

Number theory combination: natural density and SMT

Guilherme Toledo and Yoni Zohar

Bar-Ilan University, Israel

Abstract. The study of theory combination in Satisfiability Modulo Theories (SMT) involves various model theoretic properties (e.g., stable infiniteness, smoothness, etc.). We show that such properties can be partly captured by the natural density of the spectrum of the studied theories, which is the set of sizes of their finite models. This enriches the toolbox of the theory combination researcher, by providing new tools to determine the possibility of combining theories. It also reveals interesting and surprising connections between theory combination and number theory.

1 Introduction

Imagine this: you are a researcher in Satisfiability Modulo Theories (SMT) [2], studying a theory \mathcal{T} , which is the combination of theories \mathcal{T}_1 and \mathcal{T}_2 (in the same way that, say, the theory of lists of integers is the combination of the theories of lists and integers). Given algorithms for \mathcal{T}_1 and \mathcal{T}_2 , you can plug them together using theory combination methods, such as Nelson and Oppen’s method [13], polite combination [15], shiny combination [21], or gentle combination [8].

But, before you can produce a decision procedure for \mathcal{T} , you must test certain properties of \mathcal{T}_1 and \mathcal{T}_2 , or their absence, to determine their applicability to the combination method. For example, using the Nelson-Oppen method requires that both theories are stably infinite, while using the polite combination method requires that one of them is strongly polite. The obvious way of doing so is by directly applying the definitions of these properties, what can be highly non-trivial (for example, to prove a theory is strongly polite, one needs to construct a computable function satisfying an involved set of conditions).

In this paper, we give you alternative tests, based on number theoretic *natural densities* [20], computed over the *spectrum* of the theory [11]. When testing whether a theory admits or lacks a theory combination property, you can now use these tests. We provide examples for cases where this is simpler to do, compared to the direct application of the definitions. Beyond the introduction of such tools, the results of this paper relate number theory and theory combination in surprising and insightful ways. We focus on one-sorted theories, leaving many-sorted ones for future work.

Section 2 surveys relevant notions. Section 3 contains our main results: sufficient and necessary conditions for theory combination properties, in terms of the

$\#(x_1, \dots, x_n) = \bigwedge_{i=1}^{n-1} \bigwedge_{j=i+1}^n \neg(x_i = x_j)$	$\psi_{\leq n} = \exists x_1, \dots, x_n. \forall y. \bigvee_{i=1}^n y = x_i$
$\psi_{\geq n} = \exists x_1 \dots x_n. \#(x_1, \dots, x_n)$	$\psi_{=n} = \psi_{\geq n} \wedge \psi_{\leq n}$

Fig. 1: Cardinality formulas.

natural density. In Section 4 we provide generalizations to non-empty signatures. Section 5 summarizes, and gives directions for future research.

2 Preliminaries

If X is a set, $|X|$ denotes its cardinality. We denote by \aleph_0 the cardinality of \mathbb{N} , which for us contains 0; the set $\mathbb{N} \setminus \{0\}$ is denoted by \mathbb{N}^* .

2.1 First-order logic

A first-order *signature* Σ is a pair $(\mathcal{F}_\Sigma, \mathcal{P}_\Sigma)$, where: \mathcal{F}_Σ is a countable set of function symbols, each with an arity $n \in \mathbb{N}$; and \mathcal{P}_Σ is a countable set of predicate symbols, each with an arity $n \in \mathbb{N}$, containing at least the equality $=$ of arity 2. We denote by Σ_1 the signature with no function or predicate symbols other than $=$, which is therefore called *empty*. Assuming countably many variables, we define by structural induction *terms*, *literals*, *formulas*, and *sentences* (formulas without free variables) in the usual way. The set of all quantifier-free Σ -formulas is denoted by $QF(\Sigma)$; the set of all variables in φ shall be written as $vars(\varphi)$.

A Σ -*interpretation* \mathcal{A} consists of: a non-empty set $dom(\mathcal{A})$, called the *domain* of \mathcal{A} ; for each function symbol f of arity n , a function $f^{\mathcal{A}} : dom(\mathcal{A})^n \rightarrow dom(\mathcal{A})$; for each predicate symbol P of arity n , a subset $P^{\mathcal{A}}$ of $dom(\mathcal{A})^n$, where $=^{\mathcal{A}}$ is the identity; and, for every variable x , an element $x^{\mathcal{A}}$ of $dom(\mathcal{A})$. The value of a term α in \mathcal{A} is denoted by $\alpha^{\mathcal{A}}$, while for a set of terms Γ we make $\Gamma^{\mathcal{A}} = \{\alpha^{\mathcal{A}} : \alpha \in \Gamma\}$; if \mathcal{A} satisfies the formula φ , we write $\mathcal{A} \models \varphi$. Recurrent formulas include those in Figure 1, that are satisfied by an interpretation \mathcal{A} iff: \mathcal{A} has at least n elements, in the case of $\psi_{\geq n}$; \mathcal{A} has at most n elements, in the case of $\psi_{\leq n}$; and \mathcal{A} has precisely n elements, in the case of $\psi_{=n}$.

A *theory* is the class of all interpretations (thus called \mathcal{T} -interpretations, or the models of \mathcal{T}) satisfying some set of sentences $Ax(\mathcal{T})$ (which does not need to be computably enumerable), called the *axiomatization* of \mathcal{T} . A formula φ is then: (\mathcal{T} -)*satisfiable* if there is a (\mathcal{T} -)interpretation that satisfies φ ; (\mathcal{T} -)*equivalent* to a formula ψ if every (\mathcal{T} -)interpretation that satisfies one also satisfies the other; and (\mathcal{T} -)*valid* if every (\mathcal{T} -)interpretation satisfies φ , denoted $\models \varphi$ ($\models_{\mathcal{T}} \varphi$).

We denote, for $n \leq m$, the set $\{n, \dots, m\}$ by $[n, m]$; if $n = 0$, we simplify it to $[m]$. Of course, $|[n, m]| = m - n + 1$, and $|[m]| = m + 1$. Furthermore, $A \cap [1, n]$ will be denoted by A_n ; we denote $\{|dom(\mathcal{A})| : \mathcal{A} \text{ is a } \mathcal{T}\text{-interpretation}\} \cap \mathbb{N}$ by

$\text{Spec}(\mathcal{T})$, and we define $\text{Spec}_n(\mathcal{T})$ as $\text{Spec}(\mathcal{T}) \cap [1, n]$. Analogously, $\text{Spec}(\mathcal{T}, \phi)$ is the set of finite cardinalities of \mathcal{T} -interpretations that satisfy ϕ . We can then also define $\text{Spec}_n(\mathcal{T}, \phi)$ as $\text{Spec}(\mathcal{T}, \phi) \cap [1, n]$.

2.2 Number theory

The *natural density* [20] of a set $A \subseteq \mathbb{N}$ is the following real number, if it exists (and then we say the density of A is well-defined): $\mu(A) = \lim_{n \rightarrow \infty} |A \cap [n]|/|[n]|$.

Example 1. Consider the set A of even non-negative integers: we then have that $\mu(A)$ is the limit of the sequence a_n which equals $(n+2)/2(n+1)$ if n is even, and $1/2$ if it is odd, meaning that $\mu(A)$ is well-defined and equals $1/2$.

It is easy to prove that μ satisfies, for all disjoint sets A and B for which it is defined: $0 \leq \mu(A)$; $\mu(\mathbb{N}) = 1$; and $\mu(A \cup B) = \mu(A) + \mu(B)$. The subsets of the non-negative integers we shall calculate the natural density of are sets of finite cardinalities of interpretations in a theory: since they are never zero (as we assume $\text{dom}(\mathcal{A})$ is never empty), we can change $\mu(A)$ to be the limit of the ratio of $|A \cap \{1, \dots, n\}|$ to $|\{1, \dots, n\}| = n$.¹ With this, we can finally define the natural density of a theory (relative to a quantifier-free formula or not) as the natural density of its spectrum: $\mu(\mathcal{T}) = \lim_{n \rightarrow \infty} |\text{Spec}_n(\mathcal{T})|/|[1, n]|$, and $\mu(\mathcal{T}, \phi) = \lim_{n \rightarrow \infty} |\text{Spec}_n(\mathcal{T}, \phi)|/|[1, n]|$.

Definition 1. Let $r \in \mathbb{R}$. r is *computable* [27] if there are computable sequences $\{a_n\}_{n \in \mathbb{N}}$ in \mathbb{Z} and $\{b_n\}_{n \in \mathbb{N}}$ in \mathbb{N}^* with $\lim_{n \rightarrow \infty} a_n/b_n = r$.

Example 2. Every rational number p/q is computable: just take $a_n = p$ and $b_n = q$. The number $\sum_{n=1}^{\infty} 2^{-\varsigma(n)} = 0.57824\dots$ is not computable, for ς the busy beaver function [14], which maps $n \in \mathbb{N}$ to the maximum number of 1's a Turing machine with at most n states can write when it halts, assuming the tape begins with only 0's. Now, $0.57824\dots$ is the limit of $5/10, 57/100, 578/1000, \dots$. Consider then the theory \mathcal{T} with models of size 1 through 5, 11 through $52 = 57 - 5$, 101 through $521 = 578 - 57$, and so on. Its density is the limit of the fractions $5/10, 57/100, 578/1000$ and so on, i.e. $0.57824\dots$, although this number is irrational. More generally, any $0 \leq r \leq 1$ is the density of some theory.

2.3 Theory combination

In what follows, let Σ be an arbitrary signature and \mathcal{T} be a Σ -theory.

\mathcal{T} is **stably infinite** [13] if for every satisfiable quantifier-free formula ϕ , there is a \mathcal{T} -interpretation \mathcal{A} that satisfies ϕ with $|\text{dom}(\mathcal{A})| \geq \aleph_0$. \mathcal{T} is **smooth** when, for all quantifier-free formulas ϕ , \mathcal{T} -interpretations \mathcal{A} that satisfy ϕ , and cardinals $\kappa > |\text{dom}(\mathcal{A})|$, there exists a \mathcal{T} -interpretation \mathcal{B} that satisfies ϕ with $|\text{dom}(\mathcal{B})| = \kappa$. Notice that being smooth implies being stably infinite.

¹ Of course, this does not change the value of $\mu(A)$.

Example 3. The theory axiomatized by $\{\psi_{\geq 3}\}$ is smooth, as we can always add more elements to an interpretation.

\mathcal{T} is **finitely witnessable** [16] when there is a computable function wit (called a witness) from the quantifier-free formulas into themselves such that, for every quantifier-free formula ϕ : (I) ϕ and $\exists \vec{x}.wit(\phi)$ are \mathcal{T} -equivalent, where $\vec{x} = vars(wit(\phi)) \setminus vars(\phi)$; and (II) if $wit(\phi)$ is \mathcal{T} -satisfiable, then there is a \mathcal{T} -interpretation \mathcal{A} that satisfies $wit(\phi)$ and, in addition, $dom(\mathcal{A}) = vars(wit(\phi))^{\mathcal{A}}$ (that is, every element of $dom(\mathcal{A})$ is the interpretation of a variable in $wit(\phi)$). Now, given a finite set of variables V on the signature Σ , and an equivalence E on V , the **arrangement** on V induced by E , written δ_V^E or δ_V if E is clear from context, is the formula $\bigwedge_{xEy}(x = y) \wedge \bigwedge_{x\bar{E}y} \neg(x = y)$, where \bar{E} is the complement of E . Intuitively, an arrangement codifies the relationships between a finite set of variables, that is, if they should be equal or different to one another. \mathcal{T} is then **strongly finitely witnessable** [12] if it has a witness wit (that in this case will be called a strong witness) satisfying, in addition to (I) and (II), the stronger (II^{*}): for every finite set of variables V and arrangement δ_V on V , if $wit(\phi) \wedge \delta_V$ is \mathcal{T} -satisfiable, then there exists a \mathcal{T} -interpretation \mathcal{A} that satisfies that formula and, in addition, $dom(\mathcal{A}) = vars(wit(\phi) \wedge \delta_V)^{\mathcal{A}}$.

Example 4. The theory axiomatized by $\{\psi_{\leq 3}\}$ has as strong witness $wit(\phi) = \phi \wedge \bigwedge_{i=1}^3 x_i = x_i$, where x_1, x_2 and x_3 are fresh variables (i.e., not in ϕ).

\mathcal{T} has the **finite model property (FMP)** if for every \mathcal{T} -satisfiable quantifier-free formula ϕ , there is a \mathcal{T} -interpretation \mathcal{A} that satisfies ϕ with $|dom(\mathcal{A})| < \aleph_0$.² Consider $\mathbb{N}_\omega = \mathbb{N} \cup \{\aleph_0\}$. A **minimal model function** [22] for \mathcal{T} is a function $\mathbf{minmod}_{\mathcal{T}} : QF(\Sigma) \rightarrow \mathbb{N}_\omega$ such that, if ϕ is quantifier-free and \mathcal{T} -satisfiable, then $\mathbf{minmod}_{\mathcal{T}}(\phi) = n$ if, and only if: there exists a \mathcal{T} -interpretation \mathcal{A} that satisfies ϕ with $|dom(\mathcal{A})| = n$; and if \mathcal{B} is another \mathcal{T} -interpretation that satisfies ϕ , with $|dom(\mathcal{B})| \neq n$, then $|dom(\mathcal{B})| > n$.

Example 5. The theory axiomatized by $\{\psi_{\geq 3}\}$ has a computable minimal model function. To calculate it on a quantifier-free formula ϕ , take the cardinality n of the smallest interpretation in equational logic that satisfies ϕ , which can easily be found algorithmically. If $n < 3$, $\mathbf{minmod}(\phi) = 3$; otherwise $\mathbf{minmod}(\phi) = n$.

\mathcal{T} is **(strongly) polite** if it is smooth and (strongly) finitely witnessable. It is **shiny** if it is smooth, has the FMP and a computable minimal model function. \mathcal{T} is **gentle** [8] if for every quantifier-free formula ϕ , $Spec(\mathcal{T}, \phi)$ is *fully computable*, that is: (i) it is computable; (ii) it is either co-finite,³ or a finite set of finite cardinalities, and there is an algorithm with ϕ as input that tells which one is the case; (iii) if $Spec(\mathcal{T}, \phi)$ is finite, $\max(Spec(\mathcal{T}, \phi))$ is computable, and if it is infinite $\max(\mathbb{N} \setminus Spec(\mathcal{T}, \phi))$ is computable, both with ϕ as input.⁴

² A common definition for the finite model property demands this condition holds for all formulas, but in theory combination quantifier-free formulas are typically used.

³ I.e., $\mathbb{N} \setminus Spec(\mathcal{T}, \phi)$ is finite.

⁴ If $Spec(\mathcal{T}, \phi)$ or $\mathbb{N} \setminus Spec(\mathcal{T}, \phi)$ are empty, their respective maxima are 0, as usual, so \mathcal{T} must be decidable as $\max(Spec(\mathcal{T}, \phi)) = 0$ iff ϕ is not \mathcal{T} -satisfiable.

Example 6. Consider the Σ_1 -theory $\mathcal{T}_{\text{even}}^\infty$ (see [25]), with axiomatization $\{\neg\psi_{=2n+1} : n \in \mathbb{N}\}$: it is not gentle, as $x = x$ has as spectrum the set of even positive numbers, which is neither finite nor cofinite.

3 Theory combination and natural density

In this section we establish various connections between model-theoretic properties of a theory, and its natural density. We focus our investigation on the empty signature Σ_1 , that has a single sort and no function and predicate symbols other than equality. Generalizations to non-empty signatures are given in Section 4.

We start with the empty signature, because a theory on such a signature has essentially one natural density, while for the non-empty case we must consider the density with respect to both a formula and the theory (this can also be done on the empty case, but all \mathcal{T} -satisfiable formulas will give the same density). Furthermore, some results will not hold on the non-empty case, such as the third item in Theorem 1 below.

Section 3.1 deals with sufficient conditions: if the density satisfies them, then we can deduce some combination properties. Section 3.2 obtains necessary conditions: one would use the contrapositive and conclude that the theory does not have the properties at hand, and then at least one knows that a different combination method has to be used.

3.1 Sufficient conditions

In Theorem 1 we identify sufficient conditions for stable infiniteness, the finite model property and finite witnessability, properties that are needed for Nelson-Oppen combination, shiny combination, and polite combination, respectively.

Theorem 1. *If \mathcal{T} is a Σ_1 -theory with a well-defined natural density, then the positivity of $\mu(\mathcal{T})$ is sufficient for \mathcal{T} to: 1. be stably infinite; 2. have the finite model property; and 3. be finitely witnessable.*

Proof (sketch). ⁵ The proof of the third item is more involved than that of the first two, which is routine. Szemerédi’s theorem [19], which settled a well-know conjecture by Erdős and Turán, showed that each set with positive natural density contains arbitrarily long finite subsequences in arithmetic progression (i.e., the difference between two consecutive elements is constant). Item 3 is a similarly flavored result, although with a much simpler proof than that of Erdős and Turán, that will guarantee that any theory \mathcal{T} which is not finitely witnessable and has a well-defined natural density must satisfy $\mu(\mathcal{T}) = 0$.

Example 7. Fix some positive natural number n , and consider the theory $\mathcal{T}_{\geq n}$, with axiomatization $\{\psi_{\geq n}\}$. It obviously has positive density. By Theorem 1 it is stably infinite, has the finite model property, and is finitely witnessable.

⁵ Full proofs appear in the appendix.

The following example shows that all the reciprocals of Theorem 1 are false, a single counterexample being enough for all three.

Example 8. Take the Σ_1 -theory $\mathcal{T}_{=2^i}$ with axiomatization $\{\psi_{\geq 2^n} \vee \bigvee_{i=0}^n \psi_{=2^i} : n \in \mathbb{N}\}$, which has interpretations \mathcal{A} with domains whose cardinality is either infinite or a power of two. It is stably infinite, has the finite model property and is finitely witnessable,⁶ but $\mu(\mathcal{T}_{=2^i}) = \lim_{n \rightarrow \infty} \frac{|Spec_n(\mathcal{T}_{=2^i})|}{n} = \lim_{n \rightarrow \infty} \frac{\lfloor \log_2(n) \rfloor + 1}{n} = 0$.

The following example shows the sharpness of Theorem 1, in the sense that its assumption is really needed to reach its conclusions.

Example 9. The conclusion of Theorem 1 cannot hold under the assumption that $\mu(\mathcal{T}) = 0$. The theory \mathcal{T}_∞ , with axiomatization $\{\psi_{\geq n} : n \in \mathbb{N}^*\}$, has only infinite models. It has density 0 but does not have the finite model property. The theory \mathcal{T}_1 , with axiomatization $\{\psi_{=1}\}$, has a single model up to isomorphism, with a single element. It has density 0 but is not stably infinite. For item 3, a theory that is not finitely witnessable and has natural density 0 is \mathcal{T}_ς , from [26], with axiomatization $\{\psi_{\geq \varsigma(n)} \vee \bigvee_{i=2}^n \psi_{=\varsigma(i)} : n \in \mathbb{N} \setminus \{0, 1\}\}$ for $\varsigma : \mathbb{N} \rightarrow \mathbb{N}$ the busy beaver function (see Example 2). The cardinalities of its finite models are precisely the Busy Beaver numbers, that is, the elements of the image of ς . We can show that $\mu(\mathcal{T}) = 0$. In a way, item 3 of Theorem 1 shows that every theory not finitely witnessable must, like ς , "escape" all computable functions, and thus have natural density 0.

3.2 Necessary conditions

We now move on to the results establishing necessary conditions for gentleness (Section 3.2.1), smoothness and finite model property (Section 3.2.2), strong finite witnessability (Section 3.2.3), the computability of a minimal model function (Section 3.2.4), and finite witnessability (Section 3.2.5).

3.2.1 Gentleness

Theorem 2. *If \mathcal{T} is a Σ_1 -theory, then $\mu(\mathcal{T})$ being well-defined and equal to 0 or 1 is a necessary condition for \mathcal{T} to be gentle.*

Proof (sketch). By taking a tautology ϕ for a gentle Σ_1 -theory \mathcal{T} , we see that $Spec(\mathcal{T}) = Spec(\mathcal{T}, \phi)$ is either finite (and then its density is 0) or co-finite (and then its density is 1).

Example 10. Consider the theory $\mathcal{T}_{even}^\infty$ from Example 6: it's density is $1/2$, what implies by the theorem it is not gentle.

The reciprocal of Theorem 2 is false, as shown by the next example.

⁶ A witness being, if ϕ has n variables, $wit(\phi) = \phi \wedge \bigwedge_{i=1}^{2^n} x_i = x_i$, for fresh x_i s.

Example 11. Example 8 presents a theory $\mathcal{T}_{=2^i}$ that has density 0 but is not gentle (since both $\text{Spec}(\mathcal{T}_{=2^i})$ and $\mathbb{N} \setminus \text{Spec}(\mathcal{T}_{=2^i})$ are infinite). On the other hand, take the Σ_1 -theory $\mathcal{T}_{\neq 2^i}$ with axiomatization $\{\neg\psi_{=2^n} : n \in \mathbb{N}\}$, which has interpretations \mathcal{A} with either $|\text{dom}(\mathcal{A})|$ infinite, or $|\text{dom}(\mathcal{A})|$ finite but not a power of two. It is not gentle, yet $\mu(\mathcal{T}_{\neq 2^i}) = \lim_{n \rightarrow \infty} \frac{|\text{Spec}_n(\mathcal{T}_{\neq 2^i})|}{n} = \lim_{n \rightarrow \infty} \frac{n - \lfloor \log_2(n) \rfloor - 1}{n} = 1$.

Notice also that both cases of Theorem 2 are possible, namely: there are gentle theories with density 0 and gentle theories with density 1. Before showing them, let us present two useful lemmas, that relate gentleness to other properties:

Lemma 1. *If \mathcal{T} is gentle, then \mathcal{T} has a computable minimal model function and the finite model property, and therefore is finitely witnessable as well.*

Lemma 2. *Let \mathcal{T} be a Σ_1 -theory: if \mathcal{T} is not stably infinite, or if it is strongly finitely witnessable, then \mathcal{T} is gentle.*

Example 12.

1. The trivial Σ_1 -theory $\mathcal{T}_{\geq 1}$, with axiomatization $\{\psi_{\geq 1}\}$, consists of all Σ_1 -interpretations. It is strongly finitely witnessable (given its axiomatization is given by an universal formula, this is proven in [17]), and of course $\text{Spec}_n(\mathcal{T}_{\geq 1}) = [1, n]$ so $\mu(\mathcal{T}_{\geq 1}) = 1$.
2. The Σ_1 -theory \mathcal{T}_1 from Example 9 is also strongly finitely witnessable and thus gentle (Lemma 2), but $\text{Spec}_n(\mathcal{T}_1) = \{1\}$ so $\mu(\mathcal{T}_1) = 0$.
3. An example of a Σ_1 -theory that is gentle and has density 0, but is not strongly finitely witnessable, is denoted by $\mathcal{T}_{\langle m, n \rangle}$, for any fixed $m, n \in \mathbb{N}^*$. It has axiomatization $\{\psi_{=m} \vee \psi_{=n}\}$, and its models have cardinalities m or n .

3.2.2 Smoothness and finite model property The next result involves both smoothness and the finite model property.

Theorem 3. *If \mathcal{T} is a Σ_1 -theory, $\mu(\mathcal{T})$ being well-defined and equal to 1 is necessary for \mathcal{T} to simultaneously admit smoothness and the finite model property.*

Proof (sketch). The proof is dual to that of Theorem 1: if a theory is smooth and has the finite model property, it has all sufficiently large numbers as cardinalities of its models, and its density is therefore 1.

The following example not only allows one to visualize the use of Theorem 3, but will also help later in providing examples for each and all possible combination of the properties under consideration.

Example 13. Consider again $\mathcal{T}_{\text{even}}^\infty$ from Example 10, with density 1/2. It was already shown in [18] that $\mathcal{T}_{\text{even}}^\infty$ has the finite model property without being smooth, but notice that Theorem 3 perfectly encapsulates an intuition for why that is: as the theory has the finite model property, it has a finite model; were it smooth, it would have models of all larger cardinalities, and thus density 1.

Example 14. The reciprocal of Theorem 3 is false, as we can see from the theory $\mathcal{T}_{\#2^i}$ defined in Example 11, which is not smooth.

Now, Example 13 shows an example of a theory that has the finite model property but is not smooth. But all three other Boolean combinations of these two properties are possible, as seen below.

Example 15.

1. The theory $\mathcal{T}_{\geq 1}$ from Example 12 is smooth and has the finite model property.
2. One example of a smooth theory without the finite model property is the Σ_1 -theory \mathcal{T}_{∞} from Example 9. It has density 0, as it has no finite models.
3. To see one of a theory that is neither smooth nor has the finite model property, which by Theorem 1 must have density 0, fix an $n \in \mathbb{N}^*$ and consider the Σ_1 -theory $\mathcal{T}_{n,\infty}$, defined in [25] by the axiomatization $\{\psi_{=n} \vee \psi_{\geq m} : m \in \mathbb{N}^*\}$. Its finite models must have cardinality n .

3.2.3 Strong finite witnessability The following result, which is a corollary of earlier ones, is specially useful: proving a theory is not strongly finitely witnessable is quite challenging; it involves finding a quantifier-free formula, a set of variables, and an arrangement on that set which fail the conditions to be a strong witness, for every candidate for a strong witness. Checking whether the theory's density is 0 or 1 can be fairly easier.

Theorem 4. *If \mathcal{T} is a Σ_1 -theory, then $\mu(\mathcal{T})$ being well-defined and equal to 0 or 1 is a necessary condition for \mathcal{T} to be strongly finitely witnessable.*

Proof. By Lemma 2 and Theorem 2.

Example 16. The theory $\mathcal{T}_{\text{even}}^{\infty}$ from Example 10 is not strongly finitely witnessable, as proven in [18], but the proof found there is quite involved, demanding careful use of arrangements. Here, we only need to point to the fact that $\mathcal{T}_{\text{even}}^{\infty}$ has natural density $1/2$.

Example 17. The reciprocal of Theorem 4 is false: the theories $\mathcal{T}_{=2^i}$ and $\mathcal{T}_{\#2^i}$ from Examples 8 and 11 have, respectively, densities 0 and 1, but neither is strongly finitely witnessable, which follows from the fact that both are stably infinite without being smooth, together with [25, Theorem 7], which shows stably infinite, one-sorted theories that are strongly finitely witnessable are smooth.

3.2.4 Computability of minimal model functions We move now to the question of computability of a minimal model function. For this, we first establish in Proposition 1 a connection between this and the computability of the spectra.

Proposition 1. *\mathcal{T} is a Σ_1 -theory with a computable minimal model function if, and only if, $\text{Spec}(\mathcal{T})$ is computable.*

This proposition plays an important role in the proof of the theorem below:

Theorem 5. *If \mathcal{T} is a Σ_1 -theory with a well-defined density, the fact that $\mu(\mathcal{T})$ is a computable number is a necessary condition for \mathcal{T} to have a computable minimal model function. Furthermore, for every computable number $0 \leq r \leq 1$, there exists a Σ_1 -theory \mathcal{T} with $\mu(\mathcal{T}) = r$ that has a computable minimal model function and the finite model property, but is not smooth.*

The proof of the first part takes a theory \mathcal{T} with a computable minimal model function, and from Proposition 1 one sees that $\text{Spec}(\mathcal{T})$ is computable; we then prove that this implies $\mu(\mathcal{T})$ is itself a computable real number, ruling out non-computable numbers. Indeed, if A is a computable set, the sequence $\{|A_n|\}_{n \in \mathbb{N}} = \{|\{k \in A : k \leq n\}|\}_{n \in \mathbb{N}}$ is computable (and so is $\{n\}_{n \in \mathbb{N}}$, but that is obvious); if $\mu(A) = r$, we have $r = \lim_{n \rightarrow \infty} |A_n|/n$, proving r is computable.

We prove the second part by constructing in Definition 2 below, from two sequences $\{a_n\}_{n \in \mathbb{N}}$ and $\{b_n\}_{n \in \mathbb{N}}$, a function f whose image (which will also equal the spectrum of the theory to be constructed) will be a computable set and have a density associated to the mediants of the ratios a_n/b_n , where the mediant of the fractions a/b and c/d is the fraction $(a+c)/(b+d)$. Although tedious to prove, it is true that the limit of the mediants of the ratios between two sequences equals the limit of the ration, and this guarantees that the natural density of the image of f will be the limit of a_n/b_n .

Definition 2. *Given sequences $\{a_n\}_{n \in \mathbb{N}}$ and $\{b_n\}_{n \in \mathbb{N}}$ with $0 < a_n < b_n$ and $a_n, b_n \in \mathbb{N}$, for all $n \in \mathbb{N}$, we define an associated function $f : \mathbb{N}^* \rightarrow \mathbb{N}^*$ inductively as follows: $f(n) = n$ for $1 \leq n \leq a_0$, and $f(n) = a_0$ for $a_0 + 1 \leq n \leq b_0$; and, assuming $f(n)$ defined for all $1 \leq n \leq M = \sum_{i=0}^m b_i$, for any $m \geq 0$, we make $f(n) = n$ for $M + 1 \leq n \leq M + a_{m+1}$, and $f(n) = M + a_{m+1}$ for $M + a_{m+1} + 1 \leq n \leq M + b_{m+1}$.*

The construction defined in Definition 2 is outlined in Figure 2 (A). Notice the step-shape of the function, that can be computed by induction on n .

We next show an application of Theorem 5, by identifying a theory that according to this theorem has a computable minimal model function, but this fact seems difficult to prove without using this theorem.

Example 18. Take $\Omega = 0.57824\dots$ from Example 2. Take the sequence of fractions $5/10, 57/100, \dots$ converging to Ω , and define a function f as in Definition 2: so $f(n) = n$ for $1 \leq n \leq 4$, and $f(n) = 5$ for $5 \leq n \leq 10$; $f(n) = n$ for $11 \leq n \leq 66$, and $f(n) = 67$ for $67 \leq n \leq 110$, and so on. Define then a theory \mathcal{T}_Ω with axiomatization $\{\psi_{\geq f(n+1)} \vee \bigvee_{i=1}^n \psi_{=f(i)} : n \in \mathbb{N}^*\}$: it has models of size 1 through 5, 11 through 67, and so on.⁷ We can prove that it has natural density Ω , and thus does not have a computable minimal model function. It is, however unclear how one would prove this without resorting to Theorem 5 and Proposition 1.

The reciprocal of Theorem 5 is false, in the sense that a theory without a computable minimal model function can still have a computable natural density.

⁷ Notice \mathcal{T}_Ω is not the same as the theory from Example 2: indeed, the construction from Definition 2 is more general.

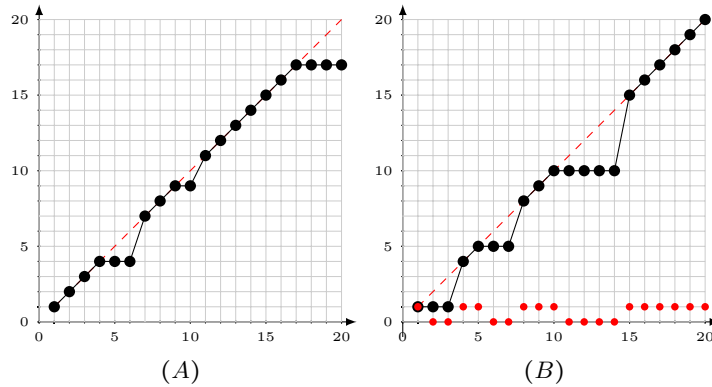


Fig. 2: (A) Initial values of the function f from Definition 2, for $a_0 = 4$, $a_1 = 3$, $a_2 = 7$, $b_0 = 6$, $b_1 = 4$ and $b_2 = 10$. (B) Initial values of the functions G (in red) and f (in black) from the case with computable r of Theorem 6, for $a_0 = 1$, $a_1 = 2$, $a_2 = 3$, $b_0 = 2$, $b_1 = 3$, $b_2 = 5$, $g(1) = 1$, $g(2) = 1$ and $g(3) = 0$.

Example 19. Consider again the theory \mathcal{T}_ζ from Example 9, whose natural density we have shown to be the computable number 0, despite the fact it does not have a computable minimal model function (see [6, Lemma 128]), what by Proposition 1 means $\text{Spec}(\mathcal{T}_\zeta)$ is not computable.

Notice that, from [25, Lemma 7] and [24, Theorem 4], the theories in Theorem 5 are also finitely witnessable without being strongly finitely witnessable. They are also not gentle in the case that $0 < \mu(\mathcal{T}) < 1$. It is still possible to come up with gentle examples for $\mu(\mathcal{T}) = 1$ or $\mu(\mathcal{T}) = 0$, as the next example shows.

Example 20. To obtain Σ_1 -theories \mathcal{T} , with $\mu(\mathcal{T}) = 1$ or $\mu(\mathcal{T}) = 0$, that have a computable minimal model function and are gentle but not strongly finitely witnessable it is enough to consider, for the first case, \mathcal{T} with axiomatization $\{\psi_{=1} \vee \psi_{\geq 3}\}$; for the second, \mathcal{T} with axiomatization $\{\psi_{=1} \vee \psi_{=3}\}$.

3.2.5 Finite witnessability The theorems so far have provided necessary conditions for a theory to be gentle, smooth strongly finitely witnessable, or have a computable minimal model function. We now show that this is as far as this goes: namely, we cannot achieve necessary conditions using natural densities for finite witnessability alone.

In fact, any real number r is the natural density of a finitely witnessable theory. If r is computable then Theorem 5 already constructs a finitely witnessable theory \mathcal{T} with $\mu(\mathcal{T}) = r$, as [24, Theorem 4] proved that a theory with a computable minimal model function is finitely witnessable. However, in the next theorem we construct such a theory also for non-computable numbers. In addition, the theorem shows that the generated theory does not need to have a computable minimal model function even if r is computable.

Theorem 6. *If \mathcal{T} is a Σ_1 -theory, nothing can be said about $\mu(\mathcal{T})$ if \mathcal{T} is only known to be finitely witnessable; that is, for every number $0 \leq r \leq 1$, there exists a Σ_1 -theory \mathcal{T} with $\mu(\mathcal{T}) = r$ that is finitely witnessable yet doesn't have a computable minimal model function.*

The proof of Theorem 6 is divided in two cases: when r is computable, and when it is not. When it is not, we write it in decimal notation, take the obvious series of decimal fractions converging to it, define the function f as in Definition 2 and take the theory whose spectrum is the image of f .

If r is computable, which is the difficult case, we take computable sequences $\{a_n\}_{n \in \mathbb{N}}$ and $\{b_n\}_{n \in \mathbb{N}}$ such that a_n/b_n converges to r , and a non-computable function $g : \mathbb{N}^* \rightarrow \{0, 1\}$. We then define an auxiliary function $G : \mathbb{N}^* \rightarrow \{0, 1\}$ by making, for $M = 2 \sum_{i=0}^m b_i$: $G(M+1) = g(m+2)$; $G(n) = 0$ for $M+2 \leq n \leq 2(b_{m+1} - a_{m+1}) + M$; $G(2(b_{m+1} - a_{m+1}) + M+1) = 0$ if $g(m+2) = 1$, and otherwise $G(2(b_{m+1} - a_{m+1}) + M+1) = 1$; and $G(n) = 1$ for $2(b_{m+1} - a_{m+1}) + M+2 \leq n \leq M+2b_{m+1}$. We make $f(n) = \max\{m \leq n : G(m) = 1\}$, and take the theory whose spectrum is the image of f , axiomatized by $\{\psi_{\geq f(n+1)} \vee \bigvee_{i=1}^n \psi_{=f(i)} : n \in \mathbb{N}^*\}$, which will have a density equal to whatever is the limit of a_n/b_n , i.e. r . An example of this construction appears in Figure 2 (B), where the red dots represent G , and the black ones represent f : for the corresponding theory we will have $\text{Spec}_{20}(\mathcal{T}) = \{1, 4, 5, 8, 9, 10, 15, 16, 17, 18, 19, 20\}$.

3.2.6 Summary We can now, as this section about necessary conditions for the empty signature comes to an end, summarize its overall arch: we have seen what are all possible values for the density of a theory given some of its theory combination properties, for all such combinations of properties.

We have considered, in total, 7 properties related to theory combination. Were we to consider all Boolean combinations of them, we would need to analyze 128 cases; [24, Theorems 5,6,7] has shown, however, that for Σ_1 there are only 8 of these possibilities, excluding gentleness. It may look like we need to analyze 16 possibilities then, but we can cut them down to 9 by using Lemmas 1 and 2.

In Table 1, **SI** stands for stably infinite; **SM** for smooth; **FW** for finitely witnessable; **SW** for strongly finitely witnessable; **FM** for the finite model property; **CF** for a computable minimal model function; and **G** for gentle. **REC** denotes the set of real computable numbers.

Each line in the table corresponds to a possible combination of properties (that remains possible after Lemmas 1 and 2). For example, the first line corresponds to theories that admit all properties, while the second line correspond to theories that are stably infinite, smooth, have a computable minimal model function, but do not admit any of the other property.

For each possible combination of properties, we list in the table the possible natural densities of theories that admit the corresponding properties. For example, theories that admit all properties must have density one.⁸

⁸ The theory found in this specific row, $\mathcal{T}_{\geq 1}$ from Example 12, is strongly finitely witnessable, and Lemma 2 then shows it is also gentle, as implied by the table.

SI	SM	FW	SW	FM	CF	G	Natural densities	Reference	Construction
T	T	T	T	T	T	T	1	Theorem 3	Example 12
		F	F	F	T	F	0	Theorem 1	Example 15
	F	T	F	T	T	T	$\{0, 1\}$	Theorem 2	Example 20
				T	F	F	$\mathbf{REC} \cap [0, 1]$	Theorem 5	Theorem 5
					F	F	$[0, 1]$	Theorem 6	Theorem 6
		F	F	T	F	F	0	Theorem 1	Example 9
				F	T	F	0	Theorem 1	Example 15
				F	T	T	0	Theorem 1	Example 12
F	F	T	F	T	T	T	0	Theorem 1	Example 12

Table 1: Classification of combinations *vis-à-vis* their natural densities.

The column titled "Reference" leads to the result in this paper proving the values are indeed restricted to the mentioned ones; and the column "Construction" refers to examples of theories having the possible natural densities shown.

4 Non-empty signatures

In this section we provide generalizations of the results of Section 3 to non-empty signatures. We are able to do so by considering $\mu(\mathcal{T}, \phi)$ for all formulas ϕ , rather than $\mu(\mathcal{T})$: this is due to the fact that in a non-empty signature we can have two quantifier-free satisfiable formulas with distinct densities.

Example 21. Take the theory \mathcal{T} on the signature with a unary function s , axiomatized by $\{\psi_{=2} \vee \forall x. s(x) = x\}$. In its models that do not have exactly two elements, s must be interpreted as the identity. For $\phi_1 = \neg(s(x) = x)$ and $\phi_2 = (s(x) = x)$ we have $\mu(\mathcal{T}, \phi_1) = 0$ and $\mu(\mathcal{T}, \phi_2) = 1$.

We start by generalizing items 1 and 2 in Theorem 1. As for the third item of Theorem 1, we show in Example 23 below that it cannot be generalized similarly.

Theorem 7. *The positivity of $\mu(\mathcal{T}, \phi)$ for every \mathcal{T} -satisfiable quantifier-free formula ϕ is sufficient for \mathcal{T} to be stably infinite and have the finite model property.*

In the next example we show how to use Theorem 7.

Example 22. Consider a signature Σ with only function symbols, and the Σ -theory \mathcal{T} of uninterpreted functions. For every quantifier-free formula ϕ and \mathcal{T} -interpretation \mathcal{A} that satisfies it, we can add an element a to its domain, from that it follows that $\mu(\mathcal{T}, \phi) = 1$. Using Theorem 7, we conclude that the theory of uninterpreted functions is both stably infinite, and has the finite model property.

Next, we generalize the result concerning gentleness to non empty signatures. The proof of the following result is, *mutatis mutandis*, the same as Theorem 2.

Theorem 8. *$\mu(\mathcal{T}, \phi)$ being well-defined and equal to 0 or 1 for all quantifier-free \mathcal{T} -satisfiable formulas ϕ , is a necessary condition for \mathcal{T} to be gentle.*

$$\begin{aligned}
 \psi_{\geq n}^{\bar{}} &= \exists x_1. \dots \exists x_n. \left[\bigwedge_{i=1}^n \neg (x_1, \dots, x_n) \wedge [s(x_i) = x_i] \right] \\
 \psi_{=n}^{\bar{}} &= \exists x_1. \dots \exists x_n. \left[\bigwedge_{i=1}^n \neg (x_1, \dots, x_n) \wedge [s(x_i) = x_i] \wedge \forall x. [[s(x) = x] \rightarrow \bigvee_{i=1}^n x = x_i] \right] \\
 \varsigma^{-1}(k) &= \min\{l : \varsigma(l+1) > \varsigma(k)\} \\
 Ax(\mathcal{T}_{\varsigma}^s) &= \{(\psi_{\geq k+1} \wedge \psi_{\geq \varsigma^{-1}(k+1)}^{\bar{}}) \vee \bigvee_{i=1}^{k+1} (\psi_{=i} \wedge \psi_{=\varsigma^{-1}(i)}^{\bar{}}) : k \in \mathbb{N}\}
 \end{aligned}$$

 Fig. 3: The theory $\mathcal{T}_{\varsigma}^s$.

The following theorem generalizes Theorem 3, and provides a necessary condition for smoothness and the finite model property for non-empty signatures.

Theorem 9. *Let \mathcal{T} be a theory. $\mu(\mathcal{T}, \phi)$ being well-defined and equal to 1 for all \mathcal{T} -satisfiable quantifier-free formulas ϕ is then a necessary condition for \mathcal{T} to simultaneously be smooth and have the finite model property.*

Theorem 9 can be used to show that the third item of Theorem 1 is not generalizable to non-empty signatures.

Example 23. Consider the function $\varsigma^{-1} : \mathbb{N} \rightarrow \mathbb{N}$ from [26], which is a left inverse of ς , and the theory $\mathcal{T}_{\varsigma}^s$ on the signature with only a single unary function s , from the same paper. Both ς^{-1} and $\mathcal{T}_{\varsigma}^s$ are given in Figure 3. It is smooth, has the finite model property, but is not finitely witnessable (see [23, Lemmas 71,72,73]), meaning $\mu(\mathcal{T}_{\varsigma}^s, \phi) = 1$ for all quantifier-free $\mathcal{T}_{\varsigma}^s$ -satisfiable formulas ϕ by Theorem 9. Thus, the obvious generalization of item 3 of Theorem 1 is not valid.

Theorem 9 is also useful to show, for example, that a variant of the SMT-LIB theory of bit-vectors is not smooth.

Example 24. Fix $n \in \mathbb{N}^*$, and consider the one-sorted fragment of the SMT-LIB theory $\mathbf{BV}[n]$ of bit-vectors [1] of length n , with the usual operations (but without concatenation and extraction). The domain of its interpretations has cardinality 2^n , and so it has the finite model property. By Theorem 9 this theory is not smooth, as for any quantifier-free formula ϕ one has $\mu(\mathbf{BV}[n], \phi) = 0$.

Next, we generalize Theorem 4 to non-empty signatures.

Theorem 10. *$\mu(\mathcal{T}, \phi)$ being well-defined and equal to 0 or 1 for every quantifier-free \mathcal{T} -satisfiable ϕ is necessary for \mathcal{T} to be strongly finitely witnessable.*

Example 21 shows tightness of Theorems 8 and 10: we can have a strongly finitely witnessable,⁹ gentle¹⁰ theory \mathcal{T} with two quantifier-free \mathcal{T} -satisfiable formulas that have densities 0 and 1. It also shows that the positivity in Theorem 7

⁹ The strong witness is $wit(\phi) = \phi \wedge \neg(x = y)$, for fresh variables x and y .

¹⁰ From the fact it is strongly finitely witnessable and Lemma 2.

Property	Empty case	Non-empty case
Stable Infiniteness	Theorem 1	Theorem 7
Finite Model Property	Theorem 1	Theorem 7
Gentleness	Theorem 2	Theorem 8
Smoothness	Theorem 3	Theorem 9
Strong Finite Witnessability	Theorem 4	Theorem 10
Comp. of Min. Mod. Fun.	Theorem 5	Theorem 11
Finite Witnessability	Theorem 6	Theorem 12, Example 23

Table 2: Summary of main results.

cannot hold for only some quantifier-free \mathcal{T} -satisfiable formulas ϕ , as the theory shown is not stably infinite.

The following two theorems generalize, respectively, Theorems 5 and 6. For Theorem 5, we need an alternative, non-empty version of Proposition 1. Indeed, it is not clear that if the sets $\text{Spec}(\mathcal{T}, \phi)$ are all computable, \mathcal{T} should have a computable minimal model function; the reciprocal, however, is true.

Proposition 2. *If \mathcal{T} is a theory with a computable minimal model function, then $\text{Spec}(\mathcal{T}, \phi)$ is computable for all quantifier-free \mathcal{T} -satisfiable formulas ϕ .*

Theorem 11. *If \mathcal{T} is a theory with well-defined densities $\mu(\mathcal{T}, \phi)$, for all quantifier-free \mathcal{T} -satisfiable formulas ϕ , the fact that all $\mu(\mathcal{T}, \phi)$ are computable is a necessary condition for \mathcal{T} to have a computable minimal model function. Furthermore, for every computable number $0 \leq r \leq 1$, there is a theory \mathcal{T} that has a computable minimal model function and a quantifier-free formula ϕ with $\mu(\mathcal{T}, \phi) = r$.*

Theorem 12. *If \mathcal{T} is a theory, and ϕ a quantifier-free \mathcal{T} -satisfiable formula, nothing can be said about $\mu(\mathcal{T}, \phi)$ if \mathcal{T} is only known to be finitely witnessable; that is, for every computable number $0 \leq r \leq 1$, there exists a theory \mathcal{T} , that is finitely witnessable, and a quantifier-free formula ϕ with $\mu(\mathcal{T}, \phi) = r$.*

5 Related work and conclusion

We have studied connections between densities and model-theoretic properties. Table 2 summarizes our main results. For each property, we refer to the theorems that characterize its possible densities, both for empty and non-empty signatures.

We conclude by reviewing related work and sketching the next steps.

5.1 Related work

0-1-laws and densities. Studies on spectra and densities go back as far as [4,9,7]. While we consider only models of a theory, these results, including the famous 0-1 laws, concern *random models*, that is, any models. 0-1 laws remain powerful

for theories with finite axiomatizations (as we can represent their axiomatizations using a conjunction), but here we consider also infinite axiomatizations. Later studies, such as [5,3] considered densities with respect to a theory, or even a (sufficiently well-behaved) class of models, but have not considered properties associated with theory combination. We focus on theories, and on the relationship between their combination properties and the behavior of their density.

Descriptive complexity. Note that we use slightly different definitions for the spectrum of a theory than those found in descriptive complexity [11]: although our definition of $\text{Spec}(\mathcal{T}, \phi)$ is the usual one for the spectrum of a formula relative to a theory, the spectrum of a theory \mathcal{T} is more commonly understood as the map from cardinals to cardinals which, given κ , returns the number of non-isomorphic models of \mathcal{T} of cardinality κ . But for the case of finite cardinalities in the empty signature, there is this map would return either 0 or 1. Then, our definition coincides with taking the pre-image of 1 in the more standard definition.

Theory combination properties. The current paper deals with, among other topics, Boolean combinations of theory combination properties (especially in Table 1), something comprehensively researched in [25,26,24]. While those papers study the combinations of properties per se, here we focus on establishing these properties (or lack of) through the analysis of a their density.

5.2 Future work: many-sorted densities

In this paper we only considered one-sorted theories, even though many-sorted theories are commonly used in SMT. The main reason for that is that densities for many-sorted theories would be defined on tuples rather than on numbers (i.e. on the cardinalities of the domains rather than on that of the single domain), and it is unclear how this generalization would materialize. We leave this investigation for future work, and briefly describe concrete options for such a generalization.

What makes the natural density so natural is the fact that it calculates the ratio of the number of elements in a set A to the number of elements in \mathbb{N} by doing that for numbers under a bound, and then letting said bound go to infinite. But there is no single way of doing that in \mathbb{N}^m , so we are forced to make a choice. Once fixed a bound n , do we, for example:

(i) bound all coordinates simultaneously by n (i.e., $\mu(A) = \lim_{n \rightarrow \infty} |A_n|/n$ for $A_n = A \cap [n]^m$)? (ii) bound the distance of a tuple to the origin by n (i.e., $A_n = A \cap B_d(n)$, where $B_d(n) = \{\mathbf{p} \in \mathbb{N}^m : d(\mathbf{0}, \mathbf{p}) \leq n\}$, for $\mathbf{0}$ the origin)? (iii) If so, what metric do we use to calculate the distance? Do we use the taxicab distance, where $d_1(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^n |p_i - q_i|$, or the generalized euclidean distances $d_m(\mathbf{p}, \mathbf{q}) = (\sum_{i=1}^n (p_i - q_i)^m)^{1/m}$, or something entirely different?

There is a plurality of "natural densities" to explore. Even more, while some generalizations will characterize properties w.r.t. the entire set of sorts $\{\sigma_1, \dots, \sigma_n\}$, others will characterize them with respect to some subset of sorts, while others will offer no characterization whatsoever.

All of this is left to a future work, but we expect that the results from the current paper will still be useful for many-sorted logic, as many of the potential many-sorted densities would rely on the separate projections to each sort.

References

1. Clark Barrett, Pascal Fontaine, and Cesare Tinelli. The SMT-LIB Standard: Version 2.6. Technical report, Department of Computer Science, The University of Iowa, 2017. Available at <http://smt-lib.org>.
2. Clark Barrett, Roberto Sebastiani, Sanjit Seshia, and Cesare Tinelli. Satisfiability modulo theories. In Armin Biere, Marijn J. H. Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability, Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, chapter 33, pages 825–885. IOS Press, February 2021.
3. J. Bell and Burris S. Compton’s method for proving logical limit laws. *Contemporary Mathematics*, 558:97–128, 2011.
4. Rudolf Carnap. *Logical Foundations of Probability*. Chicago University of Chicago Press, Chicago, 1950.
5. Kevin J. Compton, C. Ward Henson, and Saharon Shelah. Nonconvergence, undecidability, and intractability in asymptotic problems. *Annals of Pure and Applied Logic*, 36:207–224, 1987.
6. Guilherme Vicentin de Toledo and Yoni Zohar. Combining combination properties: Minimal models, 2024.
7. Ronald Fagin. Probabilities on finite models. *The Journal of Symbolic Logic*, 41(1):50–58, 1976.
8. Pascal Fontaine. Combinations of theories for decidable fragments of first-order logic. In Silvio Ghilardi and Roberto Sebastiani, editors, *Frontiers of Combining Systems*, pages 263–278, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
9. Y. V. Glebskii, D. I. Kogan, M. I. Liogon’kii, and V. A. Talanov. Volume and fraction of satisfiability of formulas of the lower predicate calculus. *Kibernetika (Kiev)*, 5:17–27, 1969.
10. Wilfrid Hodges. *Model Theory*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1993.
11. Neil Immerman. *Descriptive Complexity*. Texts in Computer Science. Springer, New York, NY, 1999 edition, November 1998.
12. Dejan Jovanovic and Clark W. Barrett. Polite theories revisited. In *LPAR (Yogyakarta)*, volume 6397 of *Lecture Notes in Computer Science*, pages 402–416. Springer, 2010.
13. Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Trans. Program. Lang. Syst.*, 1(2):245–257, October 1979.
14. T. Radó. On non-computable functions. *The Bell System Technical Journal*, 41(3):877–884, 1962.
15. Silvio Ranise, Christophe Ringeissen, and Calogero G. Zarba. Combining data structures with nonstably infinite theories using many-sorted logic. In Bernard Gramlich, editor, *5th International Workshop on Frontiers of Combining Systems - FroCoS’05*, volume 3717 of *Lecture Notes in Artificial Intelligence*, pages 48–64, Vienna, September 2005. Springer.
16. Silvio Ranise, Christophe Ringeissen, and Calogero G. Zarba. Combining data structures with nonstably infinite theories using many-sorted logic. In Bernard Gramlich, editor, *5th International Workshop on Frontiers of Combining Systems - FroCoS’05*, volume 3717 of *Lecture Notes in Artificial Intelligence*, pages 48–64, Vienna, September 2005. Springer.
17. Ying Sheng, Yoni Zohar, Christophe Ringeissen, Jane Lange, Pascal Fontaine, and Clark Barrett. Polite combination of algebraic datatypes. *Journal of Automated Reasoning*, 66(3):331–355, Aug 2022.

18. Ying Sheng, Yoni Zohar, Christophe Ringeissen, Andrew Reynolds, Clark Barrett, and Cesare Tinelli. Politeness and stable infiniteness: Stronger together. In André Platzer and Geoff Sutcliffe, editors, *Automated Deduction – CADE 28*, pages 148–165, Cham, 2021. Springer International Publishing.
19. E. Szemerédi. On sets of integers containing k elements in arithmetic progression. *Acta Arithmetica*, 27:199–245, 1975.
20. Gérald Tenenbaum. *Introduction to analytic and probabilistic number theory. Transl. from the 2nd French ed. by C.B.Thomas*, volume 46 of *Camb. Stud. Adv. Math.* Cambridge: Cambridge Univ. Press, 1995.
21. Cesare Tinelli and Calogero G. Zarba. Combining nonstably infinite theories. *J. Autom. Reason.*, 34(3):209–238, 2005.
22. Cesare Tinelli and Calogero G. Zarba. Combining nonstably infinite theories. *Journal of Automated Reasoning*, 34(3):209–238, April 2005.
23. Guilherme Toledo, Yoni Zohar, and Clark Barrett. Combining finite combination properties: Finite models and busy beavers, 2023.
24. Guilherme V. Toledo and Yoni Zohar. Combining combination properties: Minimal models. In Nikolaj Björner, Marijn Heule, and Andrei Voronkov, editors, *Proceedings of 25th Conference on Logic for Programming, Artificial Intelligence and Reasoning*, volume 100 of *EPiC Series in Computing*, pages 19–35. EasyChair, 2024.
25. Guilherme V. Toledo, Yoni Zohar, and Clark Barrett. Combining combination properties: An analysis of stable infiniteness, convexity, and politeness. In Brigitte Pientka and Cesare Tinelli, editors, *Automated Deduction – CADE 29*, pages 522–541, Cham, 2023. Springer Nature Switzerland.
26. Guilherme V. Toledo, Yoni Zohar, and Clark Barrett. Combining finite combination properties: Finite models and busy beavers. In Uli Sattler and Martin Suda, editors, *Frontiers of Combining Systems*, pages 159–175, Cham, 2023. Springer Nature Switzerland.
27. A. M. Turing. On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, s2-42(1):230–265, 1937.

A Useful Theorems

We briefly recall the Löwenheim-Skolem and compactness theorems (see, e.g., [10]).

Theorem 13. *Given a first-order signature Σ , if a set Γ of Σ -formulas is satisfiable by a Σ -interpretation with $|\text{dom}(\mathcal{A})| \geq \aleph_0$, then it is satisfied by an interpretation \mathcal{B} with $|\text{dom}(\mathcal{B})| = \aleph_0$.*

Theorem 14. *Given a first-order signature Σ , a set Γ of Σ -formulas is satisfiable if, and only if, every finite subset $\Gamma_0 \subseteq \Gamma$ is satisfiable.*

B Proof of Theorem 1

Lemma 3. *Let \mathcal{T} be a Σ_1 -theory. If ϕ is a quantifier-free Σ_1 -formula, and \mathcal{A} and \mathcal{B} are \mathcal{T} -interpretations such that \mathcal{B} satisfies ϕ , and $|\text{dom}(\mathcal{A})| \geq |\text{vars}(\phi)^{\mathcal{B}}|$, then there is a \mathcal{T} -interpretation \mathcal{A}' , differing from \mathcal{A} at most on the value assigned to $\text{vars}(\phi)$, that satisfies ϕ .*

If \mathcal{A} and \mathcal{B} are \mathcal{T} -interpretations with $|\text{dom}(\mathcal{A})| = |\text{dom}(\mathcal{B})|$, then there is a \mathcal{T} -interpretation \mathcal{A}' , differing from \mathcal{A} at most on the values assigned to variables, such that \mathcal{A}' and \mathcal{B} satisfy exactly the same, not necessarily quantifier-free, formulas.

Proof. Take an injective function $f : \text{vars}(\phi)^{\mathcal{B}} \rightarrow \text{dom}(\mathcal{A})$ (bijective function $f : \text{dom}(\mathcal{B}) \rightarrow \text{dom}(\mathcal{A})$ if $|\text{dom}(\mathcal{A})| = |\text{dom}(\mathcal{B})|$), and define an interpretation \mathcal{A}' such that: $\text{dom}(\mathcal{A}') = \text{dom}(\mathcal{A})$; $x^{\mathcal{A}'} = f(x^{\mathcal{B}})$ if $x^{\mathcal{B}}$ is in the domain of f ; and $x^{\mathcal{A}'} = x^{\mathcal{A}}$ otherwise. We now prove by structural induction that \mathcal{A}' satisfies ϕ ; in the case where $|\text{dom}(\mathcal{A})| = |\text{dom}(\mathcal{B})|$ and f is bijective, we must perform the proof simultaneously over all interpretations with the same cardinality as that of \mathcal{A} and \mathcal{B} .

1. If ψ is an atomic subformula of ϕ , since our signature is empty it must equal $x = y$, for some variables x and y ; then \mathcal{B} satisfies ψ iff $x^{\mathcal{B}} = y^{\mathcal{B}}$ and, since $x^{\mathcal{A}'} = f(x^{\mathcal{B}})$ and $y^{\mathcal{A}'} = f(y^{\mathcal{B}})$, that happens iff \mathcal{A}' satisfies ψ .
2. So, suppose \mathcal{B} satisfies the subformulas ψ and ψ_i (for $i \in \{1, 2\}$) of ϕ iff \mathcal{A}' satisfies the same formulas.
 - (a) We have that \mathcal{B} satisfies $\neg\psi$ iff it does not satisfy ψ , what in turn happens iff \mathcal{A}' does not satisfy ψ , and so \mathcal{B} satisfies $\neg\psi$ iff \mathcal{A}' satisfies $\neg\psi$.
 - (b) Analogously, \mathcal{B} satisfies $\psi_1 \vee \psi_2$ iff it satisfies ψ_1 or ψ_2 , what happens iff \mathcal{A}' satisfies ψ_1 or ψ_2 , what in turn happens iff \mathcal{A}' satisfies $\psi_1 \vee \psi_2$.
 - (c) For the case where $|\text{dom}(\mathcal{A})| = |\text{dom}(\mathcal{B})|$, and ϕ is not necessarily quantifier-free, suppose \mathcal{B} satisfies $\exists x. \psi$, and so there exists an interpretation \mathcal{B}_* , differing from \mathcal{B} at most on the value assigned to x , that satisfies ψ . We then define an interpretation \mathcal{A}'_* differing from \mathcal{A}' at most on the value assigned to x , where $x^{\mathcal{A}'_*} = f(x^{\mathcal{B}_*})$, and from the strengthened form of

the induction hypothesis we have \mathcal{A}'_* satisfies ψ , and thus \mathcal{A}' satisfies $\exists x. \psi$.

Reciprocally, if \mathcal{A}' satisfies $\exists x. \psi$, there is an interpretation \mathcal{A}'_* differing from \mathcal{A}' at most on x , and we can then define an interpretation \mathcal{B}_* , differing of \mathcal{B} at most on x where $x^{\mathcal{B}_*} = f^{-1}(x^{\mathcal{A}'_*})$, so that \mathcal{B}_* satisfies ψ and thus \mathcal{B} satisfies $\exists x. \psi$.

- (d) The cases of $\psi_1 \wedge \psi_2$ and $\psi_1 \rightarrow \psi_2$ (and $\forall x. \psi$) can be derived from those, and thus we are done.

That \mathcal{A}' is still a \mathcal{T} -interpretation follows from the fact that to obtain \mathcal{A}' we only (at most) changed the value assigned by \mathcal{A} to some variables.

Lemma 4. *If \mathcal{T} is a Σ_1 -theory with an infinite interpretation, it is stably infinite.*

Proof. Let \mathcal{A} be an infinite \mathcal{T} -interpretation, ϕ a quantifier-free formula, and \mathcal{B} a \mathcal{T} -interpretation that satisfies ϕ : since $|\text{dom}(\mathcal{A})| \geq |\text{vars}(\phi)^{\mathcal{B}}|$ we can thus apply Lemma 3 to obtain an infinite \mathcal{T} -interpretation \mathcal{A}' that satisfies ϕ .

Lemma 5. *If \mathcal{T} is a Σ_1 -theory that is not stably infinite, then there exists $M \in \mathbb{N}$ such that $|\text{dom}(\mathcal{A})| \leq M$ for all \mathcal{T} -interpretations \mathcal{A} .*

Proof. Suppose that, for every $n \in \mathbb{N}$, there exists a \mathcal{T} -interpretation \mathcal{A}_n with $|\text{dom}(\mathcal{A}_n)| \geq n$. If we define $\Gamma = Ax(\mathcal{T}) \cup \{\psi_{\geq m} : m \in \mathbb{N}\}$, for any finite subset $\Gamma_0 \subseteq \Gamma$, take the maximum n of m such that $\psi_{\geq m}$ occurs in Γ_0 . Of course, \mathcal{A}_n then satisfies Γ_0 ; it satisfies $\Gamma_0 \cap Ax(\mathcal{T})$ as it is a \mathcal{T} -interpretation; and it satisfies any $\psi_{\geq m}$ in Γ_0 , since it satisfies $\psi_{\geq n}$, and $\psi_{\geq n} \rightarrow \psi_{\geq m}$ is valid for all $m \leq n$. From Theorem 14, it follows that Γ is satisfiable, meaning there is an infinite \mathcal{T} -interpretation \mathcal{A} . The result then follows from Lemma 4.

Lemma 6. *If \mathcal{T} is a Σ_1 -theory that does not have the finite model property, then there exists $M \in \mathbb{N}$ such that $|\text{dom}(\mathcal{A})| \leq M$ for all finite \mathcal{T} -interpretations \mathcal{A} .*

Proof. If \mathcal{T} has only finite interpretations it already has the finite model property, so take a quantifier-free formula ϕ that is satisfied by no finite \mathcal{T} -interpretations, and an infinite \mathcal{T} -interpretation \mathcal{A} that satisfies ϕ . Suppose that there exists, for each $n \in \mathbb{N}$, a finite \mathcal{T} -interpretation \mathcal{A}_n with $|\text{dom}(\mathcal{A}_n)| \geq n$: then, for $n = |\text{vars}(\phi)^{\mathcal{A}}|$, Lemma 3 guarantees that by changing the value assigned to variables in \mathcal{A}_n we can make it satisfy ϕ , contradicting our hypothesis.

Theorem 1. *If \mathcal{T} is a Σ_1 -theory with a well-defined natural density, then the positivity of $\mu(\mathcal{T})$ is sufficient for \mathcal{T} to: 1. be stably infinite; 2. have the finite model property; and 3. be finitely witnessable.*

Proof. The proofs for the first and second items are the same: assuming \mathcal{T} is not stably infinite, respectively does not have the finite model property, we have from Lemma 5, respectively Lemma 6, that there is an $M \in \mathbb{N}$ such that, for all finite \mathcal{T} -interpretations \mathcal{A} , $|\text{dom}(\mathcal{A})| \leq M$. This means that $\text{Spec}_n(\mathcal{T}) \subseteq [1, M]$ for all $n \geq$

1, and so $|Spec_n(\mathcal{T})| \leq M$ in that case: therefore, $\mu(\mathcal{T}) = \lim_{n \rightarrow \infty} |Spec_n(\mathcal{T})|/n \leq \lim_{n \rightarrow \infty} M/n = 0$. Thus $\mu(\mathcal{T}) = 0$, and by the contrapositive we have the desired result.

For the third item, from the hypothesis in the theorem's statement, $\mu(\mathcal{T})$ is well-defined; suppose, in addition, that $\mu(\mathcal{T}) > 0$. Let $\{a_n\}_{n \in \mathbb{N}}$ be the increasing sequence of the elements of $Spec(\mathcal{T})$ (if this set were to be finite, it would be bounded by some $M \in \mathbb{N}$, and we would have $\mu(\mathcal{T}) \leq \lim_{n \rightarrow \infty} M/n = 0$). There are then two cases to consider: (1) there is a computable sequence $\{b_n\}_{n \in \mathbb{N}}$ such that $a_n \leq b_n$ for all $n \in \mathbb{N}$; and (2) for all computable sequences $\{b_n\}_{n \in \mathbb{N}}$ and all $N_0 \in \mathbb{N}$, there is an $n \geq N_0$ such that $a_n > b_n$. Notice that the negation of (2) implies (1), meaning these are indeed the only cases we need to consider: if there is a computable sequence $\{b_n\}_{n \in \mathbb{N}}$ and an integer N_0 such that $b_n \geq a_n$ for all $n \geq N_0$, we define the computable sequence $\{c_n\}_{n \in \mathbb{N}}$ by making $c_n = b_{N_0}$ for $n \leq N_0$, and $c_n = b_n$ for $n \geq N_0$, which satisfies the conditions in case (1). We will prove that case (1) implies \mathcal{T} is finitely witnessable; and that (2) implies $\mu(\mathcal{T}) = 0$, contradicting our assumption.

1. Let x_i be fresh variables, and define $N(\phi)$, for a quantifier-free formula ϕ , as b_n , for $n = |vars(\phi)|$. We define $wit(\phi) = \phi \wedge \bigwedge_{i=1}^{N(\phi)} x_i = x_i$: this is obviously a function from quantifier-free formulas into themselves, and is computable given b_n , and thus $N(\phi)$, are computable. Furthermore, ϕ and $\exists \vec{x}. wit(\phi)$ are clearly \mathcal{T} -equivalent, given $wit(\phi)$ is the conjunction of ϕ and a tautology and thus already equivalent to ϕ .

Finally, let ϕ be a quantifier-free formula, and \mathcal{A} a \mathcal{T} -interpretation that satisfies $wit(\phi)$ (and thus ϕ , since the two are equivalent). Take again $n = |vars(\phi)|$, a set A with $a_n - |vars(\phi)^{\mathcal{A}}|$ elements (a non-negative quantity, since $\{a_n\}_{n \in \mathbb{N}}$ is increasing and therefore $a_n \geq n = |vars(\phi)^{\mathcal{A}}|$), disjoint from $dom(\mathcal{A})$, and define an interpretation \mathcal{B} by making: $dom(\mathcal{B}) = vars(\phi)^{\mathcal{A}} \cup A$ (so $|dom(\mathcal{B})| = a_n$, making \mathcal{B} a \mathcal{T} -interpretation, as there is only one \mathcal{T} -interpretation of a given size up to isomorphism by Lemma 3); $x^{\mathcal{B}} = x^{\mathcal{A}}$ for all $x \in vars(\phi)$ (so \mathcal{B} satisfies ϕ);

$$x_i \in \{x_i : 1 \leq i \leq N(\phi)\} \mapsto x_i^{\mathcal{B}} \in dom(\mathcal{B})$$

a surjective map (what is possible, as $N(\phi) = b_n \geq a_n$); and $x^{\mathcal{B}}$ defined arbitrarily for all other variables. Because $\{x_i : 1 \leq i \leq N(\phi)\}$ is contained in $vars(wit(\phi))$, and $\{x_i : 1 \leq i \leq N(\phi)\}^{\mathcal{B}} = dom(\mathcal{B})$, we have $dom(\mathcal{B}) = vars(wit(\phi))^{\mathcal{B}}$ (since obviously the former set at least contains the latter): it then follows that wit is a witness.

2. Consider the computable sequence $b_n = n^2$: take m_0 as the minimum of the positive integers m such that $a_m > b_m$ and, inductively, m_{n+1} as the minimum of the integers $m > m_n$ such that $a_m > b_m$; this way, $a_{m_n} > b_{m_n} = m_n^2$ for all $n \in \mathbb{N}$. Of course, for all $n \in \mathbb{N}$ one finds $|Spec_{a_{m_n}}(\mathcal{T})| = m_n$, and $[[1, a_{m_n}]] = a_{m_n} > m_n^2$, thus $|Spec_{a_{m_n}}(\mathcal{T})|/[[1, a_{m_n}]] < 1/m_n$. This means that a subsequence of $|Spec_n(\mathcal{T})|/[[1, n]]$ (which converges to $\mu(\mathcal{T})$, as this value was assumed in the theorem's statement to be well-defined) converges to 0, meaning $\mu(\mathcal{T}) = 0$.

C Proof of Theorem 2

Theorem 2. *If \mathcal{T} is a Σ_1 -theory, then $\mu(\mathcal{T})$ being well-defined and equal to 0 or 1 is a necessary condition for \mathcal{T} to be gentle.*

Proof. If $\text{Spec}(\mathcal{T})$ is finite, there is an M such that $|\text{Spec}_n(\mathcal{T})| \leq |\text{Spec}(\mathcal{T})| \leq M$, and therefore $\mu(\mathcal{T}) = \lim_{n \rightarrow \infty} M/n = 0$. If $\text{Spec}(\mathcal{T})$ is co-finite instead, there is an M such that for all $n > M$ one has $|\text{Spec}_n(\mathcal{T})| \geq n - M$, and so $\mu(\mathcal{T}) \geq \lim_{n \rightarrow \infty} (n - M)/n = 1$.

D Proof of Lemma 1

Lemma 1. *If \mathcal{T} is gentle, then \mathcal{T} has a computable minimal model function and the finite model property, and therefore is finitely witnessable as well.*

Proof. Suppose \mathcal{T} is gentle. From the definition of gentleness we get there must exist a finite number in the spectrum of each \mathcal{T} -satisfiable quantifier-free formula ϕ ; this means there is a finite \mathcal{T} -interpretation that satisfies ϕ , so \mathcal{T} indeed has the finite model property.

Given a quantifier-free formula ϕ , we can decide whether $\text{Spec}(\mathcal{T}, \phi)$ is finite, and then computably calculate $\max(\text{Spec}(\mathcal{T}, \phi))$, or co-finite, and in this case calculate $\max(\mathbb{N} \setminus \text{Spec}(\mathcal{T}, \phi))$. Using that $\text{Spec}(\mathcal{T}, \phi)$ is computable, we can also algorithmically obtain $\min(\text{Spec}(\mathcal{T}, \phi))$: if $\text{Spec}(\mathcal{T}, \phi)$ is finite, we test which $n \leq \max(\text{Spec}(\mathcal{T}, \phi))$ are in $\text{Spec}(\mathcal{T}, \phi)$ and take their minimum (if $\max(\text{Spec}(\mathcal{T}, \phi))$ is 0, we just set $\min(\text{Spec}(\mathcal{T}, \phi))$ to 0 as well, for simplicity); if $\text{Spec}(\mathcal{T}, \phi)$ is co-finite, we test which $n \leq \max(\mathbb{N} \setminus \text{Spec}(\mathcal{T}, \phi)) + 1$ are in $\text{Spec}(\mathcal{T}, \phi)$, and again take their minimum.

We state that $\mathbf{minmod}_{\mathcal{T}}(\phi) = \min(\text{Spec}(\mathcal{T}, \phi))$ is a computable minimal model function, being certainly computable. Assume then ϕ is \mathcal{T} -satisfiable, and so $\min(\text{Spec}(\mathcal{T}, \phi)) > 0$. For the first direction, suppose there is a \mathcal{T} -interpretation \mathcal{A} that satisfies ϕ with $|\text{dom}(\mathcal{A})| < \min(\text{Spec}(\mathcal{T}, \phi))$; since $|\text{dom}(\mathcal{A})|$ is in $\text{Spec}(\mathcal{T}, \phi)$, this contradicts the fact $\min(\text{Spec}(\mathcal{T}, \phi))$ is the minimum element of that set. Now, since $\min(\text{Spec}(\mathcal{T}, \phi))$ is in $\text{Spec}(\mathcal{T}, \phi)$, there is a \mathcal{T} -interpretation \mathcal{A} that satisfies ϕ with $|\text{dom}(\mathcal{A})| = \min(\text{Spec}(\mathcal{T}, \phi))$, and thus we have proved $\min(\text{Spec}(\mathcal{T}, \phi))$ is indeed a minimal model function.

Finite witnessability follows from [6, Theorem 4].

E Proof of Lemma 2

Lemma 7. *A Σ_1 -theory is always decidable.*

Proof. Take a Σ_1 -theory \mathcal{T} . We divide the proof in two cases: for each of them we find an algorithm to decide whether a quantifier-free formula ϕ is \mathcal{T} -satisfiable; notice, however, that there is no decision method with \mathcal{T} as input that returns which of the described algorithms is the correct one.

1. Suppose \mathcal{T} has a maximum model \mathcal{A} , with $|dom(\mathcal{A})| = M$; we state that ϕ is \mathcal{T} -satisfiable if and only if there is an interpretation \mathcal{B} in equational logic that satisfies ϕ with $|dom(\mathcal{B})| \leq M$, meaning we are done since equational logic is decidable. Indeed, if \mathcal{B} is an interpretation in equational logic that satisfies ϕ with $|dom(\mathcal{B})| \leq M$, we can change the values assigned to the variables in ϕ by \mathcal{A} in order to obtain a \mathcal{T} -interpretation \mathcal{A}' that satisfies ϕ (see Lemma 3). And, if \mathcal{C} is a \mathcal{T} -interpretation that satisfies ϕ , given the facts that $|dom(\mathcal{C})| \leq M$ and \mathcal{C} is an interpretation in equational logic as well, we are done.
2. If \mathcal{T} doesn't have a maximum model, we state that all quantifier-free formulas satisfiable in equational logic are also \mathcal{T} -satisfiable. Indeed, if ϕ is satisfied by the interpretation \mathcal{B} of equational logic, we know there is a \mathcal{T} -interpretation \mathcal{A} with $|dom(\mathcal{A})| \geq |\mathcal{B}|$, and again by Lemma 3 we are done.

Lemma 2. *Let \mathcal{T} be a Σ_1 -theory: if \mathcal{T} is not stably infinite, or if it is strongly finitely witnessable, then \mathcal{T} is gentle.*

Proof. If \mathcal{T} is not stably infinite, by Lemma 5 \mathcal{T} has a maximum interpretation \mathcal{A} , say with $|dom(\mathcal{A})| = M$. $Spec(\mathcal{T}, \phi)$ is then always finite (see Lemma 3); $\max(Spec(\mathcal{T}, \phi)) = M$ if ϕ is \mathcal{T} -satisfiable (we use Lemma 7), and 0 otherwise; and a cardinality is in $Spec(\mathcal{T}, \phi)$ if it is both in $Spec(\mathcal{T})$ and larger than or equal to the minimum cardinality of an interpretation in equational logic that satisfies ϕ ($Spec(\mathcal{T})$ is a finite list, so it can be hardcoded). In summary, \mathcal{T} is gentle.

If \mathcal{T} is strongly finitely witnessable but not stably infinite the previous reasoning already implies \mathcal{T} is gentle, so assume \mathcal{T} is strongly finitely witnessable and stably infinite. \mathcal{T} is then also smooth by [25, Theorem 7], and so $Spec(\mathcal{T})$ is co-finite: let M be the maximum of $\mathbb{N} \setminus Spec(\mathcal{T})$. $Spec(\mathcal{T}, \phi)$ is then always co-finite, by Lemma 3 and the fact that strong finite witnessability implies the finite model property; $\max(\mathbb{N} \setminus Spec(\mathcal{T}, \phi)) = 0$ if ϕ is not \mathcal{T} -satisfiable (what can be decided from Lemma 7), and otherwise we make N the cardinality of the smallest interpretation in equational logic that satisfies ϕ , and then $\max(\mathbb{N} \setminus Spec(\mathcal{T}, \phi)) = \max\{M, N - 1\}$, both computable quantities; and finally, $n \in Spec(\mathcal{T}, \phi)$ iff $n > \max\{M, N - 1\}$, meaning $Spec(\mathcal{T}, \phi)$ is computable.

F Proof of Theorem 3

Theorem 3. *If \mathcal{T} is a Σ_1 -theory, $\mu(\mathcal{T})$ being well-defined and equal to 1 is necessary for \mathcal{T} to simultaneously admit smoothness and the finite model property.*

Proof. Let ϕ be any tautology, and from the fact that \mathcal{T} has the finite model property there exists a finite \mathcal{T} -interpretation \mathcal{B} that satisfies ϕ : from the fact \mathcal{T} is smooth it follows that for all $n \geq M = |dom(\mathcal{B})|$ there exists a \mathcal{T} -interpretation \mathcal{A}_n with $|dom(\mathcal{A}_n)| = n$, so $|Spec_n(\mathcal{T})| = |Spec_M(\mathcal{T})| + (n - M)$ for all $n \geq M$. Thus we have $\mu(\mathcal{T}) = \lim_{n \rightarrow \infty} |Spec_n(\mathcal{T})|/n \geq \lim_{m \rightarrow \infty} m/(m + M) = 1$; therefore $\mu(\mathcal{T}) = 1$, as we wished to prove.

G Proof of Theorem 4

Theorem 4. *If \mathcal{T} is a Σ_1 -theory, then $\mu(\mathcal{T})$ being well-defined and equal to 0 or 1 is a necessary condition for \mathcal{T} to be strongly finitely witnessable.*

Proof. From [26, Theorem 2] we know strong finite witnessability implies the final model property. From [25, Theorem 7], we know that if \mathcal{T} is stably infinite, then it is smooth, and so $\mu(\mathcal{T}) = 1$ from Theorem 3; if \mathcal{T} is not stably infinite, Theorem 1 guarantees that $\mu(\mathcal{T}) = 0$.

H Proof of Proposition 1

Proposition 1. *\mathcal{T} is a Σ_1 -theory with a computable minimal model function if, and only if, $\text{Spec}(\mathcal{T})$ is computable.*

Proof. For the right-to-left direction, suppose the spectrum $\text{Spec}(\mathcal{T})$ is computable: if \mathcal{T} does not have the finite model property, by Lemma 6 one gets that there is a natural number M that bounds the cardinalities of all finite models of \mathcal{T} . Thus, a minimal model function can simply be obtained by checking for all models up to M whether they satisfy the input formula, and stopping at the first one (recall that since the signature is empty a model is uniquely determined by its cardinality). If \mathcal{T} is empty (has no models) then any computable function is a minimal model function, as no formula is \mathcal{T} -satisfiable.

We may then assume that \mathcal{T} both has the finite model property and is not contradictory (so $\text{Spec}(\mathcal{T}) \neq \emptyset$) and take, if ϕ is not a contradiction (what can be determined algorithmically in equational logic): the set of variables V in ϕ ; the set $\text{Eq}(V)$ of all the equivalences E on V ; $M = \sup(\text{Spec}(\mathcal{T}))$ (which equals \aleph_0 if $\text{Spec}(\mathcal{T})$ is infinite, and is otherwise always a positive natural number since $\text{Spec}(\mathcal{T}) \neq \emptyset$), $M(\phi) = \min\{|V/E| : \delta_V^E \text{ implies } \phi\}$ (a quantity that is computable since the problem may be reduced to one in equational logic), and we state that

$$\mathbf{minmod}_{\mathcal{T}}(\phi) = \begin{cases} \aleph_0 & \text{if } \phi \text{ is contradictory,} \\ & \text{or } M(\phi) > M; \\ \min\{M(\phi) \leq k \leq M : k \in \text{Spec}(\mathcal{T})\} & \text{otherwise} \end{cases}$$

is a computable minimal model function for \mathcal{T} . Indeed, if ϕ is a contradiction, what the decision procedure for equational logic can tell us, $\mathbf{minmod}_{\mathcal{T}}(\phi) = \aleph_0$; if ϕ is not a contradiction, but it is a \mathcal{T} -contradiction, then $M(\phi) > M$ (indeed, if $M \geq M(\phi)$ there is a \mathcal{T} -interpretation \mathcal{A} with $|\text{dom}(\mathcal{A})| = M$, by appealing to Theorem 14 if necessary; and by Lemma 3 we get that ϕ is \mathcal{T} -satisfiable), and again $\mathbf{minmod}_{\mathcal{T}}(\phi) = \aleph_0$. Finally, in the case that ϕ is \mathcal{T} -satisfiable, since \mathcal{T} has the finite model property by assumption, there is a finite \mathcal{T} -interpretation \mathcal{A} that satisfies ϕ and, of course, $|\text{dom}(\mathcal{A})| \geq M(\phi)$ (because otherwise \mathcal{A} induces an equivalence δ_V^E that implies ϕ with $|V/E| < M(\phi)$): therefore the set $\{M(\phi) \leq k \leq M : k \in \text{Spec}(\mathcal{T})\}$ is not empty (it contains at least $|\text{dom}(\mathcal{A})|$), and finding its minimum is easy.

So assume ϕ is \mathcal{T} -satisfiable, and take: an $E \in Eq(V)$ with $|V/E| = M(\phi)$ (and necessarily $M(\phi) < \aleph_0$); the minimum n of $M(\phi) \leq k \leq M$ such that $k \in Spec(\mathcal{T})$; and a set X with $|X| = n - M(\phi)$ disjoint from V/E . We then create an interpretation \mathcal{B} with: $dom(\mathcal{B}) = (V/E) \cup X$; $x^{\mathcal{B}} = [x]$ for all $x \in V$, where $[x]$ is the equivalence class under E represented by x ; and $x^{\mathcal{B}}$ defined arbitrarily for variables x not in V . Because $|dom(\mathcal{B})| = M(\phi) + (n - M(\phi)) = n$, we know thanks to Lemma 3 that \mathcal{B} is a \mathcal{T} -interpretation; furthermore, \mathcal{B} satisfies δ_V^E by definition, thus satisfying ϕ , so that we indeed have a \mathcal{T} -interpretation that satisfies ϕ with cardinality n .

Assume now that there is a \mathcal{T} -interpretation \mathcal{C} that satisfies ϕ with $|dom(\mathcal{C})| = m < n$, and let F be the equivalence induced by \mathcal{C} on V , implying that \mathcal{C} satisfies δ_V^F . We then have that $M(\phi) \leq |V/F| \leq m \leq M$, and that $m \in Spec(\mathcal{T})$, despite the fact that $m < n$ and n should be the minimum such element. So \mathcal{T} indeed has a computable minimal model function.

For the reciprocal, the left-to-right direction, suppose that \mathcal{T} has a computable minimal model function $\mathbf{minmod}_{\mathcal{T}}$, and we shall consider two cases: one where $Spec(\mathcal{T})$ is finite, and one where it is infinite. If $Spec(\mathcal{T})$ is finite we have nothing left to do, as it is enough to simply hardcode these values into an algorithm that decides whether an element is in $Spec(\mathcal{T})$. If $Spec(\mathcal{T})$ is instead infinite, the formulas $\#(x_1, \dots, x_n)$ are \mathcal{T} -satisfiable for all $n \in \mathbb{N}$: indeed, since $Spec(\mathcal{T})$ is infinite, we can always find $m \in Spec(\mathcal{T})$ such that $m \geq n$, and thus there exists a \mathcal{T} -interpretation \mathcal{A} with $|dom(\mathcal{A})| = m$; by changing at most the values assigned to the variables x_1, \dots, x_n (assumed fresh), we get a \mathcal{T} -interpretation that satisfies $\#(x_1, \dots, x_n)$. We then define $f(0) = \mathbf{minmod}_{\mathcal{T}}(x = x)$ and, assuming $f(m)$ defined, $f(m+1) = \mathbf{minmod}_{\mathcal{T}}(\#(x_1, \dots, x_{f(m)+1}))$; we state that $n \in Spec(\mathcal{T})$ iff $n \in \{f(0), \dots, f(n)\}$.

That this results in a decision procedure follows from the fact that $\mathbf{minmod}_{\mathcal{T}}$ is assumed to be computable and so is producing the formulas $\#(x_1, \dots, x_{f(m)+1})$; we have left to prove that it is both sound and complete. If $n \in \{f(0), \dots, f(n)\}$, there exists an $0 \leq m < n$ such that $n = \mathbf{minmod}_{\mathcal{T}}(\#(x_1, \dots, x_{f(m)+1}))$, and so there is a \mathcal{T} -interpretation \mathcal{A} that satisfies $\#(x_1, \dots, x_{f(m)+1})$ with $|dom(\mathcal{A})| = n$; of course \mathcal{A} is a \mathcal{T} -interpretation with $|dom(\mathcal{A})| = n$, so $n \in Spec(\mathcal{T})$ and indeed the algorithm is sound. Reciprocally, if $n \in Spec(\mathcal{T})$, suppose that $f(m) < n < f(m+1)$ for some $0 \leq m < n$ (there is either such an element or $n = f(m)$ and we have nothing to prove, since $f(m+1) \geq f(m) + 1$): because n is in $Spec(\mathcal{T})$, there is a \mathcal{T} -interpretation \mathcal{A} with $|dom(\mathcal{A})| = n$; but since $n > f(m)$, $n \geq f(m) + 1$, and so \mathcal{A} satisfies $\#(x_1, \dots, x_{f(m)+1})$. This, in combination with the fact that $f(m+1) > n$, contradicts that $f(m+1) = \mathbf{minmod}_{\mathcal{T}}(\#(x_1, \dots, x_{f(m)+1}))$, proving that the algorithm is also complete and finishing the proof.

I Proof of Theorem 5

Lemma 8. *Given a sequence a_n/b_n converging to a $0 < r < 1$, the sequence of partial mediants*

$$\left\{ \dots, \frac{a_0 + \dots + a_n}{b_0 + \dots + b_n}, \frac{a_0 + \dots + a_n + 1}{b_0 + \dots + b_n + 1}, \dots, \frac{a_0 + \dots + a_n + a_{n+1}}{b_0 + \dots + b_n + a_{n+1}}, \frac{a_0 + \dots + a_n + a_{n+1} + 1}{b_0 + \dots + b_n + a_{n+1} + 1}, \dots \right\}$$

also converges to r .

Proof. Because $0 < r < 1$, we can assume that $0 < a_n < b_n$ for all $n \in \mathbb{N}$, what makes the definition of the partial mediants consistent. If we write $A(n) = \sum_{i=0}^n a_i$, and analogously for $B(n)$, and $A(m, n) = \sum_{i=n+1}^m a_i$ for $m > n$, and the same for $B(m, n)$, we see that the elements in our sequence of index $B(n) \leq i \leq B(n) + a_{n+1}$ lie between $A(n)/B(n)$ and $A(n+1)/(B(n) + a_{n+1})$, while those with index $B(n) + a_{n+1} \leq i \leq B(n+1)$ lie between $A(n+1)/(B(n) + a_{n+1})$ and $A(n+1)/B(n+1)$. So, it is enough to prove that both $A(n)/B(n)$ and $A(n+1)/(B(n) + a_{n+1})$ converge to r in order to sandwich the whole sequence of partial mediants into converging to r .

Take an $\epsilon > 0$: because $\lim_{n \rightarrow \infty} a_n/b_n = r$, there exists an $n_0 \in \mathbb{N}$ such that $|r - (a_n/b_n)| \leq \epsilon/2$ (and thus $|rb_n - a_n| \leq b_n\epsilon/2$) for all $n \geq n_0$. Then, for $n > m_0 = \max\{n_0, k\}$, where $k = \lceil 2|rB(n_0) - A(n_0)|/\epsilon \rceil$, we have

$$\begin{aligned} \left| r - \frac{A(n)}{B(n)} \right| &= \left| r - \frac{A(n_0) + A(n, n_0)}{B(n_0) + B(n, n_0)} \right| = \\ &= \left| \frac{rB(n_0) + rB(n, n_0) - A(n_0) - A(n, n_0)}{B(n_0) + B(n, n_0)} \right| \leq \\ &= \left| \frac{rB(n_0) - A(n_0)}{B(n_0) + B(n, n_0)} \right| + \left| \frac{rB(n, n_0) - A(n, n_0)}{B(n_0) + B(n, n_0)} \right| = \\ &= \left| \frac{rB(n_0) - A(n_0)}{B(n)} \right| + \left| \frac{\sum_{i=n_0+1}^n rb_i - a_i}{B(n)} \right| \leq \\ &= \frac{1}{n} |rB(n_0) - A(n_0)| + \frac{1}{B(n)} \sum_{i=n_0+1}^n |rb_i - a_i| \leq \\ &= \frac{\epsilon}{2} + \frac{1}{B(n)} \sum_{i=n_0+1}^n \frac{b_i\epsilon}{2} = \frac{\epsilon}{2} + \frac{B(n, n_0)\epsilon}{2B(n)} \leq \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon, \end{aligned}$$

thus proving $\lim_{n \rightarrow \infty} A(n)/B(n) = r$.¹¹ The proof for the sequence $A(n+1)/(B(n) + a_{n+1})$ is similar, so we only highlight the differences. $\epsilon/2$ must be replaced by $\epsilon/3$ since we get an extra term $a_{n+1}(1-r)/(B(n) + a_{n+1})$, which again can be bounded from the fact that $B(n) + a_{n+1} \geq n+1$: we just need to choose $m_0 = \max\{n_0, K\}$, where $K = \lceil 3 \max\{|rB(n_0) - A(n_0)|, a_{n+1}(1-r)\}/\epsilon \rceil$.

¹¹ Notice we have used that $B(n) \geq n$.

Lemma 9. *If A is a computable set with a well-defined density, $\mu(A)$ is a computable number.*

Proof. Since we have an algorithm for deciding whether a non-negative integer is in A or not, we also have one for the sequence a_n (where the input is n and the output is a_n), such that $a_0 = 1$ if $1 \in A$, and $a_0 = 0$ otherwise; and $a_n = a_{n-1} + 1$ if $n + 1 \in A$, and $a_n = a_{n-1}$ otherwise. Of course, the sequence $b_n = n + 1$ is also computable. Now, by definition of a_n , $a_n = |A_{n+1}|$, and therefore $\mu(A) = \lim_{n \rightarrow \infty} |A_n|/n = \lim_{n \rightarrow \infty} a_{n-1}/b_{n-1}$, so $\mu(A)$ is indeed computable.

Theorem 5. *If \mathcal{T} is a Σ_1 -theory with a well-defined density, the fact that $\mu(\mathcal{T})$ is a computable number is a necessary condition for \mathcal{T} to have a computable minimal model function. Furthermore, for every computable number $0 \leq r \leq 1$, there exists a Σ_1 -theory \mathcal{T} with $\mu(\mathcal{T}) = r$ that has a computable minimal model function and the finite model property, but is not smooth.*

Proof. The first part of the theorem follows from Proposition 1 and lemma 9.

If $r = 0$, we take the theory $\mathcal{T}_{=2^i}$ from Example 8, axiomatized by $\{\psi_{\geq 2^{n+1}} \vee \bigvee_{i=1}^n \psi_{=2^i} : n \in \mathbb{N}^*\}$; if $r = 1$, we take the theory $\mathcal{T}_{\neq 2^i}$ from Example 11, axiomatized by $\{\neg \psi_{=2^n} : n \in \mathbb{N}^*\}$. We already know $\mu(\mathcal{T}_0) = 0$ and $\mu(\mathcal{T}_1) = 1$: since their spectra are, respectively, the sets of power of two and non powers of two, both computable, we have by Proposition 1 that both have computable minimal model functions.

With these cases out of the way, we can assume that $0 < r < 1$, and thus there exist computable sequences (since r is computable) $\{a_n\}_{n \in \mathbb{N}}$ in \mathbb{N} , and $\{b_n\}_{n \in \mathbb{N}}$ in \mathbb{N}^* such that $r = \lim_{n \rightarrow \infty} a_n/b_n$ with $0 < a_n < b_n$ for all $n \in \mathbb{N}$. We then take the function f associated with $\{a_n\}_{n \in \mathbb{N}}$ and $\{b_n\}_{n \in \mathbb{N}}$, from Definition 2, and define \mathcal{T} to be axiomatized by $\{\psi_{\geq f(n+1)} \vee \bigvee_{i=1}^n \psi_{=f(i)} : n \in \mathbb{N}^*\}$, from what it is clear that $\text{Spec}(\mathcal{T}) = \{f(n) : n \in \mathbb{N}^*\}$. Since $\{a_n\}_{n \in \mathbb{N}}$ and $\{b_n\}_{n \in \mathbb{N}}$ are computable, so is f : indeed, the computability of $\{b_n\}_{n \in \mathbb{N}}$ implies $M = \sum_{i=0}^m b_i$ is computable; and then the computability of $\{a_n\}_{n \in \mathbb{N}}$ implies $f(n) = n$, for $M + 1 \leq n \leq M + a_m$, and $f(n) = n + a_m$, for $M + a_m + 1 \leq n \leq M + b_m$, are computable. Therefore $\text{Spec}(\mathcal{T})$ is computable: indeed, as $\text{Spec}(\mathcal{T})$ is the image of the non-decreasing function f , and $f(1 + \sum_{i=0}^m b_i) = 1 + \sum_{i=0}^m b_i$, it is enough, to test whether $n \in \text{Spec}(\mathcal{T})$, to check if $n \in \{f(1), \dots, f(1 + \sum_{i=0}^n b_i)\}$ (recall $\{b_n\}_{n \in \mathbb{N}}$ is a computable sequence of positive numbers, so $\sum_{i=0}^n b_i \geq n$). From Proposition 1 we have \mathcal{T} has a computable minimal model function, and thus from [24, Theorem 4] we have \mathcal{T} is finitely witnessable. Finally, $\mu(\mathcal{T})$ will equal the limit of partial mediants as defined in Lemma 8, if the latter is well-defined, what the result further shows it is and equals in turn r : indeed, assuming for an inductive argument that $|\text{Spec}_M \mathcal{T}|/M$ equals $(a_0 + \dots + a_m)/(b_0 + \dots + b_m)$ for $M = b_0 + \dots + b_m$, since $f(n) = n$ for $M + 1 \leq n \leq M + a_{m+1}$ we obtain

$$\frac{|\text{Spec}_n(\mathcal{T})|}{n} = \frac{a_0 + \dots + a_m + (n - M)}{n} = \frac{a_0 + \dots + a_m + (n - M)}{b_0 + \dots + b_m + (n - M)}$$

for these values; and for $M + a_{m+1} + 1 \leq n \leq M + b_{m+1}$ we have $f(n) = n + a_m$, so for these values

$$\frac{|Spec_n(\mathcal{T})|}{n} = \frac{a_0 + \dots + a_m + a_{m+1}}{n} = \frac{a_0 + \dots + a_m + a_{m+1}}{b_0 + \dots + b_m + (n - M - a_{m+1})},$$

and in both cases we obtain the aforementioned partial mediant.

J Proof of Theorem 6

Theorem 6. *If \mathcal{T} is a Σ_1 -theory, nothing can be said about $\mu(\mathcal{T})$ if \mathcal{T} is only known to be finitely witnessable; that is, for every number $0 \leq r \leq 1$, there exists a Σ_1 -theory \mathcal{T} with $\mu(\mathcal{T}) = r$ that is finitely witnessable yet doesn't have a computable minimal model function.*

Proof. We divide the proof in two big cases: when r is not computable, and when it is.

1. Of course $0 < r < 1$, as both 0 and 1 are computable: because r can be written in decimal notation as $0.0\dots 0d_0d_1d_2\dots$, for digits $d_i \in \{0, 1, \dots, 9\}$ and $d_0 \neq 0$, we can write r as the limit of a_n/b_n , where $a_n = d_0\dots d_n$, $b_n = 10^{M+n}$, and M is the number of zeros before d_0 (including the one before the decimal dot); this way, $0 < a_n < b_n$ for all $n \in \mathbb{N}$. By the fact that r is not computable, and since $\{b_n\}_{n \in \mathbb{N}}$ is certainly computable, $\{a_n\}_{n \in \mathbb{N}}$ cannot be so. We take the function f associated with $\{a_n\}_{n \in \mathbb{N}}$ and $\{b_n\}_{n \in \mathbb{N}}$, as in Definition 2, and the theory \mathcal{T} axiomatized by $\{\psi_{\geq f(n+1)} \vee \bigvee_{i=1}^n \psi_{=f(i)} : n \in \mathbb{N}^*\}$: the theory is not strongly finitely witnessable because it is not smooth (see [25, Theorem 7]); furthermore, since $\{b_n\}_{n \in \mathbb{N}}$ is computable but $\{a_n\}_{n \in \mathbb{N}}$ is not, f is also not computable, and since $Spec(\mathcal{T}) = \{f(n) : n \in \mathbb{N}^*\}$, we have that \mathcal{T} does not have a computable minimal model function (Proposition 1).

Is is clear that $\mu(\mathcal{T})$ is the limit of the mediant $(a_0 + \dots + a_n)/(b_0 + \dots + b_n)$, which from Lemma 8 is $\lim_{n \rightarrow \infty} a_n/b_n = r$; so we only have left to prove that \mathcal{T} has a witness. For a quantifier-free formula ϕ , let N be the smallest $n \geq M$ such that $10^N \geq |vars(\phi)|$, define $N(\phi) = \sum_{i=M}^N 10^i$, and

$$wit(\phi) = \phi \wedge \bigwedge_{i=1}^{N(\phi)} x_i = x_i$$

where the x_i are fresh variables: we claim this function is a witness, obviously mapping quantifier-free formulas into quantifier-free formulas, and being computable, given that finding $|vars(\phi)|$ and N are both computable procedures. Furthermore, for $\vec{x} = vars(wit(\phi)) \setminus vars(\phi)$, it is clear that $\exists \vec{x}. wit(\phi)$ and ϕ are \mathcal{T} -equivalent, given that in fact ϕ and $wit(\phi)$ are themselves equivalent. Suppose then that \mathcal{A} is a \mathcal{T} -interpretation that satisfies ϕ , take a set X disjoint from $dom(\mathcal{A})$ with cardinality $N(\phi) - |vars(\phi)^{\mathcal{A}}|$, and we construct an interpretation \mathcal{B} by making: $dom(\mathcal{B}) = vars(\phi)^{\mathcal{A}} \cup X$ (so $|dom(\mathcal{B})| = N(\phi)$, which is a cardinality in $Spec(\mathcal{T})$); $x^{\mathcal{B}} = x^{\mathcal{A}}$ for all

$x \in \text{vars}(\phi)$; $x \in \{x_i : 1 \leq i \leq N(\phi)\} \mapsto x^{\mathcal{B}}$ an injective function; and $x^{\mathcal{B}}$ can be set arbitrarily for all other variables. Then it is true that \mathcal{B} satisfies $\text{wit}(\phi)$ and $\text{dom}(\mathcal{B}) = \text{vars}(\text{wit}(\phi))^{\mathcal{B}}$, finishing the proof.

2. Because $0 \leq r \leq 1$ is computable, there are computable sequences $\{a_n\}_{n \in \mathbb{N}}$ and $\{b_n\}_{n \in \mathbb{N}}$ such that $0 < a_n < b_n$ and $\lim_{n \rightarrow \infty} a_n/b_n = r$. Now, take a non-computable function $g : \mathbb{N}^* \rightarrow \{0, 1\}$. We then define a function $G : \mathbb{N}^* \rightarrow \{0, 1\}$ by induction.

- (a) i. $G(1) = g(1)$.
 ii. $G(n) = 0$ for: $2 \leq n \leq 2(b_0 - a_0)$ (since $b_0 > a_0$, $2(b_0 - a_0) \geq 2$) if $g(1) = 0$; and $2 \leq n \leq 2(b_0 - a_0) + 1$ if $g(1) = 1$.
 iii. $G(n) = 1$ for: $2(b_0 - a_0) + 1 \leq n \leq 2b_0$ if $g(1) = 0$; and $2(b_0 - a_0) + 2 \leq n \leq 2b_0$ (since $b_0 > a_0$, $2b_0 \geq 2(b_0 - a_0) + 2$) if $g(1) = 1$.
 (b) Now, assume $G(n)$ has been defined for $1 \leq n \leq 2 \sum_{i=0}^m b_i$.
 i. $G(1 + 2 \sum_{i=0}^m b_i) = g(m+2)$;
 ii. $G(n) = 0$ for: $2 + 2 \sum_{i=0}^m b_i \leq n \leq 2(b_{m+1} - a_{m+1}) + 2 \sum_{i=0}^m b_i$ if $g(m+2) = 0$; and $2 + 2 \sum_{i=0}^m b_i \leq n \leq 1 + 2(b_{m+1} - a_{m+1}) + 2 \sum_{i=0}^m b_i$ if $g(m+2) = 1$.
 iii. $G(n) = 1$ for: $1 + 2(b_{m+1} - a_{m+1}) + 2 \sum_{i=0}^m b_i \leq n \leq 2 \sum_{i=0}^{m+1} b_i$ if $g(m+2) = 0$; and $2 + 2(b_{m+1} - a_{m+1}) + 2 \sum_{i=0}^m b_i \leq n \leq 2 \sum_{i=0}^{m+1} b_i$ if $g(m+2) = 1$.

We see that G is not computable because: the function $\rho : \mathbb{N}^* \rightarrow \mathbb{N}^*$ given by $\rho(1) = 1$ and $\rho(m) = 1 + 2 \sum_{i=0}^{m-1} b_i$, for $m > 1$, is computable, given that $\{b_n\}_{n \in \mathbb{N}}$ is computable; and, despite that, $G \circ \rho(m) = g(m)$, which is not computable.

We can finally define the Σ_1 -theory by the axiomatization

$$\{\psi_{\geq f(n+1)} \vee \bigvee_{i=1}^n \psi_{=f(i)} : n \in \mathbb{N}^*\},$$

where $f : \mathbb{N}^* \rightarrow \mathbb{N}^*$ is defined as $f(n) = \min\{m \geq n : G(m) = 1\}$, so that $\text{Spec}(\mathcal{T}) = \{f(n) : n \in \mathbb{N}^*\}$, and $|\text{Spec}_n(\mathcal{T})| = |\{1 \leq i \leq n : G(i) = 1\}|$: this means $|\text{Spec}_{2 \sum_{i=0}^m b_i}(\mathcal{T})| = 2 \sum_{i=0}^m a_i$, leading us again to partial mediants; from Lemma 8 we get $\mu(\mathcal{T})$ equals $\lim_{n \rightarrow \infty} 2a_n/2b_n$, which is precisely r . From the fact that G is not computable it follows that $\text{Spec}(\mathcal{T})$ is not computable, so from Proposition 1 \mathcal{T} does not have a computable minimal model function. It only remains for us to show that \mathcal{T} has a witness: so, for a quantifier-free formula ϕ , let $N(\phi)$ be the minimum of $2 \sum_{i=0}^m b_i$ greater than or equal to $|\text{vars}(\phi)|$; this value is computable given the sequences $\{a_n\}_{n \in \mathbb{N}}$ and $\{b_n\}_{n \in \mathbb{N}}$ are themselves computable. For x_i fresh variables, we define

$$\text{wit}(\phi) = \phi \wedge \bigwedge_{i=1}^{N(\phi)} x_i = x_i :$$

this obviously maps quantifier-free formulas into themselves, and is computable given that finding $N(\phi)$ from ϕ can be done algorithmically. It is also clear that ϕ and $\exists \vec{x}. \text{wit}(\phi)$ are \mathcal{T} -equivalent, for $\vec{x} = \text{vars}(\text{wit}(\phi)) \setminus \text{vars}(\phi)$, since $\text{wit}(\phi)$ is the conjunction of ϕ and a tautology. Finally, suppose that \mathcal{A} is a \mathcal{T} -interpretation that satisfies ϕ , and take a set X disjoint

from $\text{dom}(\mathcal{A})$ with $N(\phi) - |\text{vars}(\phi)^{\mathcal{A}}|$ elements. We define an interpretation \mathcal{B} by making: $\text{dom}(\mathcal{B}) = \text{vars}(\phi)^{\mathcal{A}} \cup X$ (so $|\text{dom}(\mathcal{B})| = N(\phi)$, and since $N(\phi) = 2 \sum_{i=0}^m b_i$ and $G(2 \sum_{i=0}^m b_i) = 1$, we obtain \mathcal{B} is a \mathcal{T} -interpretation); $x^{\mathcal{B}} = x^{\mathcal{A}}$ for all $x \in \text{vars}(\phi)$; $x_i \in \{x_i : 1 \leq i \leq N(\phi)\} \mapsto x_i^{\mathcal{B}} \in \text{dom}(\mathcal{B})$ an injective map; and arbitrarily for all other variables. We then have that \mathcal{B} is a \mathcal{T} -interpretation that satisfies ϕ , and $\text{dom}(\mathcal{B}) = \text{vars}(\text{wit}(\phi))^{\mathcal{B}}$.

K Proof of Theorem 7

Theorem 7. *The positivity of $\mu(\mathcal{T}, \phi)$ for every \mathcal{T} -satisfiable quantifier-free formula ϕ is sufficient for \mathcal{T} to be stably infinite and have the finite model property.*

Proof. Suppose \mathcal{T} is not stably infinite, and thus there exists a quantifier-free formula ϕ such that no infinite \mathcal{T} -interpretation satisfies ϕ . This means there exists an $M \in \mathbb{N}$ such that no \mathcal{T} -interpretation \mathcal{A} that satisfies ϕ has $|\text{dom}(\mathcal{A})| > M$: were that not true, we could obtain a sequence of \mathcal{T} -interpretations \mathcal{A}_n that satisfy ϕ , with $|\text{dom}(\mathcal{A}_n)| \geq n$; since \mathcal{A}_n satisfies the axiomatization of \mathcal{T} , ϕ and $\{\psi_{\geq 1}, \dots, \psi_{\geq n}\}$, we get by Theorem 14 that the axiomatization of \mathcal{T} , ϕ and $\{\psi_{\geq n} : n \in \mathbb{N}\}$ are simultaneously satisfiable, leading to a contradiction. So $\text{Spec}_n(\mathcal{T}, \phi) = \text{Spec}_M(\mathcal{T}, \phi)$ for $n \geq M$, giving us $\mu(\mathcal{T}, \phi) \leq \lim_{n \rightarrow \infty} M/n = 0$, and the result for stable infiniteness follows from the contrapositive.

If \mathcal{T} does not have the finite model property, there exists a quantifier-free formula ϕ such that no finite \mathcal{T} -interpretation satisfies ϕ : that means $\text{Spec}(\mathcal{T}, \phi) = \emptyset$, and so $\mu(\mathcal{T}, \phi) = 0$, the result for the finite model property easily following.

L Proof of Theorem 9

Theorem 9. *Let \mathcal{T} be a theory. $\mu(\mathcal{T}, \phi)$ being well-defined and equal to 1 for all \mathcal{T} -satisfiable quantifier-free formulas ϕ is then a necessary condition for \mathcal{T} to simultaneously be smooth and have the finite model property.*

Proof. If \mathcal{T} is smooth and has the finite model property, for every quantifier-free \mathcal{T} -satisfiable formula ϕ there exists $M \in \mathbb{N}$ such that for all $n \geq M$ there exists a \mathcal{T} -interpretation \mathcal{A} that satisfies ϕ with $|\text{dom}(\mathcal{A})| = n$. That means $[1, n] \setminus \text{Spec}_n(\mathcal{T}, \phi) \subseteq [1, M]$, for $n \geq M$, and so $\mu(\mathcal{T}, \phi) \geq \lim_{n \rightarrow \infty} (n - M)/n = 1$.

M Proof of Theorem 10

The following technical lemma is a version of [25, Theorem 7], modified to hold locally, that is, for every quantifier-free formula.

Lemma 10. *If \mathcal{T} is a strongly finitely witnessable theory, for all quantifier-free formulas ϕ , we have that: either ϕ is not satisfied by any infinite \mathcal{T} -interpretations; or there exists $M \in \mathbb{N}$ such that, for all cardinals $M \leq \kappa \leq \aleph_0$, there is a \mathcal{T} -interpretation \mathcal{A} that satisfies ϕ with $|\text{dom}(\mathcal{A})| = \kappa$.*

Proof. Suppose that ϕ is satisfied by an infinite \mathcal{T} -interpretation \mathcal{A} : We may assume that $|dom(\mathcal{A})| = \aleph_0$ by Theorem 13. Since ϕ and $\exists \vec{x}.wit(\phi)$ are \mathcal{T} -equivalent, for $\vec{x} = vars(wit(\phi)) \setminus vars(\phi)$, we may change the value assigned to some variables by \mathcal{A} in order to obtain a \mathcal{T} -interpretation \mathcal{A}' that satisfies $wit(\phi)$ which still has $|dom(\mathcal{A}')| = \aleph_0$. Let U be $vars(wit(\phi))$, and δ_U the arrangement on U such that x is associated with y iff $x^{\mathcal{A}'} = y^{\mathcal{A}'}$, meaning that \mathcal{A}' satisfies δ_U . From that, \mathcal{A}' satisfies $wit(\phi) \wedge \delta_U$, meaning there is a \mathcal{T} -interpretation \mathcal{B} (from the fact that wit is a strong witness) that satisfies $wit(\phi) \wedge \delta_U$, and thus ϕ , with $dom(\mathcal{B}) = vars(wit(\phi) \wedge \delta_U)^{\mathcal{B}} = U^{\mathcal{B}}$.

We state it is now possible to find a \mathcal{T} -interpretation \mathcal{C} that satisfies ϕ with any cardinality κ between $M = |dom(\mathcal{B})|$ and \aleph_0 . Indeed, if $\kappa = M$ we are done. Otherwise, take a fresh set of variables U' with cardinality $\kappa - M$, and define an arrangement δ_V on $V = U \cup U'$ such that x is related to y according to δ_V iff either $x = y$, or $x, y \in U$ and $x^{\mathcal{B}} = y^{\mathcal{B}}$; notice we can write

$$\delta_V = \delta_U \wedge \delta_{U'} \wedge \bigwedge_{x \in U, y \in U'} \neg(x = y),$$

where $\delta_{U'}$ corresponds to the identity on U' . Notice now that $wit(\phi) \wedge \delta_V$ is \mathcal{T} -satisfiable: indeed, \mathcal{A}' satisfies $wit(\phi) \wedge \delta_U$, and by changing at most the values assigned by \mathcal{A}' to the (fresh) variables in U' , we get a \mathcal{T} -interpretation \mathcal{A}'' that satisfies $wit(\phi) \wedge \delta_V$.

This means, again from the fact that wit is a strong witness, that there is a \mathcal{T} -interpretation \mathcal{C} that satisfies $wit(\phi) \wedge \delta_V$, with $dom(\mathcal{C}) = vars(wit(\phi) \wedge \delta_V)^{\mathcal{C}} = V^{\mathcal{C}}$. Since \mathcal{C} satisfies $wit(\phi)$, it satisfies $\exists \vec{x}.wit(\phi)$ and thus ϕ . And since $dom(\mathcal{C}) = V^{\mathcal{C}}$, where $U^{\mathcal{C}}$ must have M elements (since \mathcal{C} satisfies δ_U), and $(U')^{\mathcal{C}}$ must have another $\kappa - M$ elements (since \mathcal{C} satisfies δ_V), we get $|dom(\mathcal{C})| = \kappa$, as we wanted to show.

Theorem 10. $\mu(\mathcal{T}, \phi)$ being well-defined and equal to 0 or 1 for every quantifier-free \mathcal{T} -satisfiable ϕ is necessary for \mathcal{T} to be strongly finitely witnessable.

Proof. By Lemma 10, either ϕ is not satisfied by an infinite \mathcal{T} -interpretation, or $Spec(\mathcal{T}, \phi)$ contains all elements greater than or equal to some M : from the second case it clearly follows that $\mu(\mathcal{T}, \phi) \geq \lim_{n \rightarrow \infty} (n - M + 1)/n = 1$. In the first, we state $Spec(\mathcal{T}, \phi)$ must be finite: otherwise we should be able to get an infinite \mathcal{T} -interpretation that satisfies ϕ by using Theorem 14. This way, $Spec(\mathcal{T}, \phi)$ must be bounded by some M , and so $\mu(\mathcal{T}, \phi) \leq \lim_{n \rightarrow \infty} M/n = 0$.

N Proof of Proposition 2

Proposition 2. If \mathcal{T} is a theory with a computable minimal model function, then $Spec(\mathcal{T}, \phi)$ is computable for all quantifier-free \mathcal{T} -satisfiable formulas ϕ .

Proof. Suppose that \mathcal{T} has a computable minimal model function $\mathbf{minmod}_{\mathcal{T}}$, and we shall consider two cases: one where $Spec(\mathcal{T}, \phi)$ is finite, and one where it is infinite. If $Spec(\mathcal{T}, \phi)$ is finite we have nothing left to do, as it is enough to simply

hardcode these values into an algorithm that decides whether an element is in $\text{Spec}(\mathcal{T}, \phi)$. If $\text{Spec}(\mathcal{T}, \phi)$ is instead infinite, ϕ is \mathcal{T} -satisfiable (as $\text{Spec}(\mathcal{T}, \phi)$ is not empty) and, in addition, the formulas $\phi \wedge \#(x_1, \dots, x_n)$ are \mathcal{T} -satisfiable for all $n \in \mathbb{N}$: indeed, since $\text{Spec}(\mathcal{T}, \phi)$ is infinite, we can always find $m \in \text{Spec}(\mathcal{T}, \phi)$ such that $m \geq n$, and thus there exists a \mathcal{T} -interpretation \mathcal{A} that satisfies ϕ with $|\text{dom}(\mathcal{A})| = m$; by changing at most the value assigned to the variables x_1, \dots, x_n (assumed fresh), we get a \mathcal{T} -interpretation that satisfies $\phi \wedge \#(x_1, \dots, x_n)$. Given then n , we define $f(0) = \mathbf{minmod}_{\mathcal{T}}(\phi)$, and assuming $f(m)$ defined $f(m+1) = \mathbf{minmod}_{\mathcal{T}}(\phi \wedge \#(x_1, \dots, x_{f(m)+1}))$, and then we state that $n \in \text{Spec}(\mathcal{T}, \phi)$ iff $n \in \{f(0), \dots, f(n)\}$.

That this results in a decision procedure follows from the fact that $\mathbf{minmod}_{\mathcal{T}}$ is assumed to be computable, and so is producing the formulas $\#(x_1, \dots, x_{f(m)+1})$; we have left to prove that it is both sound and complete. If $n \in \{f(0), \dots, f(n)\}$, there exists an $0 \leq m < n$ such that $n = \mathbf{minmod}_{\mathcal{T}}(\phi \wedge \#(x_1, \dots, x_{f(m)+1}))$, and so there is a \mathcal{T} -interpretation \mathcal{A} that satisfies $\phi \wedge \#(x_1, \dots, x_{f(m)+1})$ with $|\text{dom}(\mathcal{A})| = n$; of course \mathcal{A} is a \mathcal{T} -interpretation that satisfies ϕ with $|\text{dom}(\mathcal{A})| = n$, so $n \in \text{Spec}(\mathcal{T}, \phi)$ and indeed the algorithm is sound. Reciprocally, if $n \in \text{Spec}(\mathcal{T}, \phi)$, suppose that $f(m) < n < f(m+1)$ for some $0 \leq m < n$ (there is either such an element or $n = f(m)$ and we have nothing to prove, since $f(m+1) \geq f(m) + 1$): because n is in $\text{Spec}(\mathcal{T}, \phi)$, there is a \mathcal{T} -interpretation \mathcal{A} that satisfies ϕ with $|\text{dom}(\mathcal{A})| = n$; but since $n > f(m)$, $n \geq f(m) + 1$, and so \mathcal{A} satisfies $\phi \wedge \#(x_1, \dots, x_{f(m)+1})$. This, in combination with the fact that $f(m+1) > n$, contradicts that $f(m+1) = \mathbf{minmod}_{\mathcal{T}}(\phi \wedge \#(x_1, \dots, x_{f(m)+1}))$, proving that the algorithm is also complete and finishing the proof.

O Proof of Theorem 11

Theorem 11. *If \mathcal{T} is a theory with well-defined densities $\mu(\mathcal{T}, \phi)$, for all quantifier-free \mathcal{T} -satisfiable formulas ϕ , the fact that all $\mu(\mathcal{T}, \phi)$ are computable is a necessary condition for \mathcal{T} to have a computable minimal model function. Furthermore, for every computable number $0 \leq r \leq 1$, there is a theory \mathcal{T} that has a computable minimal model function and a quantifier-free formula ϕ with $\mu(\mathcal{T}, \phi) = r$.*

Proof. If \mathcal{T} has a minimal model function, Proposition 2 guarantees all $\text{Spec}(\mathcal{T}, \phi)$ are computable, and by Lemma 9 it follows that all $\mu(\mathcal{T}, \phi)$ are computable.

Finally, for every computable r , we get a theory \mathcal{T} with a computable minimal model function and $\mu(\mathcal{T}) = r$ from Theorem 5, and it follows that for any tautology ϕ it is true that $\mu(\mathcal{T}, \phi) = r$.

P Proof of Theorem 12

Theorem 12. *If \mathcal{T} is a theory, and ϕ a quantifier-free \mathcal{T} -satisfiable formula, nothing can be said about $\mu(\mathcal{T}, \phi)$ if \mathcal{T} is only known to be finitely witnessable; that is, for every computable number $0 \leq r \leq 1$, there exists a theory \mathcal{T} , that is finitely witnessable, and a quantifier-free formula ϕ with $\mu(\mathcal{T}, \phi) = r$.*

Proof. Follows trivially from Theorem 6: taking, for any $0 \leq r \leq 1$, a finitely witnessable theory \mathcal{T} (without a computable minimal model function) with $\mu(\mathcal{T}) = r$, and a tautology ϕ , $\mu(\mathcal{T}, \phi) = r$.