

COMPUTATIONAL ASPECTS OF SATISFIABILITY IN PROBABILITY LOGIC

RUTGER KUYPER

ABSTRACT. We consider the complexity of satisfiability in ε -logic, a probability logic. We show that for the relational fragment this problem is Σ_1^1 -complete for rational $\varepsilon \in (0, 1)$, answering a question by Terwijn. In contrast, we show that satisfiability in 0-logic is decidable. The methods we employ to prove this fact also allow us to show that 0-logic is compact, while it was previously shown that ε -logic is not compact for $\varepsilon \in (0, 1)$.

1. INTRODUCTION

Many people have tried to combine logic with probability, which leads to so-called *probability logics*. In this paper we will study the computational aspects of one specific probability logic, introduced by Terwijn in [11]. This probability logic is called ε -logic, because it depends on a fixed error parameter ε . There are two important aspects of ε -logic which set it apart from other probability logics: first, the logic is closely related to and motivated by probabilistic induction and Valiant's pac-model from computational learning theory, and second, instead of using an entirely new syntax we use the same syntax as classical first-order language and only change its interpretation.

Terwijn was not the first to introduce a logic motivated by Valiant's pac-model. In fact, Valiant himself also introduced a probability logic related to his pac-model in [13]. However, as discussed in Kuyper and Terwijn [8] there are major differences between Valiant's logic and ε -logic: two of the most important ones being that Valiant only studies finite models, and that his syntax captures only a fragment of first-order logic. A logic which is actually more closely related to ε -logic is Keisler's logic $\mathcal{L}_{\omega P}$, surveyed in Keisler [6]. Instead of the classical quantifiers, this logic has quantifiers of the form $(Px \geq r)$ which should be read as "holds for at least measure r many x ". While this logic does not attempt to model probabilistic induction, and does not contain the classical universal and existential quantifiers, it turns out we can adapt some of the ideas used to prove results about $\mathcal{L}_{\omega P}$ to obtain similar results for ε -logic. In particular, in this paper we use some ideas from Hoover [4] in which he proves that validity for $\mathcal{L}_{\omega P}$ is Π_1^1 -complete. For more connections to other probability logics, we refer to the introduction of [8].

Unfortunately, it turns out that ε -logic is computationally quite hard. Previously, Terwijn [12] has shown that the set of ε -tautologies is undecidable, and in Kuyper [7] it was shown that this set is in fact even Π_1^1 -hard. So, ε -validity is computationally much harder than validity in first-order logic, and therefore we cannot hope to find an effective calculus for it.

Date: 7th November 2014.

2010 Mathematics Subject Classification. 03B48, 68Q17, 03B53.

Key words and phrases. Probability logic, Computational complexity, Satisfiability.

Research supported by NWO/DIAMANT grant 613.009.011 and by John Templeton Foundation grant 15619: 'Mind, Mechanism and Mathematics: Turing Centenary Research Project'.

In this paper, we will study the fragment of ε -logic not containing equality or function symbols, i.e. containing only relation and constant symbols. For this fragment, we will show that ε -satisfiability is, in general, Σ_1^1 -complete, refuting a conjecture by Terwijn [12, Conjecture 5.3]. At first one might think that the Σ_1^1 -hardness of ε -satisfiability already follows from the Π_1^1 -hardness of ε -validity mentioned above: indeed, in the classical case a formula φ is satisfiable if and only if its negation $\neg\varphi$ is not valid. However, because our logic is paraconsistent (i.e. both a formula φ and its negation $\neg\varphi$ can hold at the same time) this complementarity does not hold for ε -satisfiability and ε -validity. Therefore, we need to consider the complexity of these two problems separately.

For $\varepsilon = 0$ the results are vastly different: 0-validity coincides with classical validity, as shown in Terwijn [12], so it is Σ_1^0 -complete. In this paper we show that 0-satisfiability is decidable, so 0-satisfiability is even easier than classical satisfiability. This also clearly shows that the complexities of 0-validity and 0-satisfiability are not complementary, as argued above. The different complexities for ε -satisfiability and ε -validity are summarised in Table 1 below. Note that the exact complexity of ε -validity is still open, as there is no known matching upper bound for the Π_1^1 -hardness.

TABLE 1. Complexity of validity and satisfiability in ε -logic.

	$\varepsilon \in (0, 1) \cap \mathbb{Q}$	$\varepsilon = 0$
ε -satisfiability	Σ_1^1 -complete	decidable
ε -validity	Π_1^1 -hard	Σ_1^0 -complete

The paper is structured as follows. In the next section we will briefly recall the definition of ε -logic and some facts proven in earlier papers. After that, in section 3 we show that a certain weak version of ε -satisfiability is Σ_1^1 . In section 4 we briefly turn towards the problem of Skolemisation in ε -logic, which we will need for the later results in this paper. After that, in section 5 we show that the relational fragment of ε -satisfiability is Σ_1^1 . Next, in section 6 we turn to 0-satisfiability and we show that this problem is, quite surprisingly, decidable. This contrasts our result in section 7 that ε -satisfiability is Σ_1^1 -hard for rational $\varepsilon \in (0, 1)$, which completes our proof that ε -satisfiability is Σ_1^1 -complete for such ε . Finally, in section 8 we use the results from section 6 to show that 0-logic is compact, contrasting an earlier result by Kuyper and Terwijn [8, Theorem 8.2] that ε -logic is not compact for rational $\varepsilon \in (0, 1)$.

Our notation is mostly standard. We let ω denote the set of natural numbers. We denote by λ the Lebesgue measure on the unit interval $[0, 1]$. By a relational formula we mean a formula which only contains relation symbols and does not contain equality, function or constant symbols. Throughout the paper, \mathcal{D} denotes a probability measure, or synonymously a probability distribution. When we say that some property holds for \mathcal{D} -almost all x , we mean that the set of x for which the property holds has \mathcal{D} -measure 1; when \mathcal{D} is clear from the context we will omit it. For all formulas φ and ψ , the formula $\varphi \leftrightarrow \psi$ is short for $(\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$. For any measure \mathcal{D} and any $n \in \omega$, we let \mathcal{D}^n denote the product measure of n copies of \mathcal{D} , and we let \mathcal{D}^ω denote the product measure of countable infinitely many copies of \mathcal{D} . For functions $f, g \in \omega^\omega$ we denote by $f \oplus g$ the function defined by $f \oplus g(2n) = f(n)$ and $f \oplus g(2n + 1) = g(n)$.

For unexplained notions from measure theory, we refer to Bogachev [2], for descriptive set theory, we refer to Kechris [5], for model theory, we refer to Chang

and Keisler [3], and finally, for unexplained notions from computability theory we refer to Odifreddi [9].

2. ε -LOGIC

As we have mentioned above, ε -logic was introduced in Terwijn [11]. Over time, the definition has evolved into its current form, which was first given in Kuyper and Terwijn [8] and slightly differs from the original definition.

Definition 2.1. Let \mathcal{L} be a first-order language, possibly containing the equality symbol, of a countable signature. Let $\varphi = \varphi(x_1, \dots, x_n)$ be a first-order formula in the language \mathcal{L} , and let $\varepsilon \in [0, 1]$. Furthermore, let \mathcal{M} be a classical first-order model for \mathcal{M} and let \mathcal{D} be a probability measure on \mathcal{M} . Then we inductively define the notion of ε -truth, denoted by $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi$, as follows (where we leave the parameters implicit).

- (i) For every atomic formula φ :

$$(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi \text{ if } \mathcal{M} \models \varphi.$$

- (ii) We treat the logical connectives \wedge and \vee classically, e.g.

$$(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi \wedge \psi \text{ if } (\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi \text{ and } (\mathcal{M}, \mathcal{D}) \models_\varepsilon \psi.$$

- (iii) The existential quantifier is treated classically as well:

$$(\mathcal{M}, \mathcal{D}) \models_\varepsilon \exists x \varphi(x)$$

if there exists an $a \in \mathcal{M}$ such that $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi(a)$.

- (iv) The case of negation is split into sub-cases as follows:

(a) For φ atomic, $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \neg \varphi$ if $(\mathcal{M}, \mathcal{D}) \not\models_\varepsilon \varphi$.

(b) \neg distributes in the classical way over \wedge and \vee , e.g.

$$(\mathcal{M}, \mathcal{D}) \models_\varepsilon \neg(\varphi \wedge \psi) \text{ if } (\mathcal{M}, \mathcal{D}) \models_\varepsilon \neg\varphi \vee \neg\psi.$$

(c) $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \neg\neg\varphi$ if $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi$.

(d) $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \neg(\varphi \rightarrow \psi)$ if $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi \wedge \neg\psi$.

(e) $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \neg\exists x \varphi(x)$ if $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \forall x \neg\varphi(x)$.

(f) $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \neg\forall x \varphi(x)$ if $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \exists x \neg\varphi(x)$.

- (v) $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi \rightarrow \psi$ if $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \neg\varphi \vee \psi$.

- (vi) Finally, we define $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \forall x \varphi(x)$ if

$$\Pr_{\mathcal{D}}[a \in \mathcal{M} \mid (\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi(a)] \geq 1 - \varepsilon.$$

So, the most important change is that we do not treat the universal quantifier classically: instead of saying that we have $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi(a)$ for *all* elements $a \in \mathcal{M}$, we merely say that it holds for “many” of the elements, where “many” depends on the error parameter ε .

The main reason for this change is that we want our logic to be *learnable*, in the sense defined in Terwijn [11] (whose definition of learning is closely related to Valiant’s pac-model). Roughly speaking, learnable means that there is an algorithm which, given access to an oracle able to take samples from \mathcal{M} , on input φ outputs either φ or $\neg\varphi$ with the output being correct (i.e. its output holds in the model \mathcal{M}). Since the algorithm has to run in finite time, it can of course only take finitely many samples. We do not want to add the classical universal quantifier to our logic, since it is impossible to decide if a universal quantifier holds from just a finite amount of information. Therefore we take special care in defining our negation: we do not want $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \neg\exists x \varphi(x)$ to mean $(\mathcal{M}, \mathcal{D}) \not\models_\varepsilon \exists x \varphi(x)$, because the latter is equivalent to saying that $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi(x)$ holds *classically* for all $x \in \mathcal{M}$, which is exactly what we wanted to avoid. We define our negation in such a way that it

still behaves in a classical way on the propositional level, while it interchanges the existential and universal quantifiers.

One other important point is that both $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \forall x\varphi(x)$ and $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \exists x\neg\varphi(x)$ may hold simultaneously, i.e. both a formula and its negation might hold at the same time. Thus, ε -logic is *paraconsistent*. For our current work this has one important implication: it is no longer the case that φ is satisfiable if and only if its negation $\neg\varphi$ is not a tautology, as demonstrated in Example 2.5 below. Therefore, it is not clear a priori whether the complexity of these two problems are complementary to each other, and we need to study them separately.

To make sure that all necessary sets are measurable, we need to restrict ourselves to the right class of models.

Definition 2.2. Let \mathcal{L} be a first-order language of a countable signature, possibly containing the equality symbol, and let $\varepsilon \in [0, 1]$. Then an ε -model $(\mathcal{M}, \mathcal{D})$ for the language \mathcal{L} consists of a classical first-order \mathcal{L} -model \mathcal{M} together with a probability distribution \mathcal{D} over \mathcal{M} such that:

- (1) For all formulas $\varphi = \varphi(x_1, \dots, x_n)$ and all $a_1, \dots, a_{n-1} \in \mathcal{M}$, the set

$$\{a_n \in \mathcal{M} \mid (\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi(a_1, \dots, a_n)\}$$

is \mathcal{D} -measurable (i.e. all definable sets of dimension 1 are measurable).

- (2) All relations of arity n are \mathcal{D}^n -measurable (including equality, if it is in \mathcal{L}) and all functions of arity n are measurable as functions from $(\mathcal{M}^n, \mathcal{D}^n)$ to $(\mathcal{M}, \mathcal{D})$ (where \mathcal{D}^n denotes the n -fold product measure). In particular, constants are \mathcal{D} -measurable.

A *probability model* is a pair $(\mathcal{M}, \mathcal{D})$ that is an ε -model for every $\varepsilon \in [0, 1]$.

We note that (2) does not imply (1): the existential quantifier corresponds to taking a projection, and it is well-known that the projection of a measurable set need not be measurable. Nevertheless, the following result holds.

Proposition 2.3. Let $\varphi(x_1, \dots, x_n)$ be a universal formula and let $(\mathcal{M}, \mathcal{D})$ be an ε -model. Then

$$\{(a_1, \dots, a_n) \in \mathcal{M}^n \mid (\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi(a_1, \dots, a_n)\}$$

is \mathcal{D}^n -measurable.

Proof. First, one can use induction to prove that the lemma holds for propositional formulas ψ ; the base case is exactly (2). Next, let φ be a universal formula. By Proposition 2.8 below we may assume φ to be in prenex normal form; say $\varphi = \forall y_1 \dots \forall y_m \psi(x_1, \dots, x_n, y_1, \dots, y_m)$. We have just argued that the set

$$\{(a_1, \dots, a_n, b_1, \dots, b_m) \in \mathcal{M}^{n+m} \mid (\mathcal{M}, \mathcal{D}) \models_\varepsilon \psi(a_1, \dots, a_n, b_1, \dots, b_m)\}$$

is \mathcal{D}^{n+m} -measurable. The result now follows from repeatedly applying the fact that for any $1 \leq i \leq m$ and any $\mathcal{D}^{n+m-i+1}$ -measurable set X the function

$$(a_1, \dots, a_n, b_1, \dots, b_{m-i}) \mapsto \Pr_{\mathcal{D}}[b_{m-i+1} \in \mathcal{M} \mid (a_1, \dots, a_n, b_1, \dots, b_{m-i+1}) \in X]$$

is a \mathcal{D}^{n+m-i} -measurable function, see e.g. Bogachev [2, Theorem 3.4.1]. \square

In this paper, we will mostly be looking at languages not containing equality or function symbols. Functions and equality turn out to give major problems when trying to prove any results, mostly because they are not very compatible with measure-theoretic relations. For example, if f is an n -ary \mathcal{D} -measurable function and A is a \mathcal{D}^{n+1} -measurable set such that $\Pr_{\mathcal{D}^{n+1}}[\text{graph}(f) \Delta A] = 0$, then A does not need to be the graph of a function even though we cannot distinguish it from the graph of f using only the measure \mathcal{D} . We will not impose any restrictions

on our language beforehand, but will instead mention in each theorem to which languages it applies.

Definition 2.4. A formula $\varphi(x_1, \dots, x_n)$ is ε -satisfiable if there exists an ε -model $(\mathcal{M}, \mathcal{D})$ and there exist $a_1, \dots, a_n \in \mathcal{M}$ such that $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi$. Similarly, we say that φ is an ε -tautology or is ε -valid (notation: $\models_\varepsilon \varphi$) if for all probability models $(\mathcal{M}, \mathcal{D})$ and all $a_1, \dots, a_n \in \mathcal{M}$ it holds that $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi(a_1, \dots, a_n)$.

Example 2.5. Let Q be a unary predicate. Then $\varphi = \forall x Q(x) \vee \forall x \neg Q(x)$ is a $\frac{1}{2}$ -tautology. Namely, in every probability model, either the set on which Q holds or its complement has measure at least $\frac{1}{2}$. However, φ is not an ε -tautology for $\varepsilon < \frac{1}{2}$. Furthermore, both φ and $\neg\varphi$ are classically satisfiable and hence ε -satisfiable for every ε ; in particular we see that φ can be an ε -tautology while simultaneously $\neg\varphi$ is ε -satisfiable.

One might wonder why for satisfiability we only require ε -models, while for validity we look at the slightly stronger and less elegant probability models. To explain this, we will need the next theorem and proposition, which will both be used a lot throughout the paper.

Theorem 2.6. (Kuyper and Terwijn [8, Theorem 5.8]) *Let \mathcal{L} be a countable first-order language not containing equality or function symbols. Let Γ be an ε -satisfiable set of sentences. Then there exists an ε -model on $[0, 1]$ with the Lebesgue measure which ε -satisfies Γ . Furthermore, all relations in the new ε -model can be chosen to be Borel.*

Proposition 2.7. ([8, Proposition 5.1]) *Let \mathcal{M} be a first-order model that is a Polish space, and let \mathcal{D}_0 be a Borel probability measure on \mathcal{M} such that all relations and functions are \mathcal{D}_0^n -measurable. Then all definable sets are analytic. In particular, if we let \mathcal{D} be the completion of \mathcal{D}_0 , then $(\mathcal{M}, \mathcal{D})$ is a probability model.*

Now, if φ is any formula not containing function symbols which is ε -satisfiable in the sense that it has an ε -model, then Theorem 2.6 tells us that it has a model with Borel relations. We then know from Proposition 2.7 that this model is in fact a probability model. Thus: even though we did not require an ε -satisfiable formula to have a probability model, we get this model for free. Unfortunately we do not know of a similar result for ε -validity. For our results on validity in Kuyper [7] we really need probability models, hence the disparity between the definitions of ε -validity and ε -satisfiability.

The next result will also be used a lot.

Proposition 2.8. (Terwijn [11]) *Every formula φ is semantically equivalent to a formula φ' in prenex normal form; i.e. $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi \Leftrightarrow (\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi'$ for all $\varepsilon \in [0, 1]$ and all ε -models $(\mathcal{M}, \mathcal{D})$.*

We recall some more definitions from [8].

Definition 2.9. We will call a measure ν on a σ -algebra \mathcal{B} of subsets of N a *submeasure* of a measure μ on a σ -algebra \mathcal{A} of subsets of some set $M \supseteq N$ if for every $B \in \mathcal{B}$ there exists an $A_B \in \mathcal{A}$ such that $B = A_B \cap N$ and $\mu(A_B) = \nu(B)$.

Definition 2.10. An ε -submodel of an ε -model $(\mathcal{M}, \mathcal{D})$ is an ε -model $(\mathcal{N}, \mathcal{E})$ over the same language such that:

- \mathcal{N} is a submodel of \mathcal{M} in the classical sense,
- \mathcal{E} is a submeasure of \mathcal{D} .

We will denote this by $(\mathcal{N}, \mathcal{E}) \subset_\varepsilon (\mathcal{M}, \mathcal{D})$.

Definition 2.11. An *elementary ε -submodel* of an ε -model $(\mathcal{M}, \mathcal{D})$ is an ε -submodel $(\mathcal{N}, \mathcal{E})$ such that, for all formulas $\varphi = \varphi(x_1, \dots, x_n)$ and sequences $a_1, \dots, a_n \in \mathcal{N}$ we have:

$$(\mathcal{N}, \mathcal{E}) \models_{\varepsilon} \varphi(a_1, \dots, a_n) \Leftrightarrow (\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \varphi(a_1, \dots, a_n).$$

We will denote this by $(\mathcal{N}, \mathcal{E}) \prec_{\varepsilon} (\mathcal{M}, \mathcal{D})$.

Finally, we wish to mention the following result on reductions between different ε .

Theorem 2.12. (Kuyper and Terwijn [8, Theorem 7.1]) *Let \mathcal{L} be a countable first-order language not containing equality or function symbols. Then, for all rationals $0 \leq \varepsilon_0 \leq \varepsilon_1 < 1$, ε_0 -satisfiability many-one reduces to ε_1 -satisfiability.*

Theorem 2.13. ([8, Theorem 7.4]) *Let \mathcal{L} be a countable first-order language not containing equality or function symbols. Then, for all rationals $0 < \varepsilon_1 \leq \varepsilon_0 \leq 1$, ε_0 -satisfiability many-one reduces to ε_1 -satisfiability.*

Remark 2.14. As discussed in [8], we can perform these reductions per quantifier: that is, we do not need to apply the reduction to the entire formula, but can apply it on a per-quantifier basis. For example, if we let f_0 be the reduction from 0-satisfiability to $\frac{1}{2}$ -satisfiability and $f_{\frac{1}{4}}$ the reduction from 1-satisfiability to $\frac{1}{2}$ -satisfiability, then for all formulas $\varphi_0, \varphi_{\frac{1}{2}}, \varphi_{\frac{1}{4}}$ the expression “there exists a probability model $(\mathcal{M}, \mathcal{D})$ such that $(\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \varphi_{\varepsilon}$ for all $\varepsilon \in \{0, \frac{1}{4}, \frac{1}{2}\}$ ” is equivalent to “the formula $f_0(\varphi_0) \wedge f_{\frac{1}{4}}(\varphi_{\frac{1}{4}}) \wedge \varphi_{\frac{1}{2}}$ is $\frac{1}{2}$ -satisfiable”.

We note that 1-logic is fairly trivial: every formula in prenex normal form containing a universal quantifier is trivially true, so the only interesting fragment is the existential fragment, which is just the classical fragment. This paper will deal with the other two interesting cases: rational $\varepsilon \in (0, 1)$ in section 3, 5 and 7, and $\varepsilon = 0$ in section 6 and 8.

3. TOWARDS AN UPPER BOUND FOR ε -SATISFIABILITY

Our first goal is to show that (the relational fragment of) ε -satisfiability is Σ_1^1 . Together with the proof in section 7 that ε -satisfiability is Σ_1^1 -hard for rational $\varepsilon \in (0, 1)$ this will show that ε -satisfiability is Σ_1^1 -complete. There are multiple ways one could go about proving this. The first would be to reduce ε -satisfiability to Keisler’s logic $\mathcal{L}_{\omega P}$. In Hoover [4], it is shown that validity for this logic is Π_1^1 -complete, hence satisfiability for $\mathcal{L}_{\omega P}$ (which is dual to validity in the same way as for classical logic) is Σ_1^1 -complete. Keisler proves this through the use of a deduction system with infinitary deduction rules.

We will take a different, more direct approach. We will prove that there is a natural, equivalent formulation of ε -satisfiability of which we can directly see that it is Σ_1^1 . This equivalent formulation is the hidden heart of the completeness proof for $\mathcal{L}_{\omega P}$, and our method allows one to grasp the true essence of the proof. Of course, this comes at the price of not having a deduction system, but it is questionable how useful a deduction system with infinitary deduction rules is in the first place. Furthermore, because of the reduction to Keisler’s logic, this deduction system would talk about formulas containing Keisler’s quantifiers ($Px \geq r$). So, we would lose the advantage of using only the language of first-order logic. We will briefly come back to this point in Remark 3.7 and Remark 5.6.

As a first step, we will show that a certain weaker form of ε -satisfiability is Σ_1^1 .

Definition 3.1. A *weak ε -model* is a pair $(\mathcal{M}, \mathcal{D})$ which satisfies the conditions of Definition 2.2, except possibly for condition (2).

A *finitely additive model* is a pair $(\mathcal{M}, \mathcal{D})$ consisting of a first-order model \mathcal{M} and a finitely additive measure \mathcal{D} over \mathcal{M} .

We say that $\varphi(x_1, \dots, x_n)$ is *weakly ε -satisfiable* (respectively *finite additively ε -satisfiable*) if there exist a weak ε -model (respectively finitely additive model) $(\mathcal{M}, \mathcal{D})$ and $a_1, \dots, a_n \in \mathcal{M}$ such that $(\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \varphi(a_1, \dots, a_n)$.¹

For finitely additive models, the following result by Tarski turns out to be useful.

Theorem 3.2. (Tarski) *Every finitely additive measure \mathcal{D} on a set X can be extended to a finitely additive measure \mathcal{D}' on the power set $\mathcal{P}(X)$.*

Proof. See Birkhoff [1, p. 185]. □

This result also explains why we did not impose any measurability conditions on our finitely additive models: if φ is satisfied in some finitely additive model $(\mathcal{M}, \mathcal{D})$, then it is also satisfied in the model $(\mathcal{M}, \mathcal{D}')$ in which every set is measurable.

At first sight, finitely additive ε -satisfiability might seem to be much weaker than weak ε -satisfiability. Surprisingly, it turns out that they are equivalent.

Theorem 3.3. *A formula φ is weakly ε -satisfiable if and only if it is finite additively ε -satisfiable.*

Proof. Clearly any weak ε -model is also a finitely additive model, so if φ is weakly ε -satisfiable it is certainly finite additively ε -satisfiable.

For the converse, assume $(\mathcal{M}, \mathcal{D})$ is a finitely additive model ε -satisfying φ . Extend \mathcal{D} to a finitely additive measure \mathcal{D}' on $\mathcal{P}(X)$ using Theorem 3.2 and take an ultrapower $(\mathcal{N}, \mathcal{E})$ of $(\mathcal{M}, \mathcal{D}')$ (as defined in Kuyper and Terwijn [8, Definition 8.8]). Then $(\mathcal{N}, \mathcal{E})$ is a weak ε -model by [8, Corollary 8.10] and it ε -satisfies φ by [8, Theorem 8.9]. □

In [8, Example 4.5] it was shown that, in general, not every ε -satisfiable sentence has a countable ε -model (i.e. the Downward Löwenheim-Skolem theorem does not hold in the usual sense). In contrast, this does hold if we look at finitely additive models.

Definition 3.4. Let \mathcal{D} be a probability measure. We say that \mathcal{D} is *countable* if the Boolean algebra on which \mathcal{D} is defined is countable.

Theorem 3.5. (Downward Löwenheim-Skolem theorem for finitely additive ε -satisfiability) *Let \mathcal{L} be a countable first-order language. Let $(\mathcal{M}, \mathcal{D})$ be a finitely additive model and let $X \subseteq \mathcal{M}$ be countable. Then there exists a finitely additive model*

$$(\mathcal{N}, \mathcal{E}) \prec_{\varepsilon} (\mathcal{M}, \mathcal{D})$$

such that $X \subseteq \mathcal{N}$ and such that \mathcal{N} and \mathcal{E} are countable.

Proof. By Theorem 3.2, we may without loss of generality assume that the domain of \mathcal{D} is $\mathcal{P}(\mathcal{M})$. We will define a sequence $X_0 \subseteq X_1 \subseteq \dots$ of countable subsets of \mathcal{M} and let \mathcal{N} be the restriction of \mathcal{M} to $\bigcup_{n \in \omega} X_n$.

Let X_0 consist of X together with the interpretation $c^{\mathcal{M}}$ of all constants. Next, given X_n , we show how to define X_{n+1} . Let \mathcal{B}_{n+1} be the Boolean algebra generated

¹Because we did not impose any measurability conditions on our finitely additive model, it could be the case that the set occurring in case (vi) of Definition 2.1 is not measurable. For finitely additive models we therefore say that $(\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \forall x \varphi(x)$ if $\{a \in \mathcal{M} \mid (\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \varphi(a)\}$ is \mathcal{D} -measurable and

$$\Pr_{\mathcal{D}}[a \in \mathcal{M} \mid (\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \varphi(a)] \geq 1 - \varepsilon,$$

i.e. we let a universal quantifier be false if the corresponding set is not measurable.

by the ε -definable subsets of $(\mathcal{M}, \mathcal{D})$ using parameters from X_n . Since X_n is countable, \mathcal{B}_{n+1} will also be countable. Fix an element $x_B \in B$ for every non-empty $B \in \mathcal{B}_{n+1}$. Now let

$$X_{n+1} = X_n \cup \{x_B \mid B \in \mathcal{B}_{n+1}\} \cup \{f^{\mathcal{M}}(a_1, \dots, a_m) \mid f \in \mathcal{L}, f \text{ is } m\text{-ary and } a_1, \dots, a_m \in X_n\}.$$

Then X_{n+1} is countable and $X_n \subseteq X_{n+1} \subseteq \mathcal{M}$.

As announced above, we let \mathcal{N} be $\mathcal{M} \upharpoonright \bigcup_{n \in \omega} X_n$. Note that if $a_1, \dots, a_m \in \mathcal{N}$, then there is some $n \in \omega$ such that $a_1, \dots, a_m \in X_n$. Then $f^{\mathcal{M}}(a_1, \dots, a_m) \in X_{n+1} \subseteq \mathcal{N}$ so the functions on \mathcal{N} are well-defined.

For every B which is in the Boolean algebra generated by the ε -definable subsets of $(\mathcal{M}, \mathcal{D})$ using parameters from \mathcal{N} , let $\mathcal{E}(B \cap \mathcal{N}) = \mathcal{D}(B)$. Note that such a B uses only finitely parameters and hence is already ε -definable using parameters from some X_n . We show that the finitely additive measure \mathcal{E} is well-defined. To this end, assume that B_1, B_2 in the Boolean algebra described above are such that $B_1 \cap \mathcal{N} = B_2 \cap \mathcal{N}$. Let X_n be such that B_1, B_2 can be ε -defined using parameters from X_n . If $B_1 \neq B_2$, then either $B_1 \cap (\mathcal{M} \setminus B_2) \neq \emptyset$ or $(\mathcal{M} \setminus B_1) \cap B_2 \neq \emptyset$. In both cases, one of these points gets added to X_{n+1} , hence $B_1 \cap \mathcal{N} \neq B_2 \cap \mathcal{N}$, a contradiction. So, $B_1 = B_2$ and therefore $\mathcal{D}(B_1) = \mathcal{D}(B_2)$, as desired.

It remains to show that $(\mathcal{N}, \mathcal{E})$ is an elementary finitely additive ε -submodel of $(\mathcal{M}, \mathcal{D})$, i.e. that for all sequences $a_1, \dots, a_n \in \mathcal{N}$ and for all formulas $\varphi(x_1, \dots, x_n)$ we have

$$(\mathcal{N}, \mathcal{E}) \models_{\varepsilon} \varphi(a_1, \dots, a_n) \Leftrightarrow (\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \varphi(a_1, \dots, a_n).$$

We prove this using induction over the formulas in prenex normal form. For propositional formulas, this is directly clear. For the existential case, note that

$$(\mathcal{N}, \mathcal{E}) \models_{\varepsilon} \exists x \psi(x, a_1, \dots, a_n)$$

clearly implies that this also holds in $(\mathcal{M}, \mathcal{D})$. For the converse, assume

$$(\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \exists x \psi(x, a_1, \dots, a_n).$$

Let X_m be such that $a_1, \dots, a_n \in X_m$. By construction, we have added a point from the non-empty set

$$\{b \in \mathcal{M} \mid (\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \psi(b, a_1, \dots, a_n)\}$$

to \mathcal{N} . So, there exists some point $b \in \mathcal{N}$ such that $(\mathcal{N}, \mathcal{E}) \models_{\varepsilon} \psi(b, a_1, \dots, a_n)$. Using the induction hypothesis, this directly implies that

$$(\mathcal{N}, \mathcal{E}) \models_{\varepsilon} \exists x \psi(x, a_1, \dots, a_n).$$

For the universal case, let $\varphi = \forall x \psi(x, x_1, \dots, x_n)$. Let

$$B = \{x \in \mathcal{M} \mid (\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \psi(x, b_1, \dots, b_n)\},$$

$$C = \{x \in \mathcal{N} \mid (\mathcal{N}, \mathcal{E}) \models_{\varepsilon} \psi(x, b_1, \dots, b_n)\}.$$

Then by induction hypothesis we have

$$\begin{aligned} C &= \{x \in \mathcal{N} \mid (\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \psi(x, b_1, \dots, b_n)\} \\ &= B \cap \mathcal{N}. \end{aligned}$$

From this we see that $\mathcal{E}(C) = \mathcal{D}(B)$, and hence

$$(\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \forall x \psi(x, b_1, \dots, b_n) \Leftrightarrow (\mathcal{N}, \mathcal{E}) \models_{\varepsilon} \forall x \psi(x, b_1, \dots, b_n).$$

This concludes the induction. \square

Theorem 3.6. *Let \mathcal{L} be a countable language and let ε be rational. Then weak ε -satisfiability is Σ_1^1 .*

Proof. From Theorem 3.3 and Theorem 3.5 we see that a formula φ is weakly ε -satisfiable if and only if it is ε -satisfied in a countable finitely additive model. We sketch how to express the latter as a Σ_1^1 -formula. Namely, we can express this by saying:

- there exist interpretations $R^{\mathcal{M}} \subseteq \omega^n$ for the relations R and interpretations for the constants (note that this can be done through one second-order existential quantifiers, by using a pairing function);
- there exist $Q \subseteq \omega$ (representing the Boolean algebra) and $\mathcal{D} : \omega \rightarrow \mathbb{R}$;

such that:

- the sets $Q_n := \{m \in \omega \mid \langle n, m \rangle \in Q\}$ for $n \in \omega$ form a Boolean algebra;
- the function mapping Q_n to $\mathcal{D}(n)$ is a finitely additive measure;
- φ is ε -satisfied in $(\mathcal{M}, \mathcal{D})$.

It can be directly verified that all the items in the second list can be expressed as arithmetical formulas. Combining this with the existential second-order quantifiers from the first list, we obtain the desired result. \square

Remark 3.7. In Keisler's logic $\mathcal{L}_{\omega P}$, finitely additive models can be obtained by taking a maximal consistent set of formulas in the calculus he calls *weak $\mathcal{L}_{\omega P}$* (see Keisler [6, Definition 1.4.1 and Theorem 1.5.3]). By extending his logic with rules for the existential quantifier, we would be able to obtain something similar: note that $\varphi = \forall x_1 \exists x_2 \dots \forall x_n \psi(x_1, \dots, x_n)$ is ε -satisfiable if and only if $(Px_1 \geq 1 - \varepsilon) \exists x_2 \dots (Px_n \geq 1 - \varepsilon) \psi(x_1, \dots, x_n)$ is satisfiable. In fact, we can base our system on the finitary fragment of weak $\mathcal{L}_{\omega P}$ (i.e. we only need Keisler's axioms A1-A5) since our formulas correspond to formulas in Keisler's positive fragment. From this we find that weak ε -satisfiability is in fact even Π_1^0 . Since we are not interested in weak ε -satisfiability and merely see it as a stepping stone for our results on (regular) ε -satisfiability, which we have already shown to be Σ_1^1 -hard, we will not pursue this matter any further.

Our next goal is to extend Theorem 3.6 to regular ε -satisfiability, instead of just weak ε -satisfiability. Before we do so, we will first look at some results concerning Skolemisation in ε -logic.

4. SKOLEMISATION IN ε -LOGIC

For classical satisfiability, we can eliminate existential quantifiers by introducing Skolem functions. Clearly, we can also do something similar here: for example, if a statement $\forall x \exists y \varphi(x, y)$ is 0-true in some model $(\mathcal{M}, \mathcal{D})$, we can extend \mathcal{M} to a model \mathcal{N} such that $(\mathcal{N}, \mathcal{D}) \models_0 \forall x \varphi(x, f(x))$ holds. We do not even need that $\mathcal{M} \models \varphi(x, f(x))$ holds for every $x \in \mathcal{M}$ for which $\mathcal{M} \models \exists y \varphi(x, y)$ holds — it is enough if this is true for almost all such x .

However, we have required our functions to be measurable. It is not directly clear that we can also pick our Skolem functions in a measurable way. Our next result shows that this is possible.

Theorem 4.1. *Let φ be a formula not containing function symbols. Then φ is ε -satisfiable if and only if its Skolemisation is ε -satisfiable.*

Proof. We need to show that we can pick the Skolem functions in a measurable way. First, assume φ does not contain the equality symbol. Now, if φ is ε -satisfiable, then by Theorem 2.6 we may assume it is ε -satisfied in a model (\mathcal{M}, λ) on $[0, 1]$ with the Lebesgue measure, with Borel relations. Then all definable sets are analytic by Proposition 2.7.

Let $\varphi = \forall x_1 \exists y_1 \dots \forall x_n \exists y_n \psi(x_1, \dots, x_n, y_1, \dots, y_n)$ be in prenex normal form. Now, we use recursion over $1 \leq i \leq n$ to define Borel measurable Skolem functions f_1, f_2, \dots, f_n , by performing the following steps:

- (1) Use the Jankov and von Neumann Uniformisation Theorem (see e.g. Kechris [5, Theorem 18.1]²) to find a Lebesgue-measurable *uniformising function* g_i : that is, for all $x_1, x_2, \dots, x_i \in [0, 1]$, if there exists an $y_i \in [0, 1]$ such that

$$(\mathcal{M}, \lambda) \models_{\varepsilon} \forall x_{i+1} \dots \exists y_n \psi(x_1, \dots, x_n, \\ f_1(x_1), \dots, f_{i-1}(x_1, \dots, x_{i-1}), y_i, \dots, y_n),$$

then we have, if we denote $y = g_i(x_1, x_2, \dots, x_i)$:

$$(\mathcal{M}, \lambda) \models_{\varepsilon} \forall x_{i+1} \dots \exists y_n \psi(x_1, \dots, x_n, \\ f_1(x_1), \dots, f_{i-1}(x_1, \dots, x_{i-1}), y, y_{i+1}, \dots, y_n).$$

- (2) Let f_i be a Borel function which is equal to g_i almost everywhere; that such functions exist is shown in Bogachev [2, Proposition 2.1.11], or can easily be proven using Lusin's Theorem on measurable functions (see e.g. Kechris [5, Theorem 17.12]³).

In the case that φ does contain the equality symbol, we may assume it is satisfied by a model on $[0, r]$ plus atoms, see Kuyper and Terwijn [8, Theorem 5.9]. We can apply the construction above on $[0, r]$ and use AC to choose the values f_i takes in the atoms; since there are only countably many atoms, this does not alter the measurability of the f_i . \square

It turns out we can do even more. We assumed that our formula φ did not contain any function symbols, but now its Skolemisation does contain function symbols. The next proposition shows that we can also construct a kind of Skolemisation which does not need function symbols.

Theorem 4.2. *Let \mathcal{L} be a language not containing equality or function symbols (but it may contain constant symbols). Then there exists a language \mathcal{L}' only containing relation symbols and a computable function mapping each formula φ in the language \mathcal{L} to a universal formula φ' in the language \mathcal{L}' such that for every $\varepsilon \in [0, 1]$: φ is ε -satisfiable if and only if φ' is ε -satisfiable.*

Proof. The idea of the proof is roughly as follows: because universal quantifiers only talk about measure, we can always change the interpretation of a relation on a set of measure zero without affecting the truth of the universal quantifiers. Thus, we can always add witnesses of measure zero, as long as there are no contradicting statements about a witness. Let R be a relation symbol, say of arity k . To express which statements hold about which witnesses, we add a new copy of R for every atomic formula of the form $R(t_1, \dots, t_k)$ up to a permutation of the universally bound variables. We then form the formula φ' by replacing the atomic formula by this new copy of R . If φ is satisfiable, then φ' is satisfiable just by interpreting the new relation symbol by the interpretation of $R(t_1, \dots, t_k)$. Conversely, if φ' is satisfiable, we pick our witnesses in a set of measure zero (say, the Cantor set) and recombine all the relations into one relation R . We will now give the proof in full detail.

²Kechris only states that we get a partial function with as domain exactly the (analytic) set of those x_1, \dots, x_i for which an y_i as above exists, but we can easily extend this to a total Lebesgue measurable function by letting its value be 0 outside this set.

³Kechris only states this theorem for Borel measures, but it holds for the Lebesgue measure with exactly the same proof.

Let R be a relation symbol, say of arity k . First, we enlarge our language by adding a copy of R for every atomic formula of the form $R(t_1, \dots, t_k)$, up to a permutation of the universally bound variables. To this end, fix an enumeration c_0, c_1, \dots of the constant symbols in \mathcal{L} . Let A_R be the set of functions from $\{1, \dots, k\}$ to $\{0\} \cup (\{1, 2\} \times \omega)$. To every atomic formula $R(t_1, \dots, t_k)$ we assign a function $\alpha \in A_R$ by letting $\alpha(i) = 0$ if t_i equals some (universally bound) x_j , by letting $\alpha(i) = (1, j)$ if t_i equals (the existentially bound) y_j and by letting $\alpha(i) = (2, j)$ if t_i equals the constant symbol c_j . During the proof, we will say that the α constructed in this way is the α corresponding to the atomic formula $R(t_1, \dots, t_k)$. Now, for every relation symbol R and every $\alpha \in A_R$, we introduce a new relation symbol R_α , where the arity of R_α equals $|\{1 \leq i \leq k \mid \alpha(i) = 0\}| + \max\{j \in \omega \mid (1, j) \in \text{ran}(\alpha)\}$.

Let $\varphi = \forall x_1 \exists y_1 \dots \exists y_n \forall x_n \psi(x_1, \dots, x_n, y_1, \dots, y_n)$ be a formula in \mathcal{L} in prenex normal form. We describe how to form $\varphi' = \forall x_1 \dots \forall x_n \psi'(x_1, \dots, x_n)$ from φ . Let $R(t_1, \dots, t_k)$ be an atomic formula occurring in φ and let α be the function corresponding to this atomic formula, as described above. Let $m = \max\{j \in \omega \mid (1, j) \in \text{ran}(\alpha)\}$ and let $s_1, \dots, s_{k'}$ be the subsequence of t_1, \dots, t_k consisting of just those t_i which equal a universally bound variable x_j (during the proof, we will say that $s_1, \dots, s_{k'}$ is the *universal subsequence* of t_1, \dots, t_k). Then we form ψ' from ψ by replacing each $R(t_1, \dots, t_k)$ with $R_\alpha(s_1, \dots, s_{k'}, x_1, \dots, x_m)$.

We show that φ' is as desired. First, assume φ is ε -satisfiable. So, by Theorem 2.6 there exists an ε -model (\mathcal{M}, λ) with the Lebesgue measure λ and Borel relations ε -satisfying φ . By the previous proposition, for each y_i we can find a Borel Skolem function $f_i : [0, 1]^i \rightarrow [0, 1]$. Let R be a relation symbol and let $\alpha \in A_R$. We interpret $R_\alpha^{\mathcal{N}}(s_1, \dots, s_{k'}, x_1, \dots, x_m)$ as $R^{\mathcal{M}}(v_1, \dots, v_k)$, where

$$v_i := \begin{cases} s_j & \text{if } \alpha(i) = 0 \text{ and } |\{s \leq j \mid \alpha(s) = 0\}| = i \\ f_j(x_1, \dots, x_j) & \text{if } \alpha(i) = (1, j) \\ c_j & \text{if } \alpha(i) = (2, j). \end{cases}$$

We claim that $R_\alpha^{\mathcal{N}}$ is both analytic and co-analytic; then it is Borel by Souslin's theorem (see e.g. Kechris [5, Theorem 14.11]). Let β be the function such that

$$\beta_i(s_1, \dots, s_{k'}, y_1, \dots, y_m) := \begin{cases} s_j & \text{if } \alpha(i) = 0 \text{ and } |\{s \leq j \mid \alpha(s) = 0\}| = i \\ y_j & \text{if } \alpha(i) = (1, j) \\ c_j & \text{if } \alpha(i) = (2, j), \end{cases}$$

i.e. β 'restores' the atomic formula $R(t_1, \dots, t_k)$ in the sense that $R(t_1, \dots, t_k)$ is the same as $R(\beta(s_1, \dots, s_{k'}, x_1, \dots, x_m))$. Finally, let B_j be the graph of f_j . Then we have

$$R_\alpha^{\mathcal{N}} = \{(s_1, \dots, s_{k'}, x_1, \dots, x_m) \mid \exists y_1 \dots y_m (\beta(s_1, \dots, s_{k'}, y_1, \dots, y_m) \in R^{\mathcal{M}} \wedge (x_1, y_1) \in B_1 \wedge \dots \wedge (x_m, y_m) \in B_m)\}$$

so it is analytic, and

$$R_\alpha^{\mathcal{N}} = \{(s_1, \dots, s_{k'}, x_1, \dots, x_m) \mid \forall y_1 \dots y_m ((x_1, y_1) \in B_1 \wedge \dots \wedge (x_m, y_m) \in B_m) \rightarrow \beta(s_1, \dots, s_{k'}, y_1, \dots, y_m) \in R^{\mathcal{M}}\}$$

so it is also co-analytic. Therefore, (\mathcal{N}, λ) is an ε -model by Proposition 2.7.

We know that

$$(\mathcal{M}, \lambda) \models_\varepsilon \forall x_1 \dots \forall x_n \psi(x_1, \dots, x_n, f_1(x_1), \dots, f_n(x_1, \dots, x_n))$$

because the f_i are Skolem functions. Let $R(t_1, \dots, t_k)$ be an atomic formula with corresponding α and let $s_1, \dots, s_{k'}$ be the universal subsequence of t_1, \dots, t_k . Let

$R(u_1, \dots, u_k)$ be the atomic formula where $f_i(x_1, \dots, x_i)$ is substituted for each y_i . Then

$$(3) \quad \begin{aligned} & \{(x_1, \dots, x_n) \in \mathcal{M}^n \mid R^{\mathcal{M}}(u_1, \dots, u_k)\} \\ & = \{(x_1, \dots, x_n) \in \mathcal{N}^n \mid R_{\alpha}^{\mathcal{N}}(s_1, \dots, s_{k'}, x_1, \dots, x_m)\} \end{aligned}$$

by the definition of $R_{\alpha}^{\mathcal{N}}$. Then it follows using formula induction that

$$\begin{aligned} & \{(x_1, \dots, x_n) \in \mathcal{M}^n \mid \mathcal{M} \models \psi(x_1, \dots, x_m, f_1(x_1), \dots, f_n(x_1, \dots, x_n))\} \\ & = \{(x_1, \dots, x_n) \in \mathcal{N}^n \mid \mathcal{N} \models \psi'(x_1, \dots, x_m)\}, \end{aligned}$$

where the only interesting case is the atomic case, which is exactly (3). So, we see that also $(\mathcal{N}, \lambda) \models_{\varepsilon} \varphi'$ holds.

Conversely, assume we have a model (\mathcal{N}, λ) with the Lebesgue measure and Borel relations satisfying φ' . We now need to put all the $R_{\alpha}^{\mathcal{N}}$ together into one $R^{\mathcal{M}}$. The basic idea is that, because universal quantifiers only talk about measure, we can glue the relations together on a set of measure zero. We will pick all of our witnesses inside Cantor space, which has measure 0 inside $[0, 1]$. To this end, let $\zeta : \mathcal{C} \rightarrow \bigcup_{1 \leq i \leq n} [0, 1]^i$ be a Borel isomorphism of Cantor space (as subset of $[0, 1]$) with copies of unit boxes of increasing dimension (that such an isomorphism exists follows from Kechris [5, Theorem 15.6]). Also fix an injective function $\eta : \omega \hookrightarrow [0, 1] \setminus \mathcal{C}$, which we will use to interpret the constant symbols.

We construct a model \mathcal{M} on $[0, 1]$. For any constant symbol c_i , let $c_i^{\mathcal{M}} = \eta(i)$. Next we show how to define $R^{\mathcal{M}}$ of arity k . Let $a_1, \dots, a_k \in [0, 1]$. Let $\beta(i) = (1, j)$ if $a_i \in \mathcal{C}$ and $\zeta(a_i)$ has length j , let $\beta(i) = (2, j)$ if $a_i \in \text{ran}(\eta)$ and $\eta^{-1}(a_i) = j$ and finally let $\beta(i) = 0$ if neither of these cases hold. We say that this β corresponds to the sequence a_1, \dots, a_k . Let $b_1, \dots, b_{k'}$ be the subsequence of a_1, \dots, a_k obtained by taking just those a_i satisfying $\beta(i) = 0$. Let $m = \max\{j \in \omega \mid (1, j) \in \text{ran}(\beta)\}$ and let i be the least $1 \leq j \leq k$ such that $\beta(j) = (1, m)$. Now let $R^{\mathcal{M}}(a_1, \dots, a_k)$ be defined as $R_{\beta}^{\mathcal{N}}(b_1, \dots, b_{k'}, \zeta_1(a_i), \dots, \zeta_m(a_i))$. We can prove that $R^{\mathcal{N}}$ is Borel using Souslin's theorem, in the same way as above.

Next, let $R(t_1, \dots, t_k)$ be an atomic formula and let α correspond to this atomic formula. Furthermore, let $s_1, \dots, s_{k'}$ be the universal subsequence of t_1, \dots, t_k . Let $m = \max\{j \in \omega \mid (1, j) \in \text{ran}(\alpha)\}$. Let $b_1, \dots, b_k \in [0, 1] \setminus (\mathcal{C} \cup \text{ran}(\eta))$ and consider the sequence $a_i = t_i^{\mathcal{M}}[x_j := b_j, y_j := \zeta^{-1}(b_1, \dots, b_j)]$. Then it is directly verified that the β corresponding to the sequence a_1, \dots, a_k is equal to α . Thus, by the definition of $R^{\mathcal{N}}$ we see that

$$R^{\mathcal{M}}(t_1, \dots, t_k)[x_i := b_i, y_i := \zeta^{-1}(b_1, \dots, b_i)]$$

holds if and only if

$$R_{\alpha}^{\mathcal{N}}(s_1, \dots, s_{k'}, x_1, \dots, x_m)[x_i := b_i]$$

holds. So, from the construction of ψ' we see that

$$\begin{aligned} & \{\vec{x} \in ([0, 1] \setminus (\mathcal{C} \cup \text{ran}(\eta)))^n \mid \mathcal{M} \models \psi(x_1, \dots, x_n, \zeta^{-1}(x_1), \dots, \zeta^{-1}(x_1, \dots, x_n))\} \\ & = \{\vec{x} \in ([0, 1] \setminus (\mathcal{C} \cup \text{ran}(\eta)))^n \mid \mathcal{N} \models \psi'(x_1, \dots, x_n)\}. \end{aligned}$$

Because $\text{ran}(\eta)$ is countable we know that $\mathcal{C} \cup \text{ran}(\eta)$ has Lebesgue-measure 0, so using the fact that φ' holds in (\mathcal{N}, λ) we then directly see that

$$(\mathcal{M}, \lambda) \models_{\varepsilon} \forall x_1 \forall x_2 \dots \forall x_n \psi(x_1, \dots, x_n, \zeta^{-1}(x_1), \dots, \zeta^{-1}(x_1, \dots, x_n))$$

holds and therefore $(\mathcal{M}, \lambda) \models_{\varepsilon} \varphi$ holds, as desired. \square

5. ε -SATISFIABILITY IS Σ_1^1

In the previous section we saw that every formula in a language not containing equality or function symbols is equisatisfiable to a universal relational formula, through a computable transformation. Therefore, in our proof of the Σ_1^1 upper bound for ε -satisfiability without equality or function symbols, we will only need to consider the universal relational fragment.

In Section 3 we already showed that weak ε -satisfiability is Σ_1^1 . However, it is not directly clear how to extend Theorem 3.6 to regular ε -satisfiability: the relations in the ε -model we built there using an ultrapower will, in general, not be measurable. Thus, if we want a result similar to Theorem 3.3 for ε -satisfiability, we will need to impose extra conditions on our finitely additive models. Our extra conditions will be motivated by the next lemma.

Lemma 5.1. *Let \mathcal{D} be a probability measure on a set X and let R be a \mathcal{D}^n -measurable set. Then for every $m \in \omega$ there exists $k \in \omega$ and \mathcal{D} -measurable sets $X_{i,j}$ for $1 \leq i \leq k$, $1 \leq j \leq n$ such that for every function $f : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$:*

$$\Pr_{\mathcal{D}^n}[(a_1, \dots, a_n) \in \mathcal{M}^n \mid (a_{f(1)}, \dots, a_{f(n)}) \in (R^{\mathcal{M}} \Delta ((X_{1,1} \times \dots \times X_{1,n}) \cup \dots \cup (X_{k,1} \times \dots \times X_{k,n})))]) \leq \frac{1}{m}.$$

Proof. Fix an enumeration f_1, \dots, f_s of all functions from $\{1, \dots, n\}$ to $\{1, \dots, n\}$. Let $H(R)$ be the statement of the lemma, i.e. $H(R)$ states that: for every $m \in \omega$ there exists $k \in \omega$ and \mathcal{D} -measurable sets $X_{i,j}$ for $1 \leq i \leq k$, $1 \leq j \leq n$ such that for every function $f : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$:

$$\Pr_{\mathcal{D}^n}[(a_1, \dots, a_n) \in \mathcal{M}^n \mid (a_{f(1)}, \dots, a_{f(n)}) \in (R^{\mathcal{M}} \Delta ((X_{1,1} \times \dots \times X_{1,n}) \cup \dots \cup (X_{k,1} \times \dots \times X_{k,n})))]) \leq \frac{1}{m}.$$

We show that the class of subsets $R \subseteq X^n$ for which $H(R)$ holds is a monotone class containing all finite unions of boxes, from which then follows that $H(R)$ holds for all \mathcal{D}^n -measurable sets R by the monotone class theorem (see e.g. Bogachev [2, Theorem 1.9.3]).

Clearly, $H(R)$ holds if R is a finite union of boxes. Next, let $R = \bigcap_{j \in \omega} R_j$ with $R_0 \supseteq R_1 \supseteq \dots$ be a monotone decreasing sequence and assume that $H(R_j)$ holds for all these R_j . Fix $m \in \omega$. By countable additivity we can determine, for each $1 \leq i \leq s$ an $l_i \in \omega$ such that

$$\Pr_{\mathcal{D}^n} \left[(a_1, \dots, a_n) \in \mathcal{M}^n \mid (a_{f_i(1)}, \dots, a_{f_i(n)}) \in \left(\bigcap_{j \leq l_i} R_j \right) \Delta R \right] \leq \frac{1}{2m}.$$

Let l be the maximum of these l_i . By our hypotheses $H(R_j)$ we can determine for each $j \in \omega$ a $k_j \in \omega$ and measurable sets $X_{u,v}^j$ such that for every $1 \leq i \leq s$ we have

$$\Pr_{\mathcal{D}^n}[(a_1, \dots, a_n) \in \mathcal{M}^n \mid (a_{f_i(1)}, \dots, a_{f_i(n)}) \in (R_j^{\mathcal{M}} \Delta ((X_{1,1}^j \times \dots \times X_{1,n}^j) \cup \dots \cup (X_{k_j,1}^j \times \dots \times X_{k_j,n}^j)))]) \leq \frac{1}{2^{j+1}m}.$$

Now consider the set

$$A = \bigcap_{j \leq l} ((X_{1,1}^j \times \dots \times X_{1,n}^j) \cup \dots \cup (X_{k_j,1}^j \times \dots \times X_{k_j,n}^j)).$$

Then A is of the desired form (i.e. it can be written as a finite union of Cartesian products): by distributivity, A is equal to a finite union of expressions of the form

$$(X_{u_1,1}^1 \times \cdots \times X_{u_1,k_1}^1) \cap \cdots \cap (X_{u_l,1}^l \times \cdots \times X_{u_l,k_l}^l)$$

and this expression is equal to

$$(X_{u_1,1}^1 \cap \cdots \cap X_{u_l,1}^l) \times \cdots \times (X_{u_1,n}^1 \cap \cdots \cap X_{u_l,n}^l).$$

Furthermore, it is also directly verified that for every $1 \leq i \leq s$ we have

$$\Pr_{\mathcal{D}^n}[(a_1, \dots, a_n) \in \mathcal{M}^n \mid (a_{f_i(1)}, \dots, a_{f_i(n)}) \in R \Delta A] \leq \frac{1}{2m} + \sum_{j=1}^l \frac{1}{2^{j+1}m} \leq \frac{1}{m}.$$

So, this shows that $H(R)$ holds.

The case in which R is the union of a monotone increasing sequence can be proven in a similar way, which completes the proof. \square

Now, the important idea is that this property of ‘having finite approximations’ can be expressed in the language of ε -logic, and is hence preserved under taking ultrapowers. Therefore, if we take an ultrapower as in the proof of Theorem 3.3, this ultrapower will also possess these finite approximations. By taking a suitable limit of a sequence approximating the relation R we will obtain a *measurable* relation S that coincides with R almost everywhere. The next lemma expresses that such an approximation is good enough for our purposes.

Lemma 5.2. *Let $(\mathcal{M}, \mathcal{D})$ be an ε -model and let R, S be two n -ary relation symbols such that for every function $f : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ we have:*

$$(\mathcal{M}, \mathcal{D}) \models_0 \forall x_1 \dots \forall x_n (R(x_{f(1)}, \dots, x_{f(n)}) \leftrightarrow S(x_{f(1)}, \dots, x_{f(n)})).$$

Then for every universal formula φ not containing function symbols or constant symbols, if we let φ' be the formula where every occurrence of R is replaced by S : $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi$ if and only if $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi'$.

Proof. First we prove using formula induction that for every propositional formula $\psi(x_1, \dots, x_m)$:

$$(\mathcal{M}, \mathcal{D}) \models_0 \forall x_1 \dots \forall x_n (\psi(x_1, \dots, x_n) \leftrightarrow \psi'(x_1, \dots, x_n)).$$

For atomic formulas this follows from our assumption: since our language does not contain function or constant symbols, every atomic subformula of φ is of the form $R(y_1, \dots, y_n)$ where $y_1, \dots, y_n \in \{x_1, \dots, x_n\}$. Now let f be the function such that $y_i = x_{f(i)}$ and apply the hypothesis. The other cases in the formula induction are direct.

Because of our definition of \rightarrow , which is only classical on the propositional level, we see that $\alpha \leftrightarrow \beta$ which is defined as $(\alpha \rightarrow \beta) \wedge (\beta \rightarrow \alpha)$ is also only classical on the propositional level. To ease our notation, for the rest of the proof we let \Leftrightarrow be the connective expressing ε -equivalence, i.e. we define $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \alpha \Leftrightarrow \beta$ if either $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \alpha \wedge \beta$ or $(\mathcal{M}, \mathcal{D}) \not\models_\varepsilon \alpha \vee \beta$. In particular \Leftrightarrow coincides with \leftrightarrow on the propositional level, so we see that

$$(\mathcal{M}, \mathcal{D}) \models_0 \forall x_1 \dots \forall x_n (\psi(x_1, \dots, x_n) \Leftrightarrow \psi'(x_1, \dots, x_n)).$$

Next, we note that universal quantifiers distribute over \Leftrightarrow in 0-logic, i.e. we have that $(\mathcal{M}, \mathcal{D}) \models_0 \forall x (\alpha(x) \Leftrightarrow \beta(x))$ if and only if $(\mathcal{M}, \mathcal{D}) \models_0 (\forall x \alpha(x)) \Leftrightarrow (\forall x \beta(x))$. From this we see that for every propositional formula ψ we have

$$(\mathcal{M}, \mathcal{D}) \models_0 (\forall x_1 \dots \forall x_n \psi(x_1, \dots, x_n)) \Leftrightarrow (\forall x_1 \dots \forall x_n \psi'(x_1, \dots, x_n)),$$

i.e. for every universal formula φ we have that $(\mathcal{M}, \mathcal{D}) \models_0 \varphi \Leftrightarrow \varphi'$. From this fact the statement of the theorem directly follows. \square

Before we continue with the completeness proof, we note the following independently interesting corollary of the previous two lemmas.

Corollary 5.3. *If a formula φ not containing equality, function and constant symbols is ε -satisfiable, then it is also ε -satisfiable in a model on $[0, 1]$ with the Lebesgue measure where we can choose each relation to be either Π_2^0 or Σ_2^0 .*

Proof. Let φ be ε -satisfiable; for now we assume that φ is a universal relational formula. By Theorem 2.6 it is then also ε -satisfied in a model (\mathcal{M}, λ) on $[0, 1]$ with the Lebesgue measure. For every relation R of arity n and every $m \in \omega$ use Lemma 5.1 to determine $k_m \in \omega$ and Lebesgue-measurable sets $X_{m,i,j}$ such that for all functions $f : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ we have

$$\Pr_{\mathcal{D}^n}[(a_1, \dots, a_n) \in \mathcal{M}^n \mid (a_{f(1)}, \dots, a_{f(n)}) \in (R^{\mathcal{M}} \Delta ((X_{m,1,1} \times \dots \times X_{m,1,n}) \cup \dots \cup (X_{m,k_m,1} \times \dots \times X_{m,k_m,n})))] \leq \frac{1}{2^m}.$$

For every such $X_{m,i,j}$, let $Y_{m,i,j}$ be a Σ_1^0 set such that $X_{m,i,j} \Delta Y_{m,i,j}$ has Lebesgue-measure 0 (see e.g. Kechris [5, Theorem 17.10]). Now let \mathcal{N} be the model where the relations are taken to be the lim sup of these approximations; i.e. let

$$R^{\mathcal{N}} = \bigcap_{M \in \omega} \bigcup_{m \geq M} ((Y_{m,1,1} \times \dots \times Y_{m,1,n}) \cup \dots \cup (Y_{m,k_m,1} \times \dots \times Y_{m,k_m,n})).$$

Then $\Pr_{\mathcal{D}^n}[R^{\mathcal{N}} \Delta R^{\mathcal{M}}] = 0$, see e.g. Bogachev [2, Theorem 1.12.6], so $(\mathcal{M}, \mathcal{D})$ and $(\mathcal{N}, \mathcal{D})$ satisfy the same universal formulas by Lemma 5.2. Furthermore, the relations in \mathcal{N} are clearly Π_2^0 .

In the general case, let φ be any formula and let φ' be the universal relational formula from Proposition 4.2. Then φ' is satisfiable in a model with Π_2^0 relations by our argument above, and from the proof of Proposition 4.2 we directly see that in fact then also φ is satisfiable in a model with Π_2^0 relations.

Finally, if we want our relations to be Σ_2^0 , we can take the lim inf of Π_1^0 sets $Y_{m,i,j}$ instead of the lim sup of Σ_1^0 sets $Y_{m,i,j}$ we took above. \square

We will now formalise the ideas discussed above to obtain an analogue of Theorem 3.3 for regular ε -satisfiability.

Theorem 5.4. *Let φ be a universal formula in a relational language \mathcal{L} . Then the following are equivalent:*

- (i) φ is ε -satisfiable;
- (ii) *there exists a finitely additive model $(\mathcal{M}, \mathcal{D})$ such that $(\mathcal{M}, \mathcal{D}) \models_{\varepsilon} \varphi$, and such that for every $m \in \omega$ and every relation R occurring in φ there exists $k_m \in \omega$ and interpretations $X_{m,i,j}^{\mathcal{M}} \subseteq \mathcal{M}$ for $1 \leq i \leq k_m$, $1 \leq j \leq n$ which satisfy that for every function $f : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ we have that:*

$$(4) \quad (\mathcal{M}, \mathcal{D}) \models_{\frac{1}{m}} \forall x_1 \dots \forall x_n \left(R(x_1, \dots, x_n) \leftrightarrow \bigvee_{i=1}^{k_m} \bigwedge_{j=1}^n X_{m,i,j}^{\mathcal{M}}(x_{f(j)}) \right).$$

Proof. The implication from (i) to (ii) directly follows from Lemma 5.1. For the converse, we note that (4) can be expressed in the language \mathcal{L}' which consists of \mathcal{L} enlarged with countably many unary predicate symbols $X_{m,i,j}$. Thus, if we have some finitely additive model satisfying (ii), then by Theorem 3.3 we can find a weak ε -model $(\mathcal{M}', \mathcal{D}')$ which also satisfies (ii). For every $m \in \omega$, fix $k_m \in \omega$ and interpretations $X_{m,i,j}^{\mathcal{M}'} \subseteq \mathcal{M}'$ as in (ii). For ease of notation, write

$$Z_m = \left((X_{m,1,1}^{\mathcal{M}'} \times \dots \times X_{m,1,n}^{\mathcal{M}'}) \cup \dots \cup (X_{m,k_m,1}^{\mathcal{M}'} \times \dots \times X_{m,k_m,n}^{\mathcal{M}'}) \right).$$

From (4) we directly see that

$$\Pr_{\mathcal{D}^n}[Z_m \triangle Z_{m'}] \leq \left(\frac{1}{m} + \frac{1}{m'} \right)^n.$$

In particular, we see that Z_0, Z_1, \dots is a Cauchy sequence in the pseudometric $d(X, Y) = \Pr_{\mathcal{D}^n}[X \triangle Y]$. Since this metric is complete (see e.g. Bogachev [2, Theorem 1.12.6]) we can determine a \mathcal{D}^n -measurable set Z_R to which this sequence converges.

Now, let $\varepsilon' > 0$. Fix m such that $\frac{1}{m} < \frac{\varepsilon'}{2}$ and such that $d(Z_m, Z_R) < \frac{\varepsilon'}{2}$. Then from (4) it easily follows that

$$(\mathcal{M}', \mathcal{D}) \models_{\varepsilon'} \forall x_1 \dots \forall x_n (R(x_1, \dots, x_n) \leftrightarrow Z_R(x_1, \dots, x_n)).$$

So, from the proof of Terwijn [12, Proposition 3.4] it then follows that

$$(\mathcal{M}', \mathcal{D}) \models_0 \forall x_1 \dots \forall x_n (R(x_1, \dots, x_n) \leftrightarrow Z_R(x_1, \dots, x_n)).$$

Now, if we let \mathcal{N} be the model where $R^{\mathcal{N}} = Z_R$, then $(\mathcal{M}', \mathcal{D})$ and $(\mathcal{N}, \mathcal{D})$ ε -satisfy the same universal formulas by Lemma 5.2. In particular, $(\mathcal{N}, \mathcal{D})$ is an ε -model which ε -satisfies φ . \square

Theorem 5.5. *Let \mathcal{L} be a countable language not containing equality and function symbols. Then ε -satisfiability is Σ_1^1 .*

Proof. We only need to consider universal relational formulas by Theorem 4.2. From Theorem 3.5 and Theorem 5.4 we see that a universal relational formula φ is ε -satisfiable if and only if there is a countable finitely additive model as in (ii) of Theorem 5.4. However, this last statement can be expressed as a Σ_1^1 -formula, in a similar way as explained in the proof of Theorem 3.6. \square

Remark 5.6. As in Remark 3.7, we can actually make finitely additive models as in Theorem 5.4 (ii) by using a variant of Keisler's calculus. The extra requirement (4) corresponds to adding an extra deduction rule with countable infinitely many hypotheses. We will not go into more detail here, for the reasons explained in the beginning of section 3.

6. DECIDABILITY OF 0-SATISFIABILITY

In this section we will prove that 0-satisfiability (for languages not containing equality and function symbols) is not only Σ_1^1 , but is in fact decidable. This stands in stark contrast to the fact that for $\varepsilon \in (0, 1)$ we have that ε -satisfiability is Σ_1^1 -complete, as will be shown in section 7. It also contrasts the fact that 0-validity is undecidable: the 0-tautologies are exactly the classical tautologies, as shown in Terwijn [12].

As in the previous section, we only need to consider universal relational formulas because of Theorem 4.2. First, we show that 0-satisfiability corresponds to classical satisfiability in a natural way.

Proposition 6.1. *Let $m \in \omega$. If a universal relational formula*

$$\varphi = \forall x_1 \dots \forall x_n \psi(x_1, \dots, x_n)$$

with $n \leq m$ has an ε -model $(\mathcal{M}, \mathcal{D})$, and $\frac{m!}{(m-n)!}(1 - (1 - \varepsilon)^n) < 1$, then there exists a finite, classical model \mathcal{N} of size m which satisfies

$$\tilde{\varphi} = \forall x_1 \dots \forall x_n \left(\bigwedge_{i < j} x_i \neq x_j \rightarrow \psi(x_1, \dots, x_n) \right)$$

classically. In particular, this holds if $\varepsilon = 0$.

Proof. By Theorem 2.6 we may assume \mathcal{D} to be the Lebesgue measure λ . In particular the diagonals have measure 0. Since $(\mathcal{M}, \mathcal{D}) \models_\varepsilon \varphi$ we therefore have that, for each injective function $f : \{1, \dots, n\} \hookrightarrow \{1, \dots, m\}$, the set

$$A_f := \{(a_1, \dots, a_m) \in \mathcal{M}^n \mid a_i \neq a_j \text{ for } i \neq j \text{ and } \mathcal{M} \models \psi(a_{f(1)}, \dots, a_{f(n)})\}$$

has \mathcal{D}^n -measure at least $(1 - \varepsilon)^n$. Let

$$B = \bigcap_{f: \{1, \dots, n\} \hookrightarrow \{1, \dots, m\}} A_f.$$

Then we see that the set

$$\mathcal{M} \setminus B = \bigcup_{f: \{1, \dots, n\} \hookrightarrow \{1, \dots, m\}} (\mathcal{M} \setminus A_f)$$

has measure at most $\frac{m!}{(m-n)!}(1 - (1 - \varepsilon)^n)$, which is strictly smaller than 1 by assumption. So, we see that B has positive measure. In particular B is non-empty, so choose any $(a_1, \dots, a_m) \in B$. Then \mathcal{M} restricted to $\{a_1, \dots, a_m\}$ classically satisfies $\tilde{\varphi}$: indeed, if $b_1, \dots, b_n \in \{a_1, \dots, a_m\}$ are distinct elements, and we let f be the function sending each $1 \leq i \leq n$ to the unique $1 \leq j \leq m$ such that $b_i = a_j$, then $(a_1, \dots, a_m) \in A_f$ implies that $\mathcal{M} \models \psi(b_1, \dots, b_n)$. \square

We would also like to have a converse to this proposition, i.e. we would like to know if there exists an $m \in \omega$ such that if $\forall x_1 \dots \forall x_n (\bigwedge_{i < j} x_i \neq x_j \rightarrow \psi(x_1, \dots, x_n))$ has a finite classical model of size m , then φ also has a 0-model. It turns out that we can do this if we choose m big enough, by using Ramsey's theorem. In fact, the way we will use Ramsey's theorem is very similar to the original use of this theorem by Ramsey in [10], where Ramsey proved his combinatorial theorem in order to prove that the variant of the Entscheidungsproblem asking if a universal relational formula has an infinite model is decidable. First, we need a definition.

Definition 6.2. Let \mathcal{M} be a first-order model, let $X \subseteq \mathcal{M}$ and let $<$ be a linear ordering of X . Then we call $(X, <)$ a *sequence of indiscernibles* if for every $n \in \omega$ and all sequences $a_1 < a_2 < \dots < a_n$, $b_1 < b_2 < \dots < b_n$ both in X we have for every formula $\varphi(x_1, \dots, x_n)$ that $\mathcal{M} \models \varphi(a_1, \dots, a_n)$ if and only if $\mathcal{M} \models \varphi(b_1, \dots, b_n)$.

Theorem 6.3. (Ramsey [10, page 279]) *Let $\varphi = \forall x_1 \dots \forall x_n \psi(x_1, \dots, x_n)$ be a universal formula not containing function and constant symbols. Let \mathcal{M} be a classical model of size n which satisfies φ and for which there exists a linear order $<$ on \mathcal{M} which turns $(\mathcal{M}, <)$ into a sequence of indiscernibles. Let (X, \prec) be any linearly ordered set. Then there exists a model \mathcal{N} for φ on X which has (X, \prec) as a sequence of indiscernibles.⁴*

Corollary 6.4. *If for a propositional relational formula ψ the formula*

$$\tilde{\varphi} = \forall x_1 \dots \forall x_n \left(\bigwedge_{i < j} x_i \neq x_j \rightarrow \psi(x_1, \dots, x_n) \right)$$

has a model \mathcal{M} of size n for which there exists a linear order $<$ on \mathcal{M} which turns $(\mathcal{M}, <)$ into a sequence of indiscernibles, then

$$\varphi = \forall x_1 \dots \forall x_n \psi(x_1, \dots, x_n)$$

has a 0-model.

⁴This is not entirely the way in which Ramsey formulated his theorem, but in fact his proof directly yields us the result we stated here.

Proof. By the previous theorem, there exists a model on $[0, 1]$ satisfying $\tilde{\varphi}$ which has $([0, 1], <)$ (where $<$ is the usual ordering on $[0, 1]$) as a sequence of indiscernibles. Because the diagonal has Lebesgue-measure 0, we then directly see that also $(\mathcal{X}, \lambda) \models_0 \varphi$. The only thing we still need to verify is that (\mathcal{X}, λ) is a 0-model. By Proposition 2.7 it is enough to prove that the relations $R^{\mathcal{X}}$ are Borel. Because $([0, 1], <)$ is a sequence of indiscernibles we have for every relation R of arity k :

$$R^{\mathcal{X}} = \bigcup \left\{ \{(a_1, \dots, a_k) \in [0, 1]^k \mid \forall 1 \leq i, j \leq k (a_i < a_j \leftrightarrow f(i) < f(j))\} \right. \\ \left. \mid f : \{1, \dots, k\} \rightarrow \{1, \dots, k\} \text{ and } R^{\mathcal{M}}(b_{f(1)}, \dots, b_{f(k)}) \right\}.$$

Then $R^{\mathcal{X}}$ is equal to a finite union of sets of the form

$$\{(a_1, \dots, a_k) \in [0, 1]^k \mid \forall 1 \leq i, j \leq k (a_i < a_j \leftrightarrow f(i) < f(j))\}$$

which are Borel because the ordering $<$ on $[0, 1]$ is Borel. So, $R^{\mathcal{X}}$ is Borel, as desired. \square

The following result follows from the finite Ramsey theorem, together with the fact that there are computable upper bounds for the Ramsey numbers.

Theorem 6.5. (Ramsey [10, page 279]) *There exists a computable $f : \omega^3 \rightarrow \omega$ such that for every universal relational formula $\varphi = \forall x_1 \dots \forall x_n \psi(x_1, \dots, x_n)$ containing k relation symbols of arity at most m , if φ has a classical model of size at least $f(n, k, m)$, then it also has a model containing a sequence of indiscernibles of size n .*

Putting all things together, we obtain:

Theorem 6.6. *There exists a computable function $f : \text{Form} \rightarrow \omega$ such that for every universal relational formula $\varphi = \forall x_1 \dots \forall x_n \psi(x_1, \dots, x_n)$ the following are equivalent:*

- (i) φ is 0-satisfiable;
- (ii) φ is ε -satisfiable for some $\varepsilon \in [0, 1]$ satisfying $\frac{f(\varphi)!}{(f(\varphi)-n)!} (1 - (1 - \varepsilon)^k) < 1$;
- (iii) $\tilde{\varphi} = \forall x_1 \dots \forall x_n \bigwedge_{i < j} (x_i \neq x_j \rightarrow \psi(x_1, \dots, x_n))$ has a classical model of size $f(\varphi)$;
- (iv) $\tilde{\varphi}$ has a classical model containing a sequence of indiscernibles of size n .

Proof. Let \tilde{f} be the computable function from Proposition 6.5; this function directly induces a computable function $f : \text{Form} \rightarrow \omega$. We prove the equivalences.

(i) \rightarrow (ii): This is directly clear.

(ii) \rightarrow (iii): This follows from Proposition 6.1.

(iii) \rightarrow (iv): This was shown in Theorem 6.5.

(iv) \rightarrow (i): Finally, this was shown in Corollary 6.4. \square

In particular, we see:

Theorem 6.7. *0-satisfiability is decidable for languages not containing equality and function symbols.*

Proof. By Theorem 4.2, we only need to consider universal relational formulas φ . Let f be the computable function as in Theorem 6.6. If we want to check if φ is 0-satisfiable, then by Theorem 6.6 (iii) we only need to check if

$$\tilde{\varphi} = \forall x_1 \dots \forall x_n \bigwedge_{i < j} (x_i \neq x_j \rightarrow \psi(x_1, \dots, x_n))$$

has a classical model of size $f(\varphi)$, which is a decidable property. \square

There is another interesting fact which follows from Theorem 6.6: the equivalence of (ii) and (i) says that a formula φ which is ε -satisfiable for small enough $\varepsilon > 0$ is in fact also 0-satisfiable. Thus, one can see 0-satisfiability as the limit of ε -satisfiability for $\varepsilon > 0$, in the sense given below.

Corollary 6.8. $\bigcap_{\varepsilon > 0} \varepsilon\text{-SAT} = 0\text{-SAT}$.

Proof. From Theorem 6.6. □

Note that we cannot have in general that $\bigcap_{\varepsilon' > \varepsilon} \varepsilon'\text{-SAT} = \varepsilon\text{-SAT}$. Namely, in that case, one could prove using Lemma 5.1 that a universal relational sentence φ is ε -satisfiable if and only if it has a finite ε' -model for every rational $\varepsilon' > \varepsilon$. However, the latter can be expressed as a first-order arithmetical sentence, while the universal relational fragment of ε -satisfiability will be shown to be Σ_1^1 -hard in Theorem 7.13.

7. ε -SATISFIABILITY IS Σ_1^1 -HARD

In this section we will show that ε -satisfiability is Σ_1^1 -hard for rational $\varepsilon \in (0, 1)$. Together with the result from section 5 this will show that ε -satisfiability is Σ_1^1 -complete. We will prove this hardness step by step, interpreting more and more of arithmetic within ε -logic as we go. As a first step, we will look at sentences of the form $\exists Q\varphi(Q)$ in the language of arithmetic (that is, the language consisting of $S, +, \cdot, 0$ and $=$) where φ is universal, in the sense the only quantifiers occurring in φ are first-order universal quantifiers. Equivalently, we can see $\varphi(Q)$ as a first-order universal sentence in φ in the language of arithmetic enlarged with a unary predicate Q , and $\exists Q\varphi(Q)$ is satisfiable in second-order arithmetic if and only if φ is satisfiable in first-order arithmetic (i.e. there exists some interpretation $Q^{\mathbb{N}}$ such that $\mathbb{N} \models \varphi$ under this interpretation for Q). We will implicitly use this equivalence throughout this section.

Furthermore, to optimise our result and show that we do not need function or constant symbols in our language to prove hardness (i.e. to show that the relational fragment is already Σ_1^1 -hard), we will not look at $S, +, \cdot$ and 0 as functions or constants, but instead as relations $S(x) = y, x + y = z, x \cdot y = z$ and $0 = x$. It is easy to see that any universal sentence φ in the language with functions $S, +, \cdot$ and 0 can be transformed into a universal sentence in the language with relations for $S, +, \cdot$ and 0 . Henceforth, when we talk about the language of arithmetic, we will mean the language with relation symbols for $S, +, \cdot, 0$ and $=$.

Definition 7.1. The *language of arithmetic* is the language consisting of relation symbols $S(x) = y, x + y = z, x \cdot y = z, 0 = x$ and $x = y$. For a formula φ in the language of arithmetic, we say that $\mathbb{N} \models \varphi$ if φ holds in the natural numbers ω together with the usual interpretations for $S, +, \cdot, 0$ and $=$.

To formulate the first step in our interpretation, we need a few more tools.

Definition 7.2. Let φ be a formula in prenex normal form and N a unary predicate. Then φ^N , or φ *relativised to N* , is defined as the formula where each $\forall x\psi(x)$ is replaced by $\forall x(N(x) \rightarrow \psi(x))$ and each $\exists x\psi(x)$ is replaced by $\exists x(N(x) \wedge \psi(x))$.

Definition 7.3. Let $\psi(x_1, \dots, x_n)$ be a propositional formula. Then we denote by $\psi^\#(x_1, \dots, x_n)$ the formula

$$\bigwedge_{f: \{1, \dots, n\} \rightarrow \{1, \dots, n\}} \psi(x_{f(1)}, \dots, x_{f(n)}).$$

If $\varphi = \forall x_1 \dots \forall x_n \psi(x_1, \dots, x_n)$ with ψ a propositional formula, then we denote by $\varphi^\#$ the formula $\forall x_1 \dots \forall x_n \psi^\#(x_1, \dots, x_n)$.

Note that for every universal sentence φ , if $\mathcal{M} \models \varphi$ holds classically, then also $\mathcal{M} \models \varphi^\#$. It is the following converse which we will need.

Lemma 7.4. *For every propositional formula $\psi(x_1, \dots, x_n)$ and every countable model \mathcal{M} : if there exists an enumeration $\{a_0, a_1, \dots\}$ of \mathcal{M} such that for all injective functions $\pi : \{1, \dots, n\} \hookrightarrow \omega$ we have $\mathcal{M} \models \psi^\#(a_{\pi(1)}, \dots, a_{\pi(n)})$, then $\mathcal{M} \models \forall x_1 \dots \forall x_n \psi(x_1, \dots, x_n)$.*

Proof. To derive a contradiction, assume $\mathcal{M} \not\models \forall x_1 \dots \forall x_n \psi(x_1, \dots, x_n)$. Let $a_{i_1}, \dots, a_{i_n} \in \mathcal{M}$ (not necessarily distinct) be such that $\mathcal{M} \not\models \psi(a_{i_1}, \dots, a_{i_n})$. Fix a subset $A \subseteq \omega$ of size n such that $\{i_1, \dots, i_n\} \subseteq A$ and fix a bijection $\pi : \{1, \dots, n\} \rightarrow A$. Finally, let $f : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ be the function mapping $1 \leq j \leq n$ to $\pi^{-1}(i_j)$. Then $i_j = \pi(f(j))$ so $\mathcal{M} \not\models \psi(a_{\pi(f(1))}, \dots, a_{\pi(f(n))})$. Therefore $\mathcal{M} \not\models \psi^\#(a_{\pi(1)}, \dots, a_{\pi(n)})$, a contradiction. \square

Proposition 7.5. *Let \mathcal{L} be a relational language. Let $\varphi = \forall x_1 \dots \forall x_n \varphi(x_1, \dots, x_n)$ be a universal formula in the language \mathcal{L} and let $(\mathcal{M}, \mathcal{D}) \models_0 \varphi^\#$. Then there exists a countable submodel $\mathcal{N} \subseteq \mathcal{M}$ such that $\mathcal{N} \models \varphi$ classically. In fact, for \mathcal{D}^ω -almost all $(a_0, a_1, \dots) \in \mathcal{M}^\omega$ we have that $\mathcal{M} \upharpoonright \{a_0, a_1, \dots\} \models \varphi$ classically.*

Proof. For each injective function $\pi : \{1, \dots, n\} \hookrightarrow \omega$, the set

$$\{(a_0, a_1, \dots) \in \mathcal{M}^\omega \mid (\mathcal{M}, \mathcal{D}) \models_0 \psi^\#(a_{\pi(1)}, \dots, a_{\pi(n)})\}$$

is measurable (by Proposition 2.3). Furthermore, one can easily see that it in fact has measure 1, by using Fubini's theorem and the fact that $(\mathcal{M}, \mathcal{D}) \models_0 \varphi^\#$; for this it is essential that π is injective.

Therefore the set

$$B = \bigcap_{\pi: \{1, \dots, n\} \hookrightarrow \omega} \{(a_0, a_1, \dots) \in \mathcal{M}^\omega \mid (\mathcal{M}, \mathcal{D}) \models_0 \psi^\#(a_{\pi(1)}, \dots, a_{\pi(n)})\}$$

is a countable intersection of sets of measure 1, and therefore has measure 1 itself. From Lemma 7.4 we know that for every $(a_0, a_1, \dots) \in B$ we have that φ holds classically in $\mathcal{M} \upharpoonright \{a_0, a_1, \dots\}$, which completes our proof. \square

Theorem 7.6. *Let \mathcal{L} be the language consisting of unary relation symbols $0 = x$ and $N(x)$ (which will represent our set of natural numbers), binary relation symbols $S(x) = y$, $x = y$,⁵ $x < y$, $x \prec y$ (which will represent $x < S(y)$) and $R(x, y)$,⁶ ternary relation symbols $x + y = z$ and $x \cdot y = z$, and a unary relation symbol Q . Furthermore, let f_0 be the reduction from 0-satisfiability to $\frac{1}{2}$ -satisfiability from Theorem 2.12, and let $f_{\frac{1}{4}}$ and $f_{\frac{3}{4}}$ be similar reductions for $\frac{1}{4}$ - and $\frac{3}{4}$ -satisfiability. Then there exists finite theories $T_0, T_{\frac{1}{4}}, T_{\frac{1}{2}}$ and $T_{\frac{3}{4}}$ in the language \mathcal{L} , containing only universal sentences, such that for every first-order universal sentence φ in the language of arithmetic enlarged with Q ,⁷ the following are equivalent:*

- (i) $f_0((\bigwedge T_0 \wedge \varphi^N)^\#) \wedge f_{\frac{1}{4}}(\bigwedge T_{\frac{1}{4}}) \wedge f_{\frac{3}{4}}(\bigwedge T_{\frac{3}{4}}) \wedge \bigwedge T_{\frac{1}{2}}$ is $\frac{1}{2}$ -satisfiable;⁸
- (ii) $\mathbb{N} \models \exists Q \varphi(Q)$.

⁵Here we do not mean true equality, but rather a binary relation that we will use to represent equality.

⁶The intended interpretation of $R(x, y)$ is quite technical. It is best to think of our model \mathcal{M} as consisting of two copies of ω : one living inside N , one living outside N . However, all the operations will only be defined on the elements of N ; there is no extra structure on $\mathcal{M} \setminus N^{\mathcal{M}}$. Our relation R will then be a subset of $N^{\mathcal{M}} \times (\mathcal{M} \setminus N^{\mathcal{M}})$; the intended interpretation is then that $R(a, b)$ holds if $a \in N^{\mathcal{M}}$, $b \in \mathcal{M} \setminus N^{\mathcal{M}}$ and $b \neq S(a)$.

⁷Remember, as defined above this is the language consisting of (relation symbols for) $S, +, \cdot, 0, =$ and Q .

⁸Strictly speaking, $(\bigwedge T_0 \wedge \varphi^N)^\#$ is undefined because it is not in prenex normal form. To avoid the problem that prenex normal forms are not unique, we assume it has been transformed

Proof. For T_0 we take the universal axioms of Robinson's Q relativised to N , axioms specifying that our relations only hold on N , axioms for $<$ and \prec , and some special axioms for N and R . Some of the axioms will turn out to be redundant, but we have added them anyway so that all axioms of Robinson's Q are in T_0 . The reason we only add the universal axioms and we avoid the axioms involving existential quantifiers is that, while we can say that something exists (using the \exists quantifier), we cannot say that something holds with strictly positive measure. Our relation R is meant as a trick to work around this problem.

Thus, we will add the following axioms to T_0 .

All equality axioms. For example:

$$\begin{aligned} &\forall x(x = x) \\ &\forall x\forall y((N(x) \wedge x = y) \rightarrow N(y)) \end{aligned}$$

We should guarantee that 0 is in N :

$$\forall x(0 = x \rightarrow N(x))$$

We now give the axioms for the successor function:

$$\begin{aligned} &\forall x\forall y(S(x) = y \rightarrow N(x) \wedge N(y)) \\ &(\forall x\forall y\neg(S(x) = y \wedge 0 = y))^N \\ &(\forall x\forall y\forall u\forall v((S(x) = u \wedge S(y) = v) \rightarrow (x = y \leftrightarrow u = v)))^N \end{aligned}$$

We proceed with the inductive definitions of $+$ and \cdot :

$$\begin{aligned} &(\forall x\forall y\forall z(x + y = z \rightarrow (N(x) \wedge N(y) \wedge N(z)))) \\ &(\forall x\forall y\forall z(0 = y \rightarrow (x + y = z \leftrightarrow x = z)))^N \\ &(\forall x\forall y\forall u\forall v\forall w((S(y) = u \wedge x + y = w) \rightarrow (x + u = v \leftrightarrow S(w) = v)))^N \\ &(\forall x\forall y\forall z(x \cdot y = z \rightarrow (N(x) \wedge N(y) \wedge N(z))))^N \\ &(\forall x\forall y\forall z(0 = y \rightarrow (x \cdot y = z \leftrightarrow 0 = z)))^N \\ &(\forall x\forall y\forall u\forall v\forall w((S(y) = u \wedge x \cdot y = w) \rightarrow (x \cdot u = v \leftrightarrow w + x = v)))^N. \end{aligned}$$

into prenex normal form using some fixed algorithm (for example, the algorithm arising from the proof of Proposition 2.8).

Normally, we would be able to use an existential quantifier and the other relations to define $x < y$ and $x \prec y$ (the latter standing for $S(x) < y$). However, because we are trying to avoid the existential quantifier in order to make the proof work, we give axioms for $<$ and \prec which we need in our proof. It is easy to see that these axioms are sound. Because our formula φ does not contain $<$ or \prec , it is irrelevant if they completely define $<$ and \prec or not, so we only add those axioms which we need in our proof.

$$\begin{aligned} & \forall x \forall y (x < y \rightarrow N(x) \wedge N(y)) \\ & \forall x \forall y (x \prec y \rightarrow N(x) \wedge N(y)) \\ & (\forall x \forall y \neg (S(x) = y \wedge x \prec y))^N \\ & (\forall x \forall y (x < y \leftrightarrow (x \prec y \vee S(x) = y)))^N \\ & (\forall x \forall y \neg (0 = y \wedge x < y))^N \\ & (\forall x \forall y (0 = x \vee (0 = y \rightarrow y < x)))^N \\ & (\forall x \forall y \forall z (S(x) = z \rightarrow (x \prec y \leftrightarrow z < y)))^N \end{aligned}$$

Finally, we introduce a predicate R , with axioms which use $<$ and \prec . This predicate is meant to function as a sort of ‘padding’. The goal of this predicate is to force the measure of $\{x \mid x = S(y)\}$ to be exactly half of $\{x \mid x > y\}$. Since we will also add an axiom saying that the set of points equal to 0 has measure $\frac{1}{4}$, this means the set of points equal to $S^n(0)$ will have measure 2^{-n-2} . We will use this to show that the collection of ‘standard submodels’ of \mathcal{M} has measure 1. The intended interpretation was described above. First, we want the second coordinate of R to only hold outside of N .

$$\forall x \forall y (R(x, y) \rightarrow \neg N(y))$$

The next axioms for R will be in $T_{\frac{1}{2}}$ instead of in T_0 , because these need to be evaluated for $\varepsilon = \frac{1}{2}$ while the rest will be evaluated for $\varepsilon = 0$. So, this means in these axioms the universal quantifier will mean “for measure at least $\frac{1}{2}$ many” instead of the interpretation “for almost all” in T_0 .

$$\begin{aligned} & \forall x (N(x) \wedge \forall y (R(x, y) \vee S(x) = y)) \\ & \forall x (N(x) \wedge \forall y \neg (R(x, y) \vee S(x) = y)) \\ & \forall x (N(x) \wedge \forall y (R(x, y) \vee x \prec y)) \\ & \forall x (N(x) \wedge \forall y \neg (R(x, y) \vee x \prec y)). \end{aligned}$$

We also add axioms to $T_{\frac{1}{2}}$ to make sure that N has measure $\frac{1}{2}$:

$$\begin{aligned} & \forall x N(x) \\ & \forall x \neg N(x). \end{aligned}$$

Finally, we add the next two axioms to $T_{\frac{1}{4}}$ respectively $T_{\frac{3}{4}}$ to ensure that the points equal to 0 together have measure $\frac{1}{4}$:

$$\begin{aligned} & (\forall x (0 = x))^N \\ & (\forall x \neg (0 = x))^N \end{aligned}$$

We will now show that these axioms indeed do what we promised.

(ii) \rightarrow (i): Assume $\mathbb{N} \models \exists Q \varphi(Q)$. Let $Q^{\mathbb{N}}$ witness this fact. Now take the model $\mathcal{M} := \omega \times \{0, 1\}$ to be the disjoint union of two copies of ω , where we define $S, +, \cdot, \leq, \prec, 0$ on the first copy $\omega \times \{0\}$ of ω as in \mathbb{N} (remembering that $x \prec y$ should

mean $S(x) < y$ and let them be false elsewhere. For example, $(a, 0) + (b, 0) = (c, 0)$ holds if and only if $a + b = c$, and it is false in all other cases. Let $Q^{\mathcal{M}}((a, 0))$ hold if $Q^{\mathbb{N}}(a)$ holds and let it be false elsewhere. Let

$$N = \omega \times \{0\} \text{ and } R = \{(a, 0), (b, 1) \mid b \neq S(a)\}.$$

Finally, define \mathcal{D} by

$$\mathcal{D}(a, 0) = \mathcal{D}(a, 1) := \frac{1}{2^{a+2}}.$$

Then it is directly verified that

$$(\mathcal{M}, \mathcal{D}) \models_0 \left(\bigwedge T_0 \wedge \varphi^N \right)^{\#},$$

because all formulas in T_0 even hold classically in \mathcal{M} .

Furthermore, N clearly has measure $\frac{1}{2}$ and the points equal to $0 = (a_0)$ have measure $\frac{1}{4}$, so the axioms expressing this fact hold; in particular $(\mathcal{M}, \mathcal{D}) \models_{\frac{1}{4}} \bigwedge T_{\frac{1}{4}}$ and $(\mathcal{M}, \mathcal{D}) \models_{\frac{3}{4}} \bigwedge T_{\frac{3}{4}}$. Next, if we let $a \in \omega$ then we have that

$$\begin{aligned} & \Pr_{\mathcal{D}}[y \in \mathcal{M} \mid (\mathcal{M}, \mathcal{D}) \models_{\frac{1}{2}} R((a, 0), y) \vee S(a, 0) = y] \\ &= \frac{1}{2} - \frac{1}{2^{a+3}} + \frac{1}{2^{a+3}} \\ &= \frac{1}{2} \end{aligned}$$

while we also have that

$$\begin{aligned} & \Pr_{\mathcal{D}}[y \in \mathcal{M} \mid (\mathcal{M}, \mathcal{D}) \models_{\frac{1}{2}} R((a, 0), y) \vee (a, 0) < y] \\ &= \frac{1}{2} - \frac{1}{2^{a+3}} + \sum_{i=a+4}^{\infty} 2^{-i} \\ &= \frac{1}{2}. \end{aligned}$$

Thus, we also see that the axioms expressing that these two sets have measure $\frac{1}{2}$ hold, i.e. that $(\mathcal{M}, \mathcal{D}) \models_{\frac{1}{2}} \bigwedge T_{\frac{1}{2}}$. But then we see from Remark 2.14 that there is a $\frac{1}{2}$ -model $(\mathcal{N}, \mathcal{E})$ such that

$$(\mathcal{N}, \mathcal{E}) \models_{\frac{1}{2}} f_0 \left(\left(\bigwedge T_0 \wedge \varphi^N \right)^{\#} \right) \wedge f_{\frac{1}{4}} \left(\bigwedge T_{\frac{1}{4}} \right) \wedge f_{\frac{3}{4}} \left(\bigwedge T_{\frac{3}{4}} \right) \wedge \bigwedge T_{\frac{1}{2}}$$

i.e. that (i) holds.

(i) \rightarrow (ii): Assume (i) holds. Then, by Remark 2.14 there exists a probability model $(\mathcal{M}, \mathcal{D})$ such that $(\mathcal{M}, \mathcal{D}) \models_{\varepsilon} T_{\varepsilon}$ for $\varepsilon \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\}$ and $(\mathcal{M}, \mathcal{D}) \models_0 \left(\bigwedge T_0 \wedge \varphi^N \right)^{\#}$.

By Proposition 7.5 we have that the set

$$B = \{(a_0, a_1, \dots) \in \mathcal{M}^{\omega} \mid \mathcal{M} \upharpoonright \{a_0, a_1, \dots\} \models \bigwedge T_0 \wedge \varphi^N\}$$

has \mathcal{D}^{ω} -measure 1. In particular, because all universal Robinson axioms are in T_0 we see that for every $(a_0, a_1, \dots) \in B$ these universal Robinson axioms hold classically in $\mathcal{M} \upharpoonright \left(\{a_0, a_1, \dots\} \cap N^{\mathcal{M}} \right)$. The only problem is that we do not know if the two missing non-universal axioms $\forall x(0 = x \vee \exists y S(y) = x)$ and $\forall x \exists y S(x) = y$ also hold. However, if we were able to find a subsequence (b_0, b_1, \dots) of a permutation of some $(a_0, a_1, \dots) \in B$ satisfying $0 = b_0$ and $S(a_i) = a_{i+1}$ for all $i \in \omega$, these two axioms would also hold in $\mathcal{M} \upharpoonright \{b_0, b_1, \dots\}$. Furthermore, note that T_0 guarantees that $\{b_0, b_1, \dots\} \subseteq N^{\mathcal{M}}$, so we then have that $\mathcal{M} \upharpoonright \{b_0, b_1, \dots\}$ models *all* the Robinson axioms. Even stronger: we have that this model is a model isomorphic to \mathbb{N} in which

φ holds (note that φ is universal, so it holds in any submodel of $\mathcal{M} \upharpoonright \{a_0, a_1, \dots\}$), showing that (ii) holds. The rest of the proof will therefore consist of showing that we can find such a sequence in B .

First, we restrict to a subset B'' of B which still has measure 1. For every sentence $\alpha \in T_0$, say $\forall x_1 \dots \forall x_t \beta(x_1, \dots, x_t)$ (so $t \leq 3$), let $\tilde{\alpha}(y_1, y_2)$ be the formula $\forall x_t \beta(y_1, \dots, y_{t-1}, x_t)$. Then

$$B' := B \cap \bigcap_{\pi: \{1,2\} \leftrightarrow \omega} \{(a_0, a_1, \dots) \in \mathcal{M}^\omega \mid (\mathcal{M}, \mathcal{D}) \models_0 \bigwedge_{\alpha \in T_0} \tilde{\alpha}(a_{\pi(1)}, a_{\pi(2)})\}$$

is still a set of measure 1: it is an intersection of measurable sets by Proposition 2.3 and all of these sets have measure 1 because $(\mathcal{M}, \mathcal{D}) \models_0 (\bigwedge T_0 \wedge \varphi^N)^\#$. We make one more restriction. Let X be the set

$$\{x \in \mathcal{M} \mid \mathcal{M} \models \neg N(x)\} \cup \left\{x \in N^{\mathcal{M}} \mid \Pr_{\mathcal{D}}[y \in \mathcal{M} \mid R(x, y) \vee S(x) = y] = \Pr_{\mathcal{D}}[y \in \mathcal{M} \mid R(x, y) \vee x \prec y] = \frac{1}{2}\right\}.$$

Then X has measure 1: the union above is disjoint, $\neg N$ has measure exactly $\frac{1}{2}$ and by the axioms for R it follows that the second set has measure $\frac{1}{2}$. Now let $B'' = B' \cap X^\omega$. Then our axioms hold inside B'' in a strong sense (that is, the axioms of T_0 hold classically and the axioms about R hold classically in the first coordinate), which is a fact we will use shortly.

For each $n \in \omega$, let D_n be the set

$$\bigcup_{(m_0, \dots, m_n) \in \omega^{n+1}} \{(a_0, a_1, \dots) \in B'' \mid 0 = a_{m_0}, S(a_{m_0}) = a_{m_1}, \dots, S(a_{m_{n-1}}) = a_{m_n}\}.$$

Then $D_0 \supseteq D_1 \supseteq \dots$. By Lemma 7.7 below, each D_n has measure 1. Therefore, $\bigcap_{n \in \omega} D_n$ has measure 1 and in particular it is non-empty. Let $(a_0, a_1, \dots) \in \bigcap_{n \in \omega} D_n$. For each $n \in \omega$, fix an α_n such that

$$\exists i_0, \dots, i_{n-1} (0 = a_0 \wedge \dots \wedge S(a_{i_{n-2}}) = a_{i_{n-1}} \wedge S(a_{i_{n-1}}) = a_{\alpha_n}),$$

which exists because $(a_0, a_1, \dots) \in D_n$. Now let (b_0, b_1, \dots) be the sequence $(a_{\alpha_0}, a_{\alpha_1}, \dots)$. Then $0 = b_0$. Also, we claim that $S(b_0) = b_1$. We know that there exists some i_0 such that $a_{i_0} = 0$ and $S(a_{i_0}) = b_1$. Now, because $\mathcal{M} \upharpoonright \{a_0, a_1, \dots\} \models \bigwedge T_0$ classically, we know in particular that the equality axioms hold classically in this model. Thus, $a_{i_0} = b_0$ and $S(b_0) = b_1$, as desired. In the same way, we can show that $S(b_i) = b_{i+1}$ holds for any $i \in \omega$. Therefore, as discussed above the model $\mathcal{M} \upharpoonright \{b_0, b_1, \dots\}$ is isomorphic to \mathbb{N} and $\varphi(Q)$ holds in it, which shows that (ii) holds. \square

Lemma 7.7. *Let D_n be as in the proof of Theorem 7.6. Then $\Pr_{\mathcal{D}}[D_n] = 1$.*

Proof. Fix $n \in \omega$. For $k, i \in \omega$ with $0 \leq i < k$, denote by $D_n^{i,k}$ the set

$$\{(a_0, a_1, \dots) \in B'' \mid 0 = a_i \wedge S(a_i) = a_{i+k} \wedge \dots \wedge S(a_{i+(n-1)k}) = a_{i+nk}\}.$$

Then $D_n^{i,k} \subseteq D_n$. From $(\mathcal{M}, \mathcal{D}) \models_{\frac{1}{4}} T_{\frac{1}{4}}$ and $(\mathcal{M}, \mathcal{D}) \models_{\frac{3}{4}} T_{\frac{3}{4}}$ we see that $\Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models 0 = a] = \frac{1}{4}$, so together with Lemma 7.8 below this gives us

$$\Pr_{\mathcal{D}}[D_n^{i,k}] = \frac{1}{4} \frac{1}{8} \dots \frac{1}{2^{n+2}} = \frac{1}{2^{c_n}}$$

where $c_n = \frac{(n+2)(n+3)}{2} - 1 > 0$. Clearly, for $i \neq j$ the sets $D_n^{i,k}$ and $D_n^{j,k}$ are independent. So, we have

$$\Pr_{\mathcal{D}}\left[\bigcup_{i \leq k} D_n^{i,k}\right] = 1 - \left(1 - \frac{1}{2^{c_n}}\right)^k.$$

Combining this with the fact that $\bigcup_{i \leq k} D_n^{i,k} \subseteq D_n$ for every $k \in \omega$ we thus have

$$\Pr_{\mathcal{D}}[D_n] \geq \lim_{k \rightarrow \infty} 1 - \left(1 - \frac{1}{2^{c_n}}\right)^k = 1. \quad \square$$

Lemma 7.8. *Let B'' be as in the proof of Theorem 7.6. Let $(a_0, \dots, a_n) \in \mathcal{M}^n$ be any subsequence of a sequence in B'' such that $0 = a_0$ and $S(a_i) = a_{i+1}$. Then*

$$\Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models S(a_n) = a] = \frac{1}{2^{n+3}}.$$

Proof. Using induction over n . First, let $n = 0$. Then by the definition of B'' we have that

$$\begin{aligned} \Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models S(a_0) = a] &= \frac{1}{2} - \Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models R(a_0, a)] \\ &= \Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models a_0 \prec a] \end{aligned}$$

(where we use that R only holds outside N while S and \prec only hold inside N , so R is disjoint from $<$ and \prec). We then have that

$$\begin{aligned} 2 \Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models S(a_0) = a] &= \Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models S(a_0) = a] + \Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models a_0 \prec a] \\ &= \Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models a_0 < a]. \end{aligned}$$

(For this, we use the definition of B' together with the axiom $(\forall x \forall y (x < y \leftrightarrow (x \prec y \vee S(x) = y)))^N$.) Now, because $0 = a$ this is equal to

$$\Pr_{\mathcal{D}}[N] - \Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models 0 = a] = \frac{1}{2} - \frac{1}{4} = \frac{1}{4}.$$

So, $\Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models S(a_0) = a]$ is one half of that, i.e. $\frac{1}{8}$.

Next, assume

$$\Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models S(a_n) = a] = \frac{1}{2^{n+3}}.$$

In the same way as above, we find that

$$\Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models S(a_{n+1}) = a] = \Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models a_{n+1} \prec a].$$

Again, the sum of these two probabilities is

$$\Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models a_{n+1} < a].$$

This is equal to

$$\Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models a_n \prec a],$$

so by the induction hypothesis we now see that

$$\Pr_{\mathcal{D}}[a \in \mathcal{M} \mid \mathcal{M} \models S(a_{n+1}) = a] = \frac{1}{2} \frac{1}{2^{n+3}} = \frac{1}{2^{n+4}}. \quad \square$$

Next, we turn to the language of arithmetic enlarged with relation symbols representing finitely many primitive recursive functions — say g_1, \dots, g_n . So, for each of these primitive recursive function $g_i(x_1, \dots, x_{m_i})$ we add a relation symbol $g_i(x_1, \dots, x_{m_i}) = y$ to our language which represents this function. We next show that Theorem 7.6 still holds if we add these finitely many primitive recursive functions to our language.

Theorem 7.9. *Let g_1, \dots, g_n be primitive recursive functions. Let $f_0, f_{\frac{1}{4}}, f_{\frac{1}{2}}$ and $f_{\frac{3}{4}}$ be as in Theorem 7.6. Then there exists finite theories $T_0, T_{\frac{1}{4}}, T_{\frac{1}{2}}$ and $T_{\frac{3}{4}}$, containing only universal relational sentences, such that for each universal first-order sentence φ in the language of arithmetic enlarged with relation symbols Q, g_1, \dots, g_n , the following are equivalent:*

- (i) $f_0((\bigwedge T_0 \wedge \varphi^N)^\#) \wedge f_{\frac{1}{4}}(\bigwedge T_{\frac{1}{4}}) \wedge f_{\frac{3}{4}}(\bigwedge T_{\frac{3}{4}}) \wedge \bigwedge T_{\frac{1}{2}}$ is $\frac{1}{2}$ -satisfiable;
- (ii) $\mathbb{N} \models \exists Q\varphi(Q)$.

Proof. The proof is very similar to that of Theorem 7.6. We let $T_{\frac{1}{4}}, T_{\frac{1}{2}}$ and $T_{\frac{3}{4}}$ be as in the proof of that theorem and we extend T_0 with axioms for every g_i .

So, let $1 \leq i \leq n$. Fix any sequence h_0, \dots, h_k of primitive recursive functions in such that for each $1 \leq j \leq k$, either:

- $h_j = 0$, $h_j = S$, or h_j is a projection.
- h_j is the composition of h_s with h_{t_1}, \dots, h_{t_m} for some $1 \leq s, t_1, \dots, t_m < j$.
- h_j is defined by primitive recursion from h_s and h_t for some $1 \leq s, t < j$.
- $h_k = g_i$.

We add axioms to define each such h_j . It is clear what we need to do in the first case. If $h_j(x_1, \dots, x_m)$ is the composition of h_s with h_{t_1}, \dots, h_{t_m} , we take the relativisation to N of the universal closure of

$$\begin{aligned} & (h_{t_1}(x_1, \dots, x_n) = z_1 \wedge \dots \wedge h_{t_m}(x_1, \dots, x_n) = z_m) \\ & \rightarrow (h_j(x_1, \dots, x_m) = y \leftrightarrow h_s(z_1, \dots, z_m) = y). \end{aligned}$$

Similarly, if h_j is defined by primitive recursion over h_s and h_t , then we take the relativisation to N of the universal closure of

$$0 = x \rightarrow (h_j(x, y_1, \dots, y_n) = z \leftrightarrow h_s(y_1, \dots, y_n) = z)$$

and of

$$\begin{aligned} & (h_j(x, y_1, \dots, y_n) = z \wedge S(x) = u) \\ & \rightarrow (h_t(x, z, y_1, \dots, y_n) = v \leftrightarrow h_j(u, y_1, \dots, y_n) = v). \end{aligned}$$

This completes our description of the construction of T_0 .

To prove that (ii) implies (i), we can follow the proof of the previous theorem, giving every primitive recursive function g its usual interpretation on $\omega \times \{0\}$ and letting it be undefined elsewhere.

Conversely, in our proof that (i) implies (ii), we define B for our new formula $\bigwedge T_0 \wedge \varphi^N$. It will still have measure 1. Then for any $(a_0, a_1, \dots) \in B$, the axioms given above hold classically in $\mathcal{M} \upharpoonright (\{a_0, a_1, \dots\} \cap N^{\mathcal{M}})$ and therefore the primitive recursive functions h_i , so in particular g_1, \dots, g_n , have their usual or true interpretation in this restricted model. So, if we perform the construction above with this new B , then we get a model $\mathcal{M} \upharpoonright \{b_0, b_1, \dots\}$ in which φ holds which is not just isomorphic to \mathbb{N} , but is even isomorphic to \mathbb{N} enlarged with the usual interpretations of g_1, \dots, g_n . Therefore we still see that (i) implies (ii) even with the addition of these primitive recursive functions. \square

Next, we will want to include relation symbols of the form $\{e\}^Q(x_1, x_2) = y$. Here, $\{e\}^Q(x_1, x_2) = y$ means that the e th partial computable functional halts with oracle Q and input x_1, x_2 and outputs y , or in the notation of the normal form theorem for partial computable restricted functionals (see e.g. Odifreddi [9, Theorem II.3.11]), it means that $\exists z \mathcal{T}_{1,2}(e, x_1, x_2, \hat{Q}(z), z)$ and $\mathcal{U}(\mu z[\mathcal{T}_{1,2}(e, x_1, x_2, \hat{Q}(z), z)]) = y$.

To this end, we add ternary relation symbols $\{e\}^Q(x_1, x_2) = y$ to our language for all $e \in \omega$ (so, we do not see e or Q as a variable here; we add a relation symbol

for every e and the Q is purely notational). Using the normal form of partial computable restricted functionals we can now easily expand our result to include these functionals.

Theorem 7.10. *Let $f_0, f_{\frac{1}{4}}, f_{\frac{1}{2}}$ and $f_{\frac{3}{4}}$ be as in Theorem 7.6. Then there exists finite theories $T_0, T_{\frac{1}{4}}, T_{\frac{1}{2}}$ and $T_{\frac{3}{4}}$, containing only universal relational sentences, and a computable function $e \mapsto \psi_e$ mapping each $e \in \omega$ to a universal relational sentence, such that for each universal sentence φ in the language of arithmetic enlarged with relation symbols $Q, \{0\}^Q, \{1\}^Q, \dots$, the following are equivalent:*

- (i) $f_0((\bigwedge T_0 \wedge \bigwedge_{\{e\}^Q \in \varphi} \psi_e \wedge \varphi^N)^\#) \wedge f_{\frac{1}{4}}(\bigwedge T_{\frac{1}{4}}) \wedge f_{\frac{3}{4}}(\bigwedge T_{\frac{3}{4}}) \wedge \bigwedge T_{\frac{1}{2}}$ is $\frac{1}{2}$ -satisfiable;⁹
- (ii) $\mathbb{N} \models \exists Q(\varphi)$, i.e. there exists some $Q^{\mathbb{N}} \subseteq \omega$ such that $\mathbb{N} \models \varphi$ if Q is interpreted as $Q^{\mathbb{N}}$ and $\{e\}^Q(x_1, x_2) = y$ is interpreted as $\{e\}^{Q^{\mathbb{N}}}(x_1, x_2) = y$.

Proof. We let \mathcal{U} and $\mathcal{T}_{1,2}$ represent the primitive recursive functions from the normal form theorem for partial computable restricted functionals, as discussed above. Furthermore, we let $\mathcal{V}(x, y)$ represent the primitive recursive function which outputs $z + 1$ if y has length at least x and the x th digit equals z , and which outputs 0 otherwise. Now let the T_ε be as in Theorem 7.9, applied to \mathcal{U} , $\mathcal{T}_{1,2}$ and \mathcal{V} .

We let $\mathcal{W}(x, y)$ represent “the sequence x is an approximation to Q up to digit y ”, so we add to T_0 the relativisation to N of the universal closure of

$$\begin{aligned} (0 = y \wedge 0 = u \wedge S(u) = v \wedge S(v) = w) \\ \rightarrow (\mathcal{W}(x, y) \leftrightarrow ((\mathcal{V}(y, x) = v \leftrightarrow \neg Q(y)) \wedge (\mathcal{V}(y, x) = w \leftrightarrow Q(y)))) \end{aligned}$$

and of

$$\begin{aligned} (S(z) = y \wedge 0 = u \wedge S(u) = v \wedge S(v) = w) \\ \rightarrow (\mathcal{W}(x, y) \leftrightarrow (\mathcal{W}(x, z) \wedge (\mathcal{V}(y, x) = v \leftrightarrow \neg Q(y)) \wedge (\mathcal{V}(y, x) = w \leftrightarrow Q(y)))) \end{aligned}$$

Finally, for each $e \in \omega$, let ψ_e be the relativisation to N of the universal closure of

$$\mathcal{W}(u, z) \wedge \mathcal{T}_{1,2}(e, x_1, x_2, u, z) \rightarrow (\{e\}^Q(x_1, x_2) = y \leftrightarrow \mathcal{U}(z) = y).$$

The proof now proceeds in the same way as for Theorem 7.9, observing that our new axioms ensure that $\{e\}^Q$ gets the right interpretation. \square

So, now we have Π_1^0 -hardness. To get to Σ_1^1 -hardness, we first put the Σ_1^1 -formulas into a more appropriate form.

Lemma 7.11. *Let \mathcal{A} be the set of those universal (relational) formulas φ in the language of arithmetic enlarged with relation symbols $Q, \{0\}^Q, \{1\}^Q, \dots$ for which there exists some interpretation $Q^{\mathbb{N}} \subseteq \omega$ satisfying:*

- $\mathbb{N} \models \varphi$ if Q is interpreted as $Q^{\mathbb{N}}$ and $\{e\}^Q(x_1, x_2) = y$ is interpreted as $\{e\}^{Q^{\mathbb{N}}}(x_1, x_2) = y$;
- For all n there is some $m \in Q^{\mathbb{N}}$ coding a string of length n .

Then \mathcal{A} is Σ_1^1 -hard.

Proof. The set of indices of non-well-founded computable trees is Σ_1^1 -hard (for example, see Odifreddi [9, Corollary IV.2.16]). This set many-one reduces to \mathcal{A} : for every index $e \in \omega$, the expression “ $\{e\}$ computes a tree and $Q \subseteq \{e\}$ ” can be written as a universal formula $\varphi(Q)$ in the language as above. Now $\{e\}$ computes a non-well-founded tree if and only if there exists a Q as above, completing our proof. \square

⁹By $\{e\}^Q \in \varphi$ we mean that the relation symbol $\{e\}^Q$ occurs in the formula φ .

So, we need to somehow expand our theories in such a way that it only allows those Q which contain a string of every possible length. To achieve this, we will add even more relation symbols to our language, and we will use some of these as additional oracles. Let $A \subseteq \omega \times X$ for some set X . We view elements from the set X as parameters: i.e. for all $p \in X$ we will write $\{e\}^{Q \oplus A_p}(x_1, x_2) = y$ if the e th partial computable function halts with oracle $Q \oplus \{n \in \omega \mid (n, p) \in A\}$ and input x_1, x_2 , and outputs y . We will add a binary relation symbol A , a ternary relation symbol B and a 5-ary relation symbol $\{e\}^{Q \oplus A_p}(x_1, x_2) = y$ for every $e \in \omega$ to our language. Using the same method as in the proof of Theorem 7.10 above, we will guarantee that for every model $(\mathcal{M}, \mathcal{D})$ satisfying our theory we have for almost all $p \in \mathcal{M}$ that our interpretation of $\{e\}^{Q \oplus A_p}(x_1, x_2) = y$ in the model \mathcal{M} agrees with its true interpretation in \mathbb{N} . We will combine this with the construction used in Kuyper and Terwijn [8] to show that compactness does not hold for ε -logic; in some sense, compactness failing is the same as forcing infinite objects to exist, which explains the correlation with our current goal.

Theorem 7.12. *Let $f_0, f_{\frac{1}{4}}, f_{\frac{1}{2}}$ and $f_{\frac{3}{4}}$ be as in Theorem 7.6. Then there exists finite theories $T_0, T_{\frac{1}{4}}, T_{\frac{1}{2}}$ and $T_{\frac{3}{4}}$, containing only universal relational sentences, and a computable function $e \mapsto \psi_e$ mapping each $e \in \omega$ to a universal relational sentence, such that for each universal sentence φ in the language of arithmetic enlarged with relation symbols $Q, \{0\}^Q, \{1\}^Q, \dots$, the following are equivalent:*

- (i) $f_0((\bigwedge T_0 \wedge \bigwedge_{\{e\}^Q \in \varphi} \psi_e \wedge \varphi^N)^\#) \wedge f_{\frac{1}{4}}(\bigwedge T_{\frac{1}{4}}) \wedge f_{\frac{3}{4}}(\bigwedge T_{\frac{3}{4}}) \wedge \bigwedge T_{\frac{1}{2}}$ is $\frac{1}{2}$ -satisfiable;
- (ii) φ is in the set \mathcal{A} of Lemma 7.7.

Proof. We start with the theories T_ε and the universal sentences ψ_e as in Theorem 7.10. We need to expand our theory in such a way that case (i) can only hold if Q contains a string of length n for every $n \in \omega$, using the relation symbols A and B .

For this, we will make use of the construction which was used to show that ε -logic is not compact in Kuyper and Terwijn [8, Theorem 8.2]. There, sentences α_n were defined such that for any model $(\mathcal{M}, \mathcal{D})$ for some language containing a binary relation symbols symbol E we have that $(\mathcal{M}, \mathcal{D}) \models_{\frac{1}{2}} \alpha_n$ if and only if we have that:

For almost all y (i.e. measure 1 many), there exists a set C_y of measure at least $1 - \frac{1}{n}$ such that for all $y' \in C_y$ the sets $D_y = \{u \mid E(u, y)\}$ and $D_{y'} = \{u \mid E(u, y')\}$ both have measure $\frac{1}{2}$, while $D_y \cap D_{y'}$ has measure $\frac{1}{4}$ (in other words, the two sets are independent sets of measure $\frac{1}{2}$). In the proof of [8, Theorem 8.2], it was shown that every finite subset of $\{\alpha_1, \alpha_2, \dots\}$ is satisfiable, but no infinite subset is satisfiable.

Fix an index e_1 such that $\{e_1\}^Q(n, m)$ computes the indicator function of “there exists some $n' \leq n$ such that n' codes a string of length m and $n' \in Q$ ”. Our idea is to let α_n hold, with E interpreted as $\{(x, y) \mid B(m, x, y)\}$, if and only if $\{e_1\}^Q(n, m) = 0$. Then, as discussed above, only finitely many of these α_n are allowed to hold for any fixed m , so we see that there has to be some $n \in \omega$ such that $\{e\}^Q(n, m) = 1$, i.e. Q contains a string of length m . This is precisely what we want to achieve.

To this end we first need to change the definition of α_n from [8] into a more uniform one. In [8], we have a different formula for every n (which in fact grows exponentially in n). We wish to change it into one fixed formula $\tilde{\alpha}$ which has n as a parameter, i.e. we want one formula $\tilde{\alpha}$ such that α_n is equivalent to $\tilde{\alpha}(n)$ for every $n \in \omega$. This is what we will use A for. If we unfold the first two reductions in the definition of α_n from [8] for $n \geq 2$, and rename some of the relation symbols

X_i, Y_i , we obtain the following formula:

$$\begin{aligned} & \forall y Y_1(y) \wedge \forall y Y_2(y) \\ & \wedge \bigwedge_{I \subseteq \{3, \dots, 2n\}, |I|=n-1} \forall y \left(\bigvee_{i \in I} Y_i(y) \right) \\ & \wedge \bigwedge_{1 \leq i \leq 2} \forall y \left(Y_i(y) \right) \\ & \wedge \bigwedge_{3 \leq i_1 < i_2 < \dots < i_n \leq 2n} \forall y' \left((Y_{i_1}(y) \vee \dots \vee Y_{i_n}(y)) \wedge \alpha'(y, y') \right) \end{aligned}$$

where α' is a formula that does not depend on n . Here, we denoted by $Y_i(x)$ the formula $X_i(x) \wedge \bigwedge_{1 \leq j \leq 2, j \neq i} \neg X_j(x)$ for $i \in \{1, 2\}$ and $X_i(x) \wedge \bigwedge_{3 \leq j \leq 2n, j \neq i} \neg X_j(x)$ for $i \in \{3, \dots, 2n\}$, and the X_i are new predicates.

Because we want to make a more uniform version of this formula, we want to use the single binary predicate A instead of these predicates X_i and Y_i . To this end, fix an index e_2 such that for all oracles Q, C and all binary sequences σ we have that $\{e_2\}^{Q \oplus C}(\sigma, \langle n, m \rangle)$ computes the indicator function of

$$\exists i \in \sigma \left(C(\langle \langle n, m \rangle, i \rangle) \wedge \bigwedge_{3 \leq j \leq 2n, j \neq i} \neg C(\langle \langle n, m \rangle, j \rangle) \right).$$

Next, fix an index e_3 such that $\{e_3\}^Q(\sigma, n)$ computes if σ is a subset of $\{3, \dots, 2n\}$ of size $n-1$, and similarly an index e_4 which computes if σ is a subset of $\{3, \dots, 2n\}$ of size n . Finally, fix an index e_5 such that $\{e_5\}^Q(n, m)$ computes the pairing function $\langle x, y \rangle$. We can now define our uniform version $\tilde{\alpha}(n, m)$ of α_n :

$$\begin{aligned} & \forall y A(\langle \langle n, m \rangle, 1 \rangle, y) \wedge \forall y A(\langle \langle n, m \rangle, 2 \rangle, y) \\ & \wedge \forall u (\{e_3\}^Q(u, n) \rightarrow \forall y \{e_2\}^{Q, A_y}(u, \langle n, m \rangle)) \\ & \wedge \bigwedge_{1 \leq i \leq 2} \forall y \left(A(\langle \langle n, m \rangle, i \rangle, y) \right. \\ & \left. \wedge \forall u (\{e_4\}^Q(u, n) \rightarrow \forall y' (\{e_2\}^{Q, A_{y'}}(u, \langle n, m \rangle) \wedge \tilde{\alpha}'(m, y, y')) \right) \end{aligned}$$

Here the formula $\tilde{\alpha}'(m, y, y')$ is the formula obtained by replacing each occurrence of $E(x, y)$ by $B(m, x, y)$ in α' . We freely wrote $\alpha(0)$ for $\forall x (N(x) \wedge (0 = x \rightarrow \alpha(x)))$; we will interpret this formula for $\varepsilon = \frac{1}{2}$, and because N has measure $\frac{1}{2}$ this essentially means that for almost all $x \in N$ we have that $0 = x \rightarrow \alpha(x)$. We did similar things for 1, for 2 and for the pairing function $\langle x, y \rangle$, using the index e_5 for the latter.

We now add the following formula to $T_{\frac{1}{2}}$:

$$\forall n \forall m (N(n) \wedge N(m) \wedge (\{e_1\}^Q(n, m) = 0 \rightarrow \tilde{\alpha}(n, m))).$$

Furthermore, we add ψ_e to T_0 for $e \in \{e_1, e_3, e_4, e_5\}$. Like in the proof of Theorem 7.10, we add an axiom for $\mathcal{W}'(x, y, z)$ specifying that $\mathcal{W}'(x, y, z)$ holds if and only if the sequence x is an approximation to $Q \oplus A_z$, where $A_z = \{n \mid (n, z) \in A\}$, up to digit y . Finally, we add a rule for e_2 to T_0 , namely the universal closure of

$$\mathcal{W}'(u, z, v) \wedge \mathcal{T}_{1,2}(e_2, x_1, x_2, u, z) \rightarrow (\{e_2\}^{Q \oplus A_v}(x_1, x_2) = y \leftrightarrow \mathcal{U}(z) = y),$$

where all quantifiers except the one for v are relativised to N .

We prove that the theories defined above are as desired. First, assume (ii) holds. We define a model \mathcal{M} in a similar way as in the proof of Theorem 7.10. However,

this time we do not take the universe to be two copies of ω , but instead we take our universe to be the unit interval. We let our measure be the Lebesgue measure. Fix a partition $\{U_0, U_1, \dots\} \cup \{V_0, V_1, \dots\}$ of $[0, 1]$ such that all U_n, V_n are measurable with measure $\frac{1}{2^{n+2}}$. Looking back at our proof of Theorem 7.10 (and Theorem 7.6), we identify the elements of each U_n with the element n in the first copy of ω and each V_n with the element n in the second copy of ω . This allows us to define the relations in the language of arithmetic in the same way as before: for example, we let $x + y = z$ hold precisely if $x \in U_n, y \in U_m$ and $z \in U_{n+m}$ for some $n, m \in \omega$. We let $N(x)$ hold if $x \in U_n$ for some $n \in \omega$. With this in mind, it should be clear from the proof of Theorem 7.10 how to define R and all the other relations already appearing in the proof of that theorem.

The only relations of which it might not be directly clear how to define them are A and B . How to define B is explained in the proof of [8, Theorem 8.2], while how to define A follows from the proof of Theorem 2.12.

For the converse, assume (i) holds. Then, as in the proof of Theorem 7.10, we can find a model $\mathcal{M} \upharpoonright \{b_0, b_1, \dots\}$ which is isomorphic to \mathbb{N} and a relation Q on $\{b_0, b_1, \dots\}$, such that $\{e_i\}^Q, \{e_2\}^{Q \oplus A_p}$ have their usual interpretation and such that φ holds in this model. Using similar ideas as in the proof of that theorem, we can even assume that our model satisfies our new axioms in the following way:

If $\{e_1\}^Q(b_n, b_m)$ does not hold, then $\tilde{\alpha}(b_n, b_m)$ holds in the sense that all of the following hold:

$$\Pr_{\mathcal{D}}[y \in \mathcal{M} \mid \mathcal{M} \models A(\langle\langle b_n, b_m \rangle, b_1 \rangle, y)] \geq \frac{1}{2}$$

$$\Pr_{\mathcal{D}}[y \in \mathcal{M} \mid \mathcal{M} \models A(\langle\langle b_n, b_m \rangle, b_2 \rangle, y)] \geq \frac{1}{2}$$

for all $u \in \{b_0, b_1, \dots\}$, if $\mathcal{M} \models \{e_3\}^Q(u, n)$

$$\text{then } \Pr_{\mathcal{D}}[y \in \mathcal{M} \mid \mathcal{M} \models \{e_2\}^{Q \oplus A_y}(u, \langle b_m, b_m \rangle)] \geq \frac{1}{2}.$$

$$\text{for } 1 \leq i \leq 2, \Pr_{\mathcal{D}} \left[y \in \mathcal{M} \mid \mathcal{M} \models A(\langle\langle b_n, b_m \rangle, b_i \rangle, y) \right.$$

and for all $u \in \{b_0, b_1, \dots\}$, if $\mathcal{M} \models \{e_4\}^Q(u, n)$

$$\left. \text{then } \Pr_{\mathcal{D}}[y' \in \mathcal{M} \mid (\mathcal{M}, \mathcal{D}) \models_{\frac{1}{2}} \{e_2\}^{Q \oplus A_{y'}}(u, \langle n, m \rangle) \wedge \tilde{\alpha}'(m, y, y')] \geq \frac{1}{2} \right] \geq \frac{1}{2}.$$

Now one can directly verify (using the proof of Theorem 2.12) that under this interpretation, $\tilde{\alpha}(b_n, b_m)$ is equivalent to the formula α_n with $E(x, y)$ replaced by $\{(x, y) \mid B(m, x, y)\}$. However, we also have that if Q does not contain a string of length m , then $\tilde{\alpha}(b_n, b_m)$ holds for all $n \in \omega$, which is a contradiction (as explained above). So, we see that Q contains a string of every length and therefore (ii) holds. \square

Theorem 7.13. *For rational $\varepsilon \in (0, 1)$ we have that ε -satisfiability is Σ_1^1 -hard.*

Proof. From Lemma 7.7 and Theorem 7.12. \square

Theorem 7.14. *Let \mathcal{L} be a countable language not containing equality and function symbols and let $\varepsilon \in (0, 1)$ be rational. Then ε -satisfiability is Σ_1^1 -complete.*

Proof. The Σ_1^1 -hardness was shown in Theorem 7.13, while Theorem 5.5 gives the matching upper bound. \square

8. COMPACTNESS OF 0-LOGIC

We conclude this paper by showing that 0-logic is compact (when considering languages not containing equality and function symbols), which will quite directly follow from the results from the previous section. This contrasts the fact that for rational $\varepsilon \in (0, 1)$ we have that ε -logic is not compact, as shown in Kuyper and Terwijn [8, Theorem 8.2].

First, we need the following strengthening of Theorem 4.2.

Theorem 8.1. *Let \mathcal{L} be a language not containing equality and function symbols (but it may contain constant symbols) and let Γ be a countable set of formulas in \mathcal{L} . Then there exists a language \mathcal{L}' only containing relation symbols and a computable function mapping each formula $\varphi \in \Gamma$ in the language \mathcal{L} to a universal formula φ' in the language \mathcal{L}' such that for every $\varepsilon \in [0, 1]$ and every subset $\Delta \subseteq \Gamma$: Δ is ε -satisfiable if and only if $\Delta' = \{\varphi' \mid \varphi \in \Delta\}$ is ε -satisfiable.*

Proof. We prove this in almost the same way as Theorem 4.2. We construct φ' as in that theorem, but instead of using relations R_α to define φ' we use relations $R_{\alpha, \varphi}$. That is, we introduce new relation symbols for every formula φ , instead of reusing the same ones. There is one exception to this rule: when $\alpha = 0$ (i.e. the function which is constantly 0) we always use the relation R_α , which does not depend on φ .

Now let $\Delta \subseteq \Gamma$. For ease of notation we assume Δ is infinite, say $\Delta = \{\varphi_0, \varphi_1, \dots\}$; the finite case follows in the same way. If Δ is ε -satisfiable, then one can prove that Δ' is ε -satisfiable in the same way as in the proof of Theorem 4.2. For the converse we also use a similar proof, but we need to make some slight modifications. Note that, in the proof of Theorem 4.2, we used the measure 0 Cantor set $\mathcal{C} \subseteq [0, 1]$ to provide us with witnesses. Now we partition \mathcal{C} into countably many uncountable Borel measure 0 sets. For example, we can take $\mathcal{C}_i \subseteq \mathcal{C}$ to be those $x \in \mathcal{C}$ which correspond to a sequence starting with $0^i 1$. We can then take Borel isomorphisms $\zeta^i : \mathcal{C}_i \rightarrow \bigcup_{1 \leq i \leq n} [0, 1]^i$ and use ζ^i to provide the witnesses for φ_i .

More precisely, we show how to modify the definition of R^M in the proof of Theorem 4.2. Let $a_1, \dots, a_k \in [0, 1]$. If there are $1 \leq i_1 < i_2 \leq k$ such that $a_{i_1} \in \mathcal{C}_{j_1}, a_{i_2} \in \mathcal{C}_{j_2}$ but $j_1 \neq j_2$, we let $R^M(a_1, \dots, a_k)$ be false (but it does not really matter how we define R^M in this case, as long as we make sure that the resulting relation is Borel; for example, we could also define it to always be true in this case). Otherwise, we let $r \in \omega$ be such that for all $1 \leq i \leq k$ we have that $a_i \in \mathcal{C}$ implies that in fact $a_i \in \mathcal{C}_r$. Let $\alpha(i) = (1, j)$ if $a_i \in \mathcal{C}$ and $\zeta(a_i)$ has length j , let $\alpha(i) = (2, j)$ if $a_i \in \text{ran}(\eta)$ and $\eta^{-1}(a_i) = j$ and finally let $\alpha(i) = 0$ if neither of these cases hold. Let $b_1, \dots, b_{k'}$ be the subsequence of a_1, \dots, a_k obtained by taking just those a_i satisfying $\alpha(a_i) = 0$. Let $m = \max\{j \in \omega \mid (1, j) \in \text{ran}(\alpha)\}$ and let i be the least $1 \leq j \leq k$ such that $\alpha(j) = (1, m)$. Now let $R^M(a_1, \dots, a_k)$ be defined as $R_{\alpha, \varphi_r}^M(b_1, \dots, b_{k'}, (\zeta^r)_1(a_i), \dots, (\zeta^r)_m(a_i))$. The rest of the proof proceeds in the same way as for Theorem 4.2. \square

Theorem 8.2. *For countable languages \mathcal{L} not containing equality and function symbols, 0-logic is compact. That is, if Γ is any (countable) set of sentences such that any finite subset of Γ is 0-satisfiable, then Γ is 0-satisfiable.*

Proof. By Theorem 8.1 we may assume all formulas in Γ to be universal relational formulas. For every formula $\varphi = \forall x_1 \dots \forall x_n \psi(x_1, \dots, x_n)$, let

$$\tilde{\varphi} = \forall x_1 \dots \forall x_n \bigwedge_{i < j} (x_i \neq x_j \rightarrow \psi(x_1, \dots, x_n)).$$

Let $\tilde{\Gamma} = \{\tilde{\varphi} \mid \varphi \in \Gamma\}$. We claim: $\tilde{\Gamma}$ has an infinite classical model.

By classical compactness, it is enough to show that for every $m \in \omega$ and every finite subset $\tilde{\Delta} \subseteq \tilde{\Gamma}$ there is a model for $\tilde{\Delta}$ of size at least m . Let $\Delta = \{\varphi \mid \tilde{\varphi} \in \tilde{\Delta}\}$. Then by assumption we have that $\bigwedge \Delta$ is 0-satisfiable. So, by Proposition 6.1 we see that $\widetilde{\bigwedge \Delta}$ has a classical model \mathcal{M} of size at least m . However, it is directly seen that then in fact $\mathcal{M} \models \tilde{\Delta}$, so \mathcal{M} is a classical model for $\tilde{\Delta}$ of size at least m .

So, $\tilde{\Gamma}$ has an infinite classical model. Using the infinite Ramsey Theorem we then know that $\tilde{\Gamma}$ has a model on $[0, 1]$ with $([0, 1], <)$ as a sequence of indiscernibles, see e.g. Chang and Keisler [3, Theorem 3.3.10]. Because $([0, 1], <)$ is a sequence of indiscernibles we can directly see that the relations on this model are Borel (in a similar way as in the proof of Corollary 6.4), and because the diagonal has Lebesgue-measure 0 we see that $(\mathcal{M}, \lambda) \models_0 \Gamma$, as desired. \square

Acknowledgment. The author wishes to thank Sebastiaan Terwijn for helpful discussions on the subject.

REFERENCES

1. G. Birkhoff, *Lattice theory*, revised ed., American Mathematical Society, 1948.
2. V. I. Bogachev, *Measure theory*, vol. 1, Springer, 2007.
3. C. C. Chang and H. J. Keisler, *Model theory*, third ed., Studies in Logic and the Foundations of Mathematics, vol. 73, North Holland, 1990.
4. D. N. Hoover, *Probability logic*, *Annals of Mathematical Logic* **14** (1978), 287–313.
5. A. S. Kechris, *Classical descriptive set theory*, Springer-Verlag, 1995.
6. H. J. Keisler, *Probability quantifiers*, *Model-Theoretic Logics* (J. Barwise and S. Feferman, eds.), Perspectives in Mathematical Logic, vol. 8, Springer-Verlag, 1985.
7. R. Kuyper, *Computational hardness of validity in probability logic*, *Proceedings of Logical Foundations of Computer Science* (S. Artemov and A. Nerode, eds.), *Lecture Notes in Computer Science*, vol. 7734, Springer, 2013, pp. 252–265.
8. R. Kuyper and S. A. Terwijn, *Model theory of measure spaces and probability logic*, *The Review of Symbolic Logic* **6** (2013), no. 3, 367–393.
9. P. G. Odifreddi, *Classical recursion theory*, *Studies in Logic and the Foundations of Mathematics*, vol. 125, North-Holland, 1989.
10. F. P. Ramsey, *On a problem of formal logic*, *Proceedings of the London Mathematical Society, Series 2* **30** (1930), 264–286.
11. S. A. Terwijn, *Probabilistic logic and induction*, *Journal of Logic and Computation* **15** (2005), no. 4, 507–515.
12. ———, *Decidability and undecidability in probability logic*, *Proceedings of Logical Foundations of Computer Science* (S. Artemov and A. Nerode, eds.), *Lecture Notes in Computer Science*, vol. 5407, Springer, 2009, pp. 441–450.
13. L. G. Valiant, *Robust logics*, *Artificial Intelligence* **117** (2000), 231–253.

(Rutger Kuyper) RADBOUD UNIVERSITY NIJMEGEN, DEPARTMENT OF MATHEMATICS, P.O. BOX 9010, 6500 GL NIJMEGEN, THE NETHERLANDS.

E-mail address: r.kuyper@math.ru.nl