ourselves to finite sample spaces (finite sets of possible worlds, for instance). But infinity is at the heart of what we are concerned with: we need to take account of infinitely many stages. In the context that we are concerned with, it is desirable to have a rule that tells us how the probability value of a sentence of the form $\exists x\varphi(x)$ supervenes on the probability of the sentences $\varphi(0), \varphi(1), \varphi(2), \dots$, and this requires an infinite additivity principle. This seems to be particularly important for the *semantic* notion of probability that we are developing here.

The definition of PrHyp relates to the notion PrApprox in the following way:

Proposition 9 $\text{PrApprox}(f) \approx \text{PrHyp}(f)$, i.e. $\text{PrApprox}(f)$ is the standard part of $\text{PrHyp}(f)$.

Proof Due to the definition of PrApprox, we have:

$$\{X \mid \text{RelFreq}_X(f) \approx_{1/n} \text{PrApprox}(f)\} \in \mathcal{U}$$

for all n and we can thus conclude that $\text{PrHyp}(f) \approx_{1/n} \text{PrApprox}(f)$ for all n , i.e. $\text{PrHyp}(f) \approx \text{PrApprox}(f)$. \square

6 Constraints on Ultrafilters

The construction we have presented depends on the choice of an ultrafilter \mathcal{U} . We might wish to study what happens with particular choices of ultrafilters. In this section we consider additional restrictions one may wish to impose on the choice of ultrafilter. Using such restricted ultrafilters then allows us to read off additional properties of the probabilities assigned. For example the property of Banachness imposed on an ultrafilter ensures that the probability of the liar sentence is $1/2$.

This section considers five properties of ultrafilters, Non-Principality, Fineness, Stability, Banachness and GrandLoop, and the features of the resulting probability functions that they lead to. At the end of the section we show to what extent the properties are compatible.

For the reader who wishes to skip technical details we provide a summary of the principles and what they lead to in Section 6.6.

¹¹The notions of perfect additivity are defined for probabilities as applied to *events* instead of truth profiles or sentences. But one can see how to apply these definitions in our context by using Definition 25.

6.1 Non-Principal

There is a good reason to think that we should require that \mathcal{U} be a non-principal ultrafilter as otherwise the probability is determined by a single finite snapshot.

Definition 10 A filter \mathcal{U} on $[\kappa]^{<\omega}$ is *non-principal* if there is no $X \in [\kappa]^{<\omega}$ such that $\mathcal{U} = \{\mathcal{A} \subseteq [\kappa]^{<\omega} \mid X \in \mathcal{A}\}$ (or equivalently $\{X\} \in \mathcal{U}$).

In fact we can say something stronger about principal *ultrafilters*: they are determined by a sample consisting of a *single* stage: i.e. if \mathcal{U} is a principal ultrafilter then there is some α such that $\mathcal{U} = \{\mathcal{A} \subseteq [\kappa]^{<\omega} \mid \{\alpha\} \in \mathcal{A}\}$.

Principal ultrafilters count a single sample, in fact a singleton sample, as being enough, leading to an unnatural notion of ‘enough’. They also lead to undesirable probability assignments: If the probability is defined using a principal ultrafilter, then we will get that $\text{PrHyp}(f) = \text{RelFreq}_{\{\alpha\}}(f)$ for the α which generates \mathcal{U} . We will thus have that $\text{PrHyp}(f) = f(\alpha)$. So probability values will always be 0 or 1 depending on whether the sentence is satisfied at α or not. This is undesirable. We will thus require that our ultrafilter is non-principal.

Non-principal ultrafilters, along with all the others we mention in this section, will always exist, but we leave this result to Section 6.6.

6.2 Fine

There is a school of thought that holds that it is desirable for a probability function to be so fine-grained as to distinguish between impossible and possible events. In other words, only the probability of the empty event should be 0. This requirement is called *regularity* (Lewis [11], p. 267).¹² In our context, regularity amounts to the requirement that $\text{PrHyp}(f) = 0$ only if the sentence is always false up to stage κ . We obtain such regularity if we restrict our attention to *fine* ultrafilters:

Definition 11 For $\beta < \kappa$ define:

$$A_{\text{Fine}}^{\beta} := \{X \in [\kappa]^{<\omega} \mid \beta \in X\}$$

An ultrafilter \mathcal{U} is *fine* if $A_{\text{Fine}}^{\beta} \in \mathcal{U}$ for all $\beta < \kappa$.

Hyperrational probability functions based on fine ultrafilters have interesting properties:

Proposition 12 *If the ultrafilter \mathcal{U} on which PrHyp is based is fine, then PrHyp is regular (in the sense given above).*

Proof Brickhill and Horsten [4, Proposition 2.5]. □

¹²A philosophical discussion of the requirement of regularity can be found in Benci et al. [3, section 5.2].

It is easy to see that fineness entails non-principality. Moreover, fineness entails *uniformity*, i.e., that the probability of any truth profile that has a 1 only at a single ordinal is the same. Or, in terms of events, the probability of any singleton event is the same as that of any other singleton.

Proposition 13 *If the ultrafilter \mathcal{U} on which PrHyp is based is fine, then PrHyp is uniform.*

Proof Brickhill and Horsten [4, Proposition 2.6]. □

In the present setting, given perfect additivity, uniformity means that every model in the revision sequence makes an equal contribution to the probability of any sentence of \mathcal{L}_T .

Alternatively, a *weight function* can be added to the construction of PrHyp, as in Benci et al. [2]. This will enable us, for instance, to assign more weight to ordinals that are further away from limit ordinals (as one might think that stages that are further away from a limit are better). This would allow us to get regularity without uniformity.

It is easy to see that if PrHyp is constructed from a fine ultrafilter, then the sentence $\neg T^\top 0 = 0^\top$, which has truth profile $(1, 0, 0, 0, 0, \dots)$ if we start from the empty hypothesis, is given an infinitesimally small but non-zero probability value by it. Proposition 9, which relates the real-valued probability function PrApprox to the hyperrational probability function PrHyp, then entails that $\text{PrApprox}(\neg T^\top 0 = 0^\top) = 0$. So PrApprox will not in general be regular even if it is based on a fine ultrafilter. Note that this is not necessarily a bad thing: it seems good to *completely* ignore the only model in the revision sequence that makes the grounded sentence $\neg T^\top 0 = 0^\top$ true even though it *should* come out false! We will return to the question of to what extent regularity is desirable in Section 6.6.

6.3 Stability

The next property of ultrafilters that we consider is that of being a *stability ultrafilter*. A stability ultrafilter only cares about what happens in the ‘final’ part of the truth profile of a sentence. The motivation for it therefore directly conflicts with the motivation for the fineness property, and indeed we will see in Section 6.6 that the property of stability we introduce is incompatible with that of fineness, at least if these properties are required unrestrictedly.

The intuition behind the stability notion is that if a sentence ends up always being false from some stage onwards, then a stability ultrafilter will give that probability 0 as it doesn’t care what happened in the initial part of the truth profile. This will get us, for example, that $\text{PrHyp}(T^\top 0 = 0^\top) = 0$, whereas regularity got us that it was non-zero if our starting hypothesis was empty.

Definition 14 For $\beta < \kappa$ define:

$$A_{\text{Stability}}^\beta := \{X \in [\kappa]^{<\omega} \mid \text{for all } \alpha \in X, \alpha \geq \beta\}.$$

We say \mathcal{U} is a *stability-at- β ultrafilter* if $A_{\text{Stability}}^\beta \in \mathcal{U}$.

This allows one to ignore initial stages of the revision sequences, as made precise by the following proposition:

Proposition 15 *Suppose $A_{\text{Stability}}^\beta \in \mathcal{U}$. Suppose f and f' are truth profiles such that for $\kappa > \alpha \geq \beta$, $f(\alpha) = f'(\alpha)$. Then $\text{PrHyp}(f) = \text{PrHyp}(f')$.*

Proof Let f and f' agree from β onwards. Then for each $X \subseteq [\beta, \kappa)$, $\text{RelFreq}_X(f) = \text{RelFreq}_X(f')$, and thus $\{X \mid \text{RelFreq}_X(f) = \text{RelFreq}_X(f')\} \supseteq A_{\text{Stability}}^\beta \in \mathcal{U}$, so by the superset property of ultrafilters $\{X \mid \text{RelFreq}_X(f) = \text{RelFreq}_X(f')\} \in \mathcal{U}$. Therefore $\text{PrHyp}(f) = \text{PrHyp}(f')$. \square

Taking an ultrafilter on $[\kappa]^{<\omega}$ that is a stability-at- β ultrafilter is equivalent to taking an ultrafilter on (finite samples of ordinals in) $[\beta, \kappa)$, i.e., where one ignores the ordinals up to β . So considering a stability-at- β ultrafilter is essentially the same as just applying all the considerations to a revision sequence that starts at the point β .

6.4 Banachness

Leitgeb asked how one might take *Banach limits* to obtain probability values based on relative frequencies in long revision sequences (Leitgeb [10], Section 2.3, Question 2). In this section we introduce a property of ultrafilters that yields Leitgeb’s sought-after generalisation of his construction.

We have seen in Section 4.2 that Leitgeb’s construction leads to a shift invariant probability function: assigning the same probability values to an ω -length truth profile, f , and its shift Sf . One of the aims of the present paper is to base probability judgements on revision sequences longer than ω . In the hyperrational approach, given regularity, it is clear that we cannot have, for all truth profiles f , $\text{PrHyp}(f) = \text{PrHyp}(Sf)$: otherwise the probability value assigned would ignore the truth value at stage 0, but regularity rules that out. But it might be reasonable to ask in this context for *almost shift invariance*, i.e., $\text{PrHyp}(f) \approx \text{PrHyp}(Sf)$. This would entail that for all $\varphi \in \mathcal{L}_T$, $\text{PrHyp}(\varphi) \approx \text{PrHyp}(T^\top \varphi^\top)$. In other words, such hyperrational probability functions would *approximately* satisfy Probabilistic Convention T, and the associated real probability functions, PrApprox , would *fully* satisfy Probabilistic Convention T.

On longer finite initial intervals, the relative frequencies associated with a truth profile f and its shift Sf approximate each other ever closer. So the rough idea for constructing a shift invariant probability measure is to let larger and larger finite intervals belong to the ultrafilter on which the probability function is based.¹³

¹³A simpler requirement would be to define $A_{\text{Banach}}^n := \{I_\alpha^k \mid n < k\}$. But this would cause conflicts with Fineness so we take this weaker definition which is also sufficient for obtaining shift-invariance.

Since we will be using it a lot, it will be useful to give a shorthand for intervals of length $k \in \mathbb{N}$ starting at α .

Definition 16 Let $I_\alpha^k := [\alpha, \alpha + k) = \{\gamma \mid \alpha \leq \gamma < \alpha + k\}$.

Definition 17 For each $n \in \mathbb{N}$ define

$$A_{\text{Banach}}^n := \left\{ \bigcup_{i=1}^m I_{\alpha_i}^{k_i} \in [\kappa]^{<\omega} \mid \begin{array}{l} m \in \mathbb{N}, \\ n < k_1, \dots, k_m \in \mathbb{N}, \\ \alpha_1, \dots, \alpha_m < \kappa, \end{array} \right\}.$$

Say that \mathcal{U} is *Banach* if each $A_{\text{Banach}}^n \in \mathcal{U}$.

This will do the job we require: such Banach ultrafilters will generate shift invariant PrApprox measures, and almost shift invariant PrHyp measures. We thus have provided our answer to Leitgeb’s open problem.¹⁴

Recall that for any truth profile $f : \kappa \rightarrow \{0, 1\}$, the shift of f , Sf , is defined by $Sf(\alpha) = f(\alpha + 1)$.

Proposition 18 *If \mathcal{U} is a Banach ultrafilter, then PrApprox is shift-invariant, i.e.*

$$\text{PrApprox}(f) = \text{PrApprox}(Sf).$$

And PrHyp approximates the equality (i.e. replace $=$ by \approx).

Proof Let n be chosen. We want to show that

$$\{X \in [\kappa]^{<\omega} \mid \text{RelFreq}_X(f) \approx_{1/n} \text{RelFreq}_X(Sf)\} \in \mathcal{U}.$$

We will show that for each $X \in A_{\text{Banach}}^{2n}$, $\text{RelFreq}_X(f) \approx_{1/n} \text{RelFreq}_X(Sf)$. The result will then follow by the assumption that $A_{\text{Banach}}^{2n} \in \mathcal{U}$.

Suppose X is some member of A_{Banach}^{2n} . Then X is of the form $\bigcup_{1 \leq i \leq m} I_{\alpha_i}^{k_i}$ where $k_i > 2n$. Without loss of generality we can choose these intervals to be disjoint.

Since $k_i > 2n$,

$$\text{RelFreq}_{I_{\alpha_i}^{k_i}}(f) - \text{RelFreq}_{I_{\alpha_i}^{k_i}}(Sf) = \frac{f(\alpha_i) - f(\alpha_i + k_i)}{k_i} < \frac{2}{k_i} < \frac{1}{n}$$

¹⁴Leitgeb actually asks how one can take Banach limits, which are *linear functionals* that are shift-invariant. We have just defined our probabilities on truth profiles, which are sequences of 1s and 0s. If we instead defined everything to apply a “probability” value to bounded sequences of reals then our PrApprox would be a linear functional that is shift-invariant, i.e. a Banach limit. This alternative definition would also allow a more direct proof of Proposition 20 from Proposition 18.

So X is a finite union of intervals, where the relative frequency of f and Sf are closer than $1/n$ in each interval, and thus the relative frequency of f and Sf are closer than $1/n$ in the whole of X , as required. \square

Corollary 19 *PrApprox satisfies Probabilistic Convention T. I.e.*

$$\text{PrApprox}(\varphi) = \text{PrApprox}(\text{T}^\top \varphi^\top)$$

and PrHyp approximates it (i.e. replace = by \approx).

Proof From Proposition 18, as in the proof of Theorem 1(2) in Leitgeb [9, p. 221]. \square

This constitutes a significant improvement on the probability measures proposed in Leitgeb [9]: we have now satisfied Probabilistic Convention T whilst having at the same time assigned ‘correct’ probability values to sentences that take a long time to stabilise (such as the sentence $\forall n \text{T}^n \top 0 = 0^\top$ and its more sophisticated variants).

The key feature of Banach ultrafilters is that they ensure that the ordering of the truth values plays a role in the final probability assignment. If a truth profile falls into a repeating pattern, for example, 1, 1, 0, 1, 1, 0, 1, 1, 0, . . . , then we might want to ensure that the probability assigned to this is $2/3$. By ensuring that the ultrafilter is Banach we obtain this result.

Proposition 20 *Suppose f falls into a repeating pattern of length N that starts at stage M ($N, M \in \mathbb{N}$),¹⁶ then*

$$\text{PrApprox}(f) = \text{RelFreq}_{I_M^N}(f)$$

and thus $\text{PrHyp}(f) \approx \text{RelFreq}_{I_M^N}(f)$.

Proof We actually only need to show the result for the case when $M = 0$ (i.e. when f is N -periodic) and the result for other M holds by shift invariance (Proposition 18) as $\text{PrApprox}(f) = \text{PrApprox}(S^M f)$ and $S^M f$ is a sequence that just consists of copies of the repeating sequence.

The argument works in a similar manner to that of Proposition 18. A sketch is as follows: Consider any union of intervals and each longer than $N \cdot K$. The collection of all such intervals is a member of \mathcal{U} by our Banachness assumption. We show that any such union has $\text{RelFreq}_X(f) \approx_{1/K} \text{RelFreq}_{I_0^N}(f)$, because the majority of X is formed of N -length intervals. This proves our result. \square

¹⁵More carefully: for $X = \bigcup_{1 \leq i \leq m} I_{\alpha_i}^{k_i}$,

$$\begin{aligned} & \text{RelFreq}_X(f) - \text{RelFreq}_X(Sf) \\ &= \frac{\sum_{1 \leq i \leq m} k_i (\text{RelFreq}_{I_{\alpha_i}^{k_i}}(f) - \text{RelFreq}_{I_{\alpha_i}^{k_i}}(Sf))}{\sum_{1 \leq i \leq m} k_i} \\ &\leq \frac{\sum_{1 \leq i \leq m} k_i \cdot 1/n}{\sum_{1 \leq i \leq m} k_i} \leq 1/n \end{aligned}$$

¹⁶This will later be called N -periodic after M : see Definition 23 for a careful definition of this.

If we consider, in addition to Banachness, the property of stability, we will be able to obtain the result that the probability of sentences whose truth profile falls into a repeating finite-length sequence after some transfinite stage is reduced to a finite relative frequency. This therefore extends what we obtained in Proposition 20 where we only required Banach ultrafilters, whereas here we need Banachness *and* stability.

Proposition 21 *Suppose f falls into a repeating pattern of length N that starts at stage β ($N \in \mathbb{N}$, $\beta < \kappa$).*

Let $\alpha \geq \beta$. If \mathcal{U} is Banach and $A_{\text{Stability}}^\alpha \in \mathcal{U}$ then

$$\text{PrApprox}(f) = \text{RelFreq}_{I_\alpha^N}(f)$$

and thus $\text{PrHyp}(f) \approx \text{RelFreq}_{I_\alpha^N}(f)$.

Proof The argument works as in Proposition 20. □

By finite applications of shift-invariance, we see that the Banach property of shift-invariance entails the property of *finite* shift invariance, so for each $n \in \mathbb{N}$, $\text{PrHyp}(f) \approx \text{PrHyp}(S^n(f))$. So we can obtain almost finite-shift-invariance of our probability functions for transfinite revision sequences. It is natural to ask whether stronger forms of almost infinite-shift-invariance can also be obtained. One might wonder, for instance, whether we can obtain probability functions such that

$$\text{PrHyp}(f) \approx \text{PrHyp}(f + \omega),$$

where $(f + \omega)(\alpha) := f(\alpha + \omega)$. Indeed, if we choose $\kappa = \omega_1$, one might wonder if we can have hyperrational probability functions for \mathcal{L}_T that are almost β -shift-invariant for every $\beta < \omega_1$. We will not pursue this theme further in this article.¹⁷

6.5 Grand Loop

It is known from the literature on the revision theory of truth that if we use the Herzberger limit rule then every sentence falls into a repeating pattern. This eventual periodicity of the entire revision sequence is known as the *Grand Loop* phenomenon. This phenomenon also appears using limit rules such as the ‘constant bootstrapping policies’ of Gupta, but notably not for the permissive limit criterion which allows one to make choices at each limit ordinal. For this section we focus on the Herzberger limit rule, though any other rule that has the Grand Loop phenomenon would be equally applicable here.

There are thus ordinals ζ and λ such that for any sentence, φ , after ζ the truth profile of the sentence just consists of copies of a single λ -length sequence. One moral to be taken from this Grand Loop phenomenon is that after going through the stages $[\zeta, \zeta + \lambda)$, one might as well stop, for from then on the pattern will just

¹⁷The question of shift-invariance has been investigated to some extent in the context of *Numerosity Theory* (see Benci et al. [1]), which is from a mathematical point of view closely related to the generalised probability functions that are investigated here.

repeat indefinitely. We might want the Grand Loop phenomenon to be reflected in our probability measure. In other words, we might want the probability of the sentence to be able to be determined by its action on $[\zeta, \zeta + \lambda)$. In this section we will find such a result: we will impose a particular constraint on the ultrafilters that ensures that the probability of a sentence is the probability of the sentence on the interval $[\zeta, \zeta + \lambda)$, or in fact in any loop. We will formalise this idea in terms of conditional probabilities.

In this section we will impose a further constraint on κ : that it is closed under λ -addition, i.e. if $\alpha < \kappa$ then $\alpha + \lambda < \kappa$. If κ is ω_1 , for example, then this will hold for any $\lambda < \omega_1$.

To simplify the presentation we will first work as if the model associated with stage ζ , where the grand loop starts, is the initial stage 0 of our revision theory. To ensure that our ultrafilter property works for $\zeta > 0$, we will later additionally impose the property of stability-at- ζ , which as mentioned in Section 6.3 is equivalent to just considering an ultrafilter on $[\zeta, \kappa)$.

There is a little bit of work to do to formalise the informal idea of looping sequences when the looping length is transfinite. For this one needs the notion of ‘modulus’:¹⁸

Definition 22 Define $\alpha \bmod \lambda$ as the unique $\gamma < \lambda$ for which there is some δ with $\alpha = \lambda \cdot \delta + \gamma$.

We can now formalise the idea of the loop:

Definition 23 We say that a truth profile f is λ -periodic if for all α , $f(\alpha) = f(\alpha \bmod \lambda)$.

Or, equivalently, if for all α, β , if $\alpha \bmod \lambda \equiv \beta \bmod \lambda$, then $f(\alpha) = f(\beta)$. A truth profile is λ -periodic after ζ if the above holds restricted to $\alpha, \beta > \zeta$.¹⁹

The Grand Loop phenomenon²⁰ then says:

Theorem 24 There are ordinals ζ and λ such that for all $\varphi \in \mathcal{L}_T$, the truth profile of φ is λ -periodic after ζ .

So what we are looking for is some property of ultrafilters that will ensure that the probability of the sentence is the same as the probability just focusing on the interval $[0, \lambda)$ —recall that we are first assuming for simplicity that the beginning of the grand loop is at stage 0.

¹⁸Rivello [12] also uses this setup. See his paper for more analysis of the Grand Loop phenomenon.

¹⁹Note that our definition differs slightly from Rivello [12] as it is slightly more convenient for us, but note that the definitions are equivalent if ζ is a multiple of λ .

²⁰Results in this section are perhaps still relevant to someone who is interested in alternative limit criteria which need not exhibit this Grand Loop phenomenon as they might tell one about the action of the sentences that do happen to become periodic. However one would have to reconsider Theorem 30 with varying λ .

In order to state this we first will need to define this notion of ‘probability just looking at $[0, \lambda)$ ’, which we do by means of defining *conditional probability*:

Definition 25 We can identify a set $E \subseteq \kappa$, or an *event* with a truth-profile: its characteristic function.²¹ We will abuse notation and refer just to E when we mean its truth profile. We will also assume that intersections of truth profiles are defined as would be expected.²²

If $\text{PrHyp}(E) > 0$ (possibly infinitesimal), then

$$\text{PrHyp}(f|E) := \frac{\text{PrHyp}(f \cap E)}{\text{PrHyp}(E)}.$$

Which property of ultrafilters will get us what we want? It will be such that the collection of $X \in [\kappa]^{<\omega}$ where

$$\text{RelFreq}_X(f) = \frac{\text{RelFreq}_X(f \cap [0, \lambda))}{\text{RelFreq}_X([0, \lambda))} = \text{RelFreq}_{X \cap [0, \lambda)}(f)$$

is a member of \mathcal{U} . So we just need to ensure that there is some $A_{\text{GrandLoop}}^{0,\lambda} \in \mathcal{U}$ where all the $X \in A_{\text{GrandLoop}}^{0,\lambda}$ have that property. (The 0-superscript is for later generalisation).²³ So our question becomes: Which $X \in [\kappa]^{<\omega}$ have the feature that if f is λ -looping then

$$\text{RelFreq}_X(f) = \text{RelFreq}_{X \cap [0, \lambda)}(f),$$

i.e., that the relative frequency in the whole of X is the same as in $X \cap [0, \lambda)$?

A sufficient condition on X is that how it looks in $[0, \lambda)$ is the same as it looks at the other loops.²⁴ Such a condition will ensure that if f is λ -periodic, its relative frequency in X will be the same as the relative frequency in $X \cap [0, \lambda)$.

To define this we will first introduce a more general notion for λ -length intervals starting at ordinals θ . We only present the definition for θ multiples of λ ,²⁵ as they’ll be the starting points of the loops and will be all we need for our definition.

²¹ $\text{TruthProfile}_E(\alpha) := \begin{cases} 1 & \alpha \in E \\ 0 & \text{otherwise} \end{cases}$

²²For truth profiles f and g , $(f \cap g)(\alpha) = \begin{cases} 1 & f(\alpha) = 1 \text{ and } g(\alpha) = 1 \\ 0 & \text{otherwise} \end{cases}$.

²³It will also need to be that for all $X \in A_{\text{GrandLoop}}^{0,\lambda}$, $X \cap [0, \lambda) \neq \emptyset$ (to make the denominator > 0), but this will hold for our choice of $A_{\text{GrandLoop}}^{0,\lambda}$.

²⁴A weaker requirement is

$$X = \left\{ \bigcup_{\gamma \in \Gamma} \Sigma_\gamma \in [\kappa]^{<\omega} \mid \begin{array}{l} \Gamma \text{ is a finite set of ordinals each } < \lambda \\ \text{For each } \delta \in \Sigma_\gamma, \delta \bmod \lambda = \gamma \\ \text{For each } \gamma \in \Gamma, \Sigma_\gamma \cap I_\xi^\lambda \neq \emptyset \\ \text{There is some } m \in \mathbb{N} \text{ such that } |\Sigma_\gamma| = m \text{ for each } \gamma \in \Gamma \end{array} \right\}$$

which allows the required repetitions of the members of $X \cap I_\xi^\lambda$ to not come in blocks but be spread out. We present the stronger requirement to keep the arguments easier to follow, but one can check that everything goes through with this alternative requirement.

²⁵By which we mean $\theta \bmod \lambda = 0$.

Definition 26 For θ a multiple of λ ,²⁶ define the λ -interval starting at θ as:

$$I_\theta^\lambda := [\theta, \theta + \lambda) = \{\alpha \mid \theta \leq \alpha < \theta + \lambda\}.$$

We also present the definition not just for the 0 case, but for other ordinals ξ that are starting points of loops.

Definition 27 For ξ a multiple of λ define $A_{\text{GrandLoop}}^{\xi, \lambda} :=$

$$\left\{ \begin{array}{l} \bigcup_{\theta \in \Theta} \Delta_\theta \in [k]^{<\omega} \\ \left. \begin{array}{l} \Theta \text{ is a finite collection of multiples of } \lambda, \text{ and contains } \xi \\ \Delta_\theta \subseteq I_\theta^\lambda \\ \text{Each } \Delta_\theta \text{ is a } \lambda\text{-copy of } \Delta_\xi, \\ \text{i.e. } \{\xi + \alpha \bmod \lambda \mid \alpha \in \Delta_\theta\} = \Delta_\xi \end{array} \right\} \end{array} \right.$$

Say \mathcal{U} is a GrandLoop_λ ultrafilter if $A_{\text{GrandLoop}}^{\xi, \lambda} \in \mathcal{U}$ for all ξ multiples of λ .

This set has been constructed so that if the ultrafilter contains it then the probability of any λ -periodic sequence is the same as its conditional probability on $[0, \lambda)$.

Theorem 28 If f is a λ -periodic sequence, and $A_{\text{GrandLoop}}^{\xi, \lambda} \in \mathcal{U}$, then

$$\text{PrHyp}(f) = \text{PrHyp}(f \mid I_\xi^\lambda).$$

Proof

$$\begin{aligned} \text{RelFreq}_{\Delta_\theta}(f) &= \frac{\sum_{\alpha \in \Delta_\theta} f(\alpha)}{|\Delta_\theta|} = \frac{\sum_{\alpha \in \Delta_\theta} f(\xi + \alpha \bmod \lambda)}{|\Delta_\theta|} \\ &= \frac{\sum_{\beta \in \Delta_\xi} f(\beta)}{|\Delta_\xi|} = \text{RelFreq}_{\Delta_\xi}(f) \end{aligned}$$

Since $X \in A_{\text{GrandLoop}}^{\xi, \lambda}$ is a finite union of such Δ_θ , and $\Delta_\xi = X \cap I_\xi^\lambda$, we have that

$$\text{RelFreq}_X(f) = \text{RelFreq}_{\Delta_\xi}(f) = \text{RelFreq}_{X \cap I_\xi^\lambda}(f) = \frac{\text{RelFreq}_X(f \cap I_\xi^\lambda)}{\text{RelFreq}_X(I_\xi^\lambda)}. \quad \square$$

As we mentioned in the introduction, the grand loop result for the revision theory only says that sentences will be periodic *after some ordinal* ζ . By ensuring that our ultrafilter is a stability-at- ζ ultrafilter we can essentially get it to ignore all stages before ζ , which are the cases that would mess up with the relative frequencies. Another way of getting the same result is of course simply to take the first repeated point ζ as the starting hypothesis of the revision sequence. If we do that, then we do not have to impose an additional stability condition on our ultrafilter.

²⁶The definition without this restriction would be: If $\theta = \lambda \cdot \delta + \gamma$, for $\gamma < \lambda$, then:

$$I_\theta^\lambda := [\lambda \cdot \delta + \gamma, \lambda \cdot (\delta + 1) + \gamma).$$

And the following theorems would then still hold.

Theorem 29 *Suppose f is λ -periodic after ζ , and ξ is a multiple of λ and is $> \zeta$. If $A_{Stability}^\zeta \in \mathcal{U}$ and $A_{GrandLoop}^{\xi,\lambda} \in \mathcal{U}$, then*

$$\text{PrHyp}(f) = \text{PrHyp}(f|I_\xi^\lambda).$$

Thus we have

$$\text{PrHyp}(\varphi) = \text{PrHyp}(\varphi|I_\xi^\lambda)$$

for any $\varphi \in \mathcal{L}_T$, with ζ and λ as in Theorem 24.

Proof The argument exactly follows as in Theorem 28 observing that now by our stability assumption the θ can be restricted to ordinals $> \zeta$, where the relative frequencies work properly. □

By imposing stability-at- ζ we have thus essentially ignored all models up to stage ζ and this can be seen to be motivated by the grand loop picture: For every stage $\alpha < \zeta$, there are non-pathological true sentences (such as $\forall n T^n \top = 0$) that are made false by model T_α . So any such model is determinately deficient.

In contrast, none of the models T_α for $\alpha \geq \zeta$ assigns the wrong truth value to any non-pathological sentence. So it is likewise motivated to require that every T_α for $\alpha \geq \zeta$ makes an infinitesimal non-zero contribution to the probability value of sentences. To get this coming out in our probability assignments one could impose the constraint of fineness at all stages after ζ , which is possible because the ultrafilter properties of fineness and grand loop cohere, as is shown in the next section.

6.6 How These Properties of Ultrafilters Relate

We have now presented all the properties of ultrafilters that we will consider. We first just summarise these properties:

- **Non-Principal:** We want the ultrafilter to be non-principal because otherwise the probability value of a sentence would always reduce to the truth value at some fixed stage of the revision sequence, and therefore would not in any sense capture how often the sentence is true.
- **Fineness:** This allows PrHyp to be regular and uniform, treating each stage of the revision sequence equally.
- **Stability:** This allows one to ignore any initial segment of the revision sequence, capturing the idea that the later stages are better and should be initial stages are just ‘junk’ that can be thrown away.
- **Banachness:** This allows us to obtain Probabilistic Convention T (at least for PrApprox): $\text{PrApprox}(\varphi) = \text{PrApprox}(T^\top \varphi^\top)$. This is the principle that motivated Leitgeb’s original introduction of these probabilities of the revision sequences.
- **GrandLoop:** This is motivated by specific revision theory considerations: revision sequences (at least those where the limit is determined by a *rule*) will settle into a periodic pattern — the Grand Loop phenomenon. This GrandLoop property of ultrafilters that we considered ensures that the probability of a sentence is given by just considering its action within one of the loops of the Grand Loop.

Each of these properties we find interesting, and we find ourselves attracted to them in varying degrees. We find non-principality essential, and GrandLoop we think of as an interesting optional feature. One can note, though, that not all of them can be asked to simultaneously hold unrestrictedly. In particular we see that the motivations for Fineness and Stability are inconsistent. Fineness is motivated by treating all stages equally, whereas Stability allows us to ignore initial stages and treat later stages as more important. However, restricted versions are consistent: so long as one only asks for Stability up to some point and Fineness after that these will be compatible. There is similarly a conflict between GrandLoop and Stability as GrandLoop says we can just consider a particular loop, whereas Stability tells us that from some later perspective that loop should be ignored. However, again they will be consistent if we restrict Stability up to some stage and GrandLoop after that.

Our next two results make the above more precise:

Theorem 30 *For any $\lambda < \kappa$, there is an ultrafilter that has all the above-mentioned properties except for Stability, i.e. Non-Principality, Fineness, Banachness and GrandLoop_λ .*

Stability can be added if we restrict Fineness and GrandLoop:

Theorem 31 *For any $\lambda < \kappa$, $\zeta \leq \kappa$, ζ a multiple of λ ,²⁷ there is an ultrafilter, \mathcal{U} satisfying Non-Principality, Fineness-after- ζ , Stability-before- ζ , Banachness, and GrandLoop_λ -after- ζ . I.e.:*

- \mathcal{U} is Non-Principal,
- $A_{\text{Fine}}^\alpha \in \mathcal{U}$ for all $\alpha > \zeta$,
- $A_{\text{Stability}}^\alpha \in \mathcal{U}$ for all $\alpha < \zeta$,
- \mathcal{U} is Banach, i.e. $A_{\text{Banach}}^n \in \mathcal{U}$ for all n ,
- $A_{\text{GrandLoop}}^{\xi, \lambda} \in \mathcal{U}$ for all ξ multiples of λ that are $> \zeta$.

Proof Note that Theorem 30 follows from Theorem 31, so we prove the latter.

We note that an ultrafilter is non-principal iff $A \in \mathcal{U}$ for every cofinite A .²⁸

We use the standard fact that if $\mathcal{A} \subseteq \wp[\kappa]^{<\omega}$ has the finite intersection property, i.e. the intersection of any finite $\mathcal{B} \subseteq \mathcal{A}$ is non-empty, then there is an ultrafilter extending \mathcal{A} .

So take finitely many from each of our classes, and we will find X that is in all of these classes.

- Non-Principal: $X \in A_1 \cap \dots \cap A_n$, where each A_i is cofinite.
- Fineness: $X \in A_{\text{Fine}}^{\gamma_1} \cap \dots \cap A_{\text{Fine}}^{\gamma_n}$ with $\gamma_i > \zeta$
 - So we need to ensure $\gamma_1, \dots, \gamma_n \in X$.

²⁷If $\zeta = \kappa$ then Stability is required unrestrictedly, and GrandLoop and Fineness are dropped.

²⁸I.e. $[\kappa]^{<\omega} \setminus A$ is finite.

- **Stability:** $X \in A_{\text{Stability}}^{\beta_1} \cap \dots \cap A_{\text{Stability}}^{\beta_n}$ with $\beta_i < \zeta$
 - So $X \subseteq [\beta_i, \kappa)$.
- **Banachness:** $X \in A_{\text{Banach}}^{m_1} \cap \dots \cap A_{\text{Banach}}^{m_n}$
 - So X is a finite union of finite intervals, each at least m_i long. So we just need to check that it's a finite union of intervals each at least $m := \max\{m_1, \dots, m_n\}$ long.
- **GrandLoop:** $X \in A_{\text{GrandLoop}}^{\xi_1, \lambda} \cap \dots \cap A_{\text{GrandLoop}}^{\xi_n, \lambda}$ with each ξ_i a multiple of λ and $> \zeta$.
 - So X is a finite union of $\Delta_\theta \subseteq I_\theta^\lambda$, each being a λ -copy of Δ_{ξ_i} .

We will define infinitely many sets that are members of each of the sets, except for the non-principal ones. Then since the non-principality condition can only rule out finitely many sets, one such set must be in $A_1 \cap \dots \cap A_n$ as that is a cofinite set. Take any $M \geq m$.

$$Y_M := \bigcup_{i \leq n} I_{\gamma_i}^M.$$

This would already be enough to get all of our properties except for GrandLoop. For GrandLoop we need to get the copies working properly. The initial idea would be to ensure than anything that is a λ -copy of something in Y_M gets thrown into our set. That will end up with an infinite set, so instead we just add the additional members that are in the λ -intervals of interest, i.e., the ones that are non-empty. More carefully, then, we define the following: Θ will be the collection of starting points of our intervals of interest, and Σ_M will describe how each copy of interest looks: Let

$$\Theta := \{\theta \mid \theta \text{ is a multiple of } \lambda \text{ and } I_\theta^\lambda \cap Y_M \neq \emptyset\} \cup \{\xi_1, \dots, \xi_n\}.$$

$$\Sigma_M := \{\alpha \bmod \lambda \mid \alpha \in Y_M\}$$

These sets are both finite.²⁹

Now define

$$X_M := \{\theta + \alpha \mid \theta \in \Theta \text{ and } \alpha \in \Sigma_M\}.$$

Choose some M' with $X_{M'} \in A_1 \cap \dots \cap A_n$. This is possible because $A_1 \cap \dots \cap A_n$ is cofinite. Then let

$$X = X_{M'}.$$

We need to check that X is a member of all our required sets.

- **Non-Principal:** We have chosen X to be some $X_{M'}$ that lies in $A_1 \cap \dots \cap A_n$.
- **Fineness:** $\gamma_i \in I_{\gamma_i}^m \subseteq Y_{M'} \subseteq X$
- **Stability:** Because $\gamma_i > \zeta$, each $\alpha \in Y_{M'}$ is $> \zeta$. And because ζ is a multiple of λ , each θ with $I_\theta^\lambda \cap Y_{M'} \neq \emptyset$ must be $> \zeta$. Also we have each $\xi_i > \zeta$. So each $\alpha \in \Theta_{M'}$ is $> \zeta$. And thus each $\alpha \in X$ is $> \zeta$. Since we required our β_i be $< \zeta$ we therefore have that each $\alpha \in X$ is $> \beta_i$.

²⁹Note that we could equally have defined Θ to depend on M , but this would not change anything.

- Banachness: We started with $Y_{M'}$ being a union of M' -length intervals, and we only added further λ -copies of M' -length intervals, so we will never have obtained an interval that is shorter than $M' > m$.
- GrandLoop: Let $\Delta_\theta := X \cap I_\theta^\lambda = \{\theta + \alpha \mid \alpha \in \Sigma_{M'}\}$, and see that $X = \bigcup_{\theta \in \Theta} \Delta_\theta$. Observe that $\{\xi_i + \alpha \bmod \lambda \mid \alpha \in \Delta_\theta\} = \{\xi_i + \alpha \mid \alpha \in \Gamma_{M'}\} = \Delta_{\xi_i}$. □

One could retain commitment to Stability in its full generality and drop Fineness and GrandLoop. Commitment to GrandLoop is interesting but we think it is not as well motivated as the alternative properties, and it is only interesting if one fixes attention on limit rules that lead to the Grand Loop phenomenon. One might worry that by not requiring Fineness we can ignore certain features of the sequence that should not be ignored. A weaker idea that is consistent with the Stability idea is that we should care about all the cofinal hypotheses. That does not mean caring about any particular ordinal at which they appear, but just some ordinal at which the hypothesis reappears. It is an open question whether we can impose a property on the ultrafilters that would get us this feature.

From a philosophical point of view, the most attractive way to go might instead be to take that the global constraints on ultrafilters that we have discussed are *to some extent* correct. From our discussion of Stability, we take on board that the stages before the initial ordinal ζ should be disregarded completely because the hypothesis of every such stage is definitely defective. From our discussion of GrandLoop, we take on board that the stages from the first repeating ordinal Z onwards can all be disregarded completely because they are merely repeating hypotheses that we have seen before. From the discussion of Fineness, we take on board that each of the stages between ζ and Z should be counted and should be counted equally, because we have no convincing reasons for disqualifying any of those hypotheses or finding any of them less plausible than another.³⁰ And Banachness is important for the reasons that are discussed in Leitgeb [9].

This means that from a philosophical point of view, probability measures determined by fine Banach ultrafilters on the finite subsets of the half open interval $[\zeta, Z)$ might be of particular interest. Nonetheless, *all* the global restrictions on ultrafilters that we have discussed are of interest in themselves. We hope that our discussion of them illustrates the power and flexibility of the ultrafilter technique for building hyperirrational probability functions that satisfy specific desirable properties.

7 Further Research

In the revision theory of truth, the truth value of a sentence of a language that includes a type-free truth predicate is revised over and over. Some sentences eventually settle on a truth value that remains unchanged throughout the remainder of the revision sequence. These might be called the *grounded sentences*. But many sentences never settle on a truth value. These are the *paradoxical sentences*. In the spirit of Leitgeb

³⁰Fineness entails non-principality, so this suggestion would also retain commitment to non-principality.

[9], the suggestion of this article is that paradoxical sentences do not have a truth value as their semantic value, but have a *probability value* as their semantic value.

Building on Leitgeb [9] we have formulated and discussed two new proposals for associating relative frequency-inspired probability values to sentences of a type-free truth language on the basis of a revision sequence. In one proposal, the resulting probability function is real-valued; in the second proposal, the probability values are hyperrational-valued. But the two proposals are closely related: the real-valued probability functions can be seen as approximations of hyperrational-valued probability functions.

We could at this point extend our proposal in the spirit of Leitgeb [10] to obtain theories of *self-referential probability* based on a revision sequence. This would involve adding a probability function symbol pr to the language \mathcal{L}_T , and developing revision models for that. Our construction for defining hyperrational probability functions allows us to extract probability assignments at each stage of the revision process and can use this to add probability to the object language itself: we let the probability function symbol pr to the object language be interpreted at each stage β by the corresponding Pr_β .

There are subtleties to be dealt with for this version. For one thing, properties of ultrafilters that we have discussed, which lead to nice properties at limits, might not be desirable at successor stages. For another, it seems overkill to choose (Axiom of Choice!) a new ultrafilter at *each* stage α of the revision sequence in order to construct Pr_α . Fortunately, it turns out that it suffices to choose *one* ultrafilter on $[\kappa]^{<\omega}$ for the whole revision sequence of length κ , and obtain the ultrafilters on the (finite subsets of the) initial segments of the whole sequence by *uniformly restricting* the ultrafilter on $[\kappa]^{<\omega}$ to finite subsets of smaller ordinals.³¹ However, we leave the details of this for a future occasion.

Acknowledgements We would like to thank Philip Welch and Johannes Stern for helpful discussions on this topic, as well as anonymous referees for this journal for their helpful suggestions. Catrin Campbell-Moore would also like to thank Corpus Christi College, Cambridge, where she was based as a Stipendiary Research Fellow while this paper was being written.

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. Benci, V., Di Nasso, M., Forti, M. (2006). An Aristotelian notion of size. *Annals of Pure and Applied Logic*, 143(1-3), 43–53.
2. Benci, V., Horsten, L., Wenmackers, S. (2013). Non-Archimedean probability. *Milan Journal of Mathematics*, 81(1), 121–151.
3. Benci, V., Horsten, L., Wenmackers, S. (2018). Infinitesimal probabilities. *The British Journal for the Philosophy of Science*, 69(2).

³¹Though this needs to be applied with care as restrictions of ultrafilters need not be ultrafilters, though one can show that they will be so long as \mathcal{U} is not a stability-at- β ultrafilter for any $\beta < \kappa$.

4. Brickhill, H., & Horsten, L. (2018). Triangulating non-Archimedean probability. *The Review of Symbolic Logic*, 1–28. <https://doi.org/10.1017/S1755020318000060>.
5. Campbell-Moore, C. (2018). *Limits in the revision theory: more than just definite verdicts*. Journal of Philosophical Logic, this edition.
6. Gupta, A., & Belnap, N.D. (1993). *The revision theory of truth*. Cambridge: MIT Press.
7. Horsten, L., Leigh, G.E., Leitgeb, H., Welch, P. (2012). Revision revisited. *The Review of Symbolic Logic*, 5(4), 642–664.
8. Joyce, J. (2009). *Accuracy and coherence prospects for an alethic epistemology of partial beliefs*. Berlin: Springer.
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 Restall, G., & Russel, G. (Eds.) *New Waves in philosophical logic* (pp. 84–94). Palgrave Macmillan.
11. Lewis, D. (1980). A subjectivist's guide to objective chance. In Jeffrey, R. (Ed.) *Studies in inductive logic and probability*, (Vol. 2 pp. 263–293). Berkeley and Los Angeles: University of California Press.
12. Rivello, E. (2015). Periodicity and reflexivity in revision sequences. *Studia Logica*, 103(6), 1279–1302.
13. Roeper, P., & Leblanc, H. (1999). *Probability theory and probability logic*. University of Toronto Press.
14. Welch, P.D. (2014). Some observations on truth hierarchies. *The Review of Symbolic Logic*, 7(1), 1–30.
15. Williams, J.R.G. (2014). Probability and non-classical logic. In Hitchcock, C., & Hájek, A. (Eds.) *Oxford Handbook of Probability and Philosophy* (pp. 248–276). Oxford University Press.