Being polite is not enough (and other limits of theory combination)

Guilherme V. Toledo[1], Benjamin Przybocki[2], and Yoni Zohar[1]

[1] Bar-Ilan University, Israel
[2] University of Cambridge, UK

Abstract. In the Nelson–Oppen combination method for satisfiability modulo theories, the combined theories must be stably infinite; in gentle combination, one theory has to be gentle, and the other has to satisfy a similar yet weaker property; in shiny combination, only one has to be shiny (smooth, with a computable minimal model function and the finite model property); and for polite combination, only one has to be strongly polite (smooth and strongly finitely witnessable). For each combination method, we prove that if any of its assumptions are removed, then there is no general method to combine an arbitrary pair of theories satisfying the remaining assumptions. We also prove new theory combination results that weaken the assumptions of gentle and shiny combination.[3]

1 Introduction

Let us start at the middle. Polite theory combination [17] was not the first method to combine two theories (see, e.g., [15,23]). It was also not the last (see, e.g., [10,14,4]). However, it is one of the most influential approaches to theory combination. In fact, it has found its way to the implementation of the state-of-the-art SMT-solver cvc5 [1] (and also CVC4 [2] and CVC3 [3]).

The history of polite combination is illustrated in Figure 1, that focuses on decidable theories (marked by the large rectangle). The left circle corresponds to decidable theories that can be combined with any other decidable theory over a disjoint signature. We call such theories *combinable*. In [17], it was argued that a sufficient condition for combinability is *politeness*, a technical notion that concerns cardinalities of models. In other words, [17] claimed that the red-hatched region of Figure 1 is empty. Then, the paper [11] discovered a bug in the proof from [17], and offered to replace politeness by a seemingly stronger notion, *strong politeness* (the name is due to [7]). It was proved in [11] that strongly polite theories are combinable, which positions the small circle that represents strongly polite theories completely within the circle representing combinable theories.

While [11] found a bug in the *proof* of [17], it left two questions open: (i) does politeness imply combinability, as [17] claimed, only with a different proof? and (ii) do polite theories that are not strongly polite exist? In terms of Figure 1:

[3] This work was funded by NSF-BSF grant 2020704, ISF grant 619/21, and the Colman-Soref fellowship.

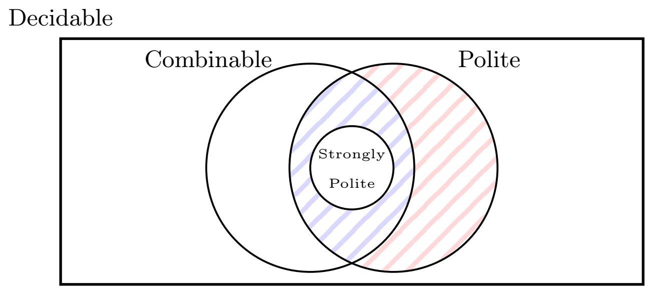

Fig. 1: An illustration of the contributions of [17,11,25] and the current paper.

(i) is the red-hatched region empty? and (ii) is the entire hatched region empty? Question (ii) was recently resolved in [21,25]: a theory named $\mathcal{T}_f$ was found wandering around the hatched region. However, question (i) remained unanswered, as it was unclear whether $\mathcal{T}_f$ resided in the red hatched region or the blue one.

In this paper we solve question (i), by placing $\mathcal{T}_f$ in the red-hatched region: it is polite but uncombinable. To show this, we introduce a new decidable theory, named $\mathcal{T}_=$, over a disjoint signature, and prove that its combination with $\mathcal{T}_f$ is *undecidable*.

$\mathcal{T}_f$ and $\mathcal{T}_=$ are not merely mustard watches [19]. They show that being polite is not enough (for theory combination), finally closing the question of politeness vs. combinability, that remained open since [11]. Foundationally, they show that the fix of [11] was indeed necessary. And practically, they justify the implementation overhead of adopting the more complicated definition of strong politeness.

The existence of $\mathcal{T}_f$ and $\mathcal{T}_=$ can be seen as a *limitation theorem*: they show that the polite combination method cannot be applied if strong politeness is weakened to politeness. We present similar limitation theorems for other combination methods. For the Nelson–Oppen method [15], we show that if only one of the theories is assumed to be stably infinite, then the combination method fails; this was previously proven in [5], but that study did not broach any combination methods other than Nelson–Oppen. We provide a similar treatment for gentle combination [10]. Finally, for a theory to be shiny [22], it has to satisfy three requirements. For each one, we show its necessity as well. Surprisingly, we can reuse $\mathcal{T}_f$ and $\mathcal{T}_=$ for almost all limitation theorems, except for one. Therefore, we use two more theories. The first, $\mathcal{T}_\infty$, is also taken from [25]. The second, $\mathcal{T}_\leq$, is new, and can be seen as a generalization of the theory from [5] that was used for the Nelson-Oppen limitation theorem. For all theories (old and new), we prove that they satisfy the required properties for each limitation theorem.

Additionally, we prove two new combination theorems, based on the gentle and shiny combination methods. These theorems relax some of the requirements for theory combination, while ensuring the decidability of the combined theories. In a sense, the new theorems remedy the limitation theorems that we prove.

This paper is organized as follows. Section 2 surveys notions regarding theory combination. Section 3 proves our main theorems, exhibiting limits of common combination methods. Section 4 improves the proofs of Section 3 by only us-

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

$$\psi_{\geq n} = \exists\, x_1 \ldots x_n.\, \neq(x_1,\ldots,x_n) \qquad \psi_{=n} = \psi_{\geq n} \wedge \psi_{\leq n}$$

Fig. 2: Cardinality formulas.

ing theories over finite signatures. Section 5 proves new combination theorems. Section 6 concludes and provides directions for future research.

2 Preliminaries

We use $\mathbb{N}$ to denote the set of naturals including 0, and $\mathbb{N}^*$ to denote $\mathbb{N} \setminus \{0\}$. If X is a set, $|X|$ is its cardinality, and $|\mathbb{N}| = \aleph_0$.

2.1 First-order logic

A *signature* is a pair $\Sigma = (\mathcal{F}_\Sigma, \mathcal{P}_\Sigma)$ where: $\mathcal{F}_\Sigma$ is a set of function symbols, each with arity $n \in \mathbb{N}$; and $\mathcal{P}_\Sigma$ is a set of predicate symbols, each with arity $m \in \mathbb{N}$, containing at least the equality $=$, of arity 2. Σ is said to be *empty* if it has no function and predicate symbols other than $=$. Two signatures are said to be *disjoint* if the only symbol they share is $=$. We define *terms*, *literals*, *clauses* (i.e., disjunctions of literals), *cubes* (i.e., conjunctions of literals), *formulas* and *sentences* in the usual way. If t is a unary function symbol and x a variable, we define by induction the terms $t^0(x) := x$ and $t^{n+1}(x) := t(t^n(x))$. The set of variables in a formula φ is denoted by $vars(\varphi)$.

Σ*-interpretations* $\mathcal{A}$ are defined as usual: $dom(\mathcal{A})$ is the *domain* of $\mathcal{A}$; for a function symbol f of arity n, $f^{\mathcal{A}} : dom(\mathcal{A})^n \to dom(\mathcal{A})$; for a predicate symbol P of arity m, $P^{\mathcal{A}} \subseteq dom(\mathcal{A})^m$; for a variable x, $x^{\mathcal{A}} \in dom(\mathcal{A})$. For a term α, $\alpha^{\mathcal{A}}$ is its value in $\mathcal{A}$, and for a set Γ of terms, $\Gamma^{\mathcal{A}} = \{\alpha^{\mathcal{A}} : \alpha \in \Gamma\}$. If $\mathcal{A}$ satisfies φ, we write $\mathcal{A} \vDash \varphi$. Formulas from Figure 2 are satisfied by $\mathcal{A}$ when $|dom(\mathcal{A})|$ is: at least n (for $\neq(x_1,\ldots,x_n)$ and $\psi_{\geq n}$); at most n (for $\psi_{\leq n}$); exactly n (for $\psi_{=n}$).

A *theory* $\mathcal{T}$ is a class of all interpretations that satisfy a (finite or infinite) set of sentences $Ax(\mathcal{T})$ called the *axiomatization* of $\mathcal{T}$; φ is said to be $\mathcal{T}$*-satisfiable* when there is a $\mathcal{T}$-interpretation satisfying φ; it is $\mathcal{T}$*-valid* when every $\mathcal{T}$-interpretation satisfies φ. Two formulas φ and ψ are $\mathcal{T}$*-equivalent* when a $\mathcal{T}$-interpretation satisfies φ if and only if it satisfies ψ. $\mathcal{T}$ is *decidable* if the set of $\mathcal{T}$-satisfiable quantifier-free formulas is decidable.

2.2 Theory combination theorems

In what follows, Σ, Σ_1, Σ_2 are signatures, and $\mathcal{T}$, $\mathcal{T}_1$, $\mathcal{T}_2$ are Σ, Σ_1, Σ_2-theories, respectively. We assume Σ_1 and Σ_2 are disjoint, and $\mathcal{T}_1$ and $\mathcal{T}_2$ are decidable.

$\Sigma_1 \cup \Sigma_2$ is the signature obtained by collecting all function and predicates symbols from Σ_1 and Σ_2. $\mathcal{T}_1 \oplus \mathcal{T}_2$ is the $\Sigma_1 \cup \Sigma_2$-theory axiomatized by $Ax(\mathcal{T}_1) \cup Ax(\mathcal{T}_2)$.

We start with Nelson–Oppen. $\mathcal{T}$ is *stably infinite* if for every quantifier-free $\mathcal{T}$-satisfiable formula φ there is an infinite $\mathcal{T}$-interpretation $\mathcal{A}$ with $\mathcal{A} \models \varphi$.

Theorem 1 ([15]). *$\mathcal{T}_1 \oplus \mathcal{T}_2$ is decidable, if both $\mathcal{T}_1$ and $\mathcal{T}_2$ are stably infinite.*

Next, we define politeness. $\mathcal{T}$ is *smooth* if for every quantifier-free $\mathcal{T}$-satisfiable formula φ, $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies φ, and cardinal $\kappa > |dom(\mathcal{A})|$, there is a $\mathcal{T}$-interpretation $\mathcal{B}$ that satisfies φ with $|dom(\mathcal{A})| = \kappa$ (notice smoothness implies stable infiniteness, as we can choose an infinite κ). $\mathcal{T}$ is *finitely witnessable* if there exists a function *wit* (called a witness) from the quantifier-free formulas of Σ into themselves such that, for every quantifier-free formula φ, one has that: (I) φ and $\exists\, \overrightarrow{x}.\, wit(\varphi)$ are $\mathcal{T}$-equivalent, where $\overrightarrow{x} = vars(wit(\varphi)) \setminus vars(\varphi)$; (II) if $wit(\varphi)$ is $\mathcal{T}$-satisfiable there exists a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies $wit(\varphi)$ with $dom(\mathcal{A}) = vars(wit(\varphi))^{\mathcal{A}}$. $\mathcal{T}$ is *polite* if it is both smooth and finitely witnessable. The following was stated as a theorem in [18], but its proof was later refuted in [11]. It therefore it remained a conjecture, which essentially states that politeness is enough for theory combination.

Conjecture 1 ([18]). $\mathcal{T}_1 \oplus \mathcal{T}_2$ is decidable, provided that $\mathcal{T}_2$ is polite.

Next: strong politeness. Given a finite set of variables V and an equivalence relation E on V, the *arrangement* induced by E on V, denoted by δ_V^E or δ_V if E is clear, is the conjunction, for $x, y \in V$, of all formulas $x = y$, if xEy, or $\neg(x = y)$ otherwise. $\mathcal{T}$ is *strongly finitely witnessable* if it is finitely witnessable, with witness *wit*, which in addition satisfies: (II') for every quantifier-free formula φ, finite set of variables V, and arrangement δ_V on V, if $wit(\varphi) \wedge \delta_V$ is $\mathcal{T}$-satisfiable then there exists a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies $wit(\varphi) \wedge \delta_V$ with $dom(\mathcal{A}) = vars(wit(\varphi) \wedge \delta_V)^{\mathcal{A}}$. In that case *wit* is called a strong witness. $\mathcal{T}$ is *strongly polite* if it is smooth and strongly finitely witnessable (it was shown in [16, Theorem 2] that in this definition, smoothness can be replaced by stable infiniteness).

Theorem 2 ([11]). *$\mathcal{T}_1 \oplus \mathcal{T}_2$ is decidable, provided that $\mathcal{T}_2$ is strongly polite.*

We continue to shiny combination. $\mathcal{T}$ has the *finite model property* if, for every quantifier-free $\mathcal{T}$-satisfiable formula φ[4] and $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies φ, there exists a $\mathcal{T}$-interpretation $\mathcal{B}$ that satisfies φ with $|dom(\mathcal{B})| < \aleph_0$. Let $\mathbb{N}_\omega = \mathbb{N} \cup \{\aleph_0\}$. The *minimal model function* **minmod** of $\mathcal{T}$ is a function from the quantifier-free formulas of Σ to $\mathbb{N}_\omega$ such that for every quantifier-free $\mathcal{T}$-satisfiable formula φ: (I) there exists a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies φ with $|dom(\mathcal{A})| = \mathbf{minmod}(\varphi)$; (II) if $\mathcal{B}$ is a $\mathcal{T}$-interpretation that satisfies φ, $\mathbf{minmod}(\varphi) \leq |dom(\mathcal{B})|$.[5] $\mathcal{T}$ is *shiny* if it is smooth, and it has both the finite model property and a computable minimal model function. Note that [6,7]

[4] This notion is often not restricted to a quantifier-free φ, but in SMT it usually is.

[5] The function **minmod** is only guaranteed to exist if $\mathcal{F}_\Sigma \cup \mathcal{P}_\Sigma$ is countable.

showed that shininess is equivalent to strong politeness for decidable theories. In our context, however, we disassemble these notions to their more rudimentary ingredients, and when doing so, the equivalence does not necessarily hold.

Theorem 3 **([23]).** *$\mathcal{T}_1 \oplus \mathcal{T}_2$ is decidable, provided that $\mathcal{T}_2$ is shiny.*

The *spectrum* $Spec(\mathcal{T}, \varphi)$ of $\mathcal{T}$ w.r.t. a quantifier-free formula φ is the set of countable cardinalities of $\mathcal{T}$-interpretations that satisfy φ. Roughly, $\mathcal{T}$ is gentle if, given a conjunction φ of literals, $Spec(\mathcal{T}, \varphi)$ can be computed, and is either a finite set of finite cardinalities or a co-finite[6] set of cardinalities. Formally, $\mathcal{T}$ is *gentle* if there is an algorithm that, for every conjunction φ of literals, outputs a pair (b, S), with b a boolean and $S \subset \mathbb{N}^*$ is finite, such that (i) if b is true, then $Spec(\mathcal{T}, \varphi) = S$ and (ii) if b is false, then $Spec(\mathcal{T}, \varphi) = \mathbb{N}_\omega \setminus S$. Note that a gentle theory is decidable, because φ is $\mathcal{T}$-satisfiable if and only if $Spec(\mathcal{T}, \varphi) \neq \emptyset$.

Theorem 4 **([10, Theorem 3]).** *$\mathcal{T}_1 \oplus \mathcal{T}_2$ is decidable, when $\mathcal{T}_1$ is gentle, and $\mathcal{T}_2$ is either: (i) gentle, (ii) finitely axiomatizable, or (iii) there is an algorithm that, for a conjunction φ of Σ_2-literals, outputs a finite $S \subset \mathbb{N}_\omega$ with $Spec(\mathcal{T}, \varphi) = S$.*[7]

3 Limitations of theory combination methods

In this section, we examine what is the outcome of dropping each assumption on the theories from Theorems 1 to 4. We show that each of these theorems break if we drop any of the assumptions it makes regarding the combined theories.

All of these theorems have the following form: if $\mathcal{T}_1$ and $\mathcal{T}_2$ are over disjoint signatures and are decidable, and in addition, $\mathcal{T}_1$ admits some properties, and $\mathcal{T}_2$ admits some properties, then $\mathcal{T}_1 \oplus \mathcal{T}_2$ is also decidable. Thus, our limitation proofs always consist of examples for theories $\mathcal{T}_1$ and $\mathcal{T}_2$ that admit all but one of the properties, such that $\mathcal{T}_1 \oplus \mathcal{T}_2$ is undecidable.

In Section 3.2 we show that the Nelson–Oppen combination method fails if we drop the requirement of stable infiniteness from one of the theories. We show a similar result for gentle combination. In Section 3.3, dedicated to polite combination, we show that it fails if we drop any of the requirements for polite combination from $\mathcal{T}_2$. This includes dropping strong finite witnessability in exchange for finite witnessability, namely, replacing strong politeness by politeness. We also show that dropping smoothness from the polite combination method results in failure. In Section 3.4, a similar investigation is carried out for shiny theories. For each of the three components of shininess, we show that it is critical for the possibility of combination.

But first, we introduce the theories that will be used to demonstrate the limits of the various combination theorems in Section 3.1.

---

[6] A set S is co-finite if its complement $\mathbb{N} \setminus S$ is finite.
[7] Notice that in (i), $S \subseteq \mathbb{N}^*$, and in (ii) $S \subseteq \mathbb{N}_\omega$.

Name	Function Symbols	Predicate Symbols
Σ_1	$\emptyset$	$\emptyset$
Σ_s	$\{s\}$	$\emptyset$
Σ_P^n	$\emptyset$	$\{P_n \mid n \in \mathbb{N}^*\}$

Table 1: Signatures. Predicate symbols are 0-ary. The function symbol is unary.

3.1 The theories that we use

Since the Nelson–Oppen and gentle methods require one property each (stable infiniteness and gentleness, respectively), the polite method requires 2 properties (smoothness and strong finite witnessability), and the shiny method requires 3 properties (smoothness, the finite model property, and the computability of the minimal model function), we have $1+1+2+3=7$ variants to consider, each removing exactly one property as an assumption from a combination theorem. For each such variant, we need to provide 2 theories, $\mathcal{T}_1$ and $\mathcal{T}_2$ for which the variant fails. So, in total, we need to produce $7 \cdot 2 = 14$ theories as examples.

Remarkably, we are able to cover all the aforementioned variants using only 4 theories, that are defined over 3 signatures. Out of these 4 theories, only 2 are used for all but one of the variants. From these 4 theories, we create 3 ordered pairs of theories ($\mathcal{T}_1$ and $\mathcal{T}_2$). Two pairs are used to show the limits of 3 combination approaches each, and the third pair is used for one limit. Clearly, even if we were only concerned with shininess, 3 distinct ordered pairs would have been necessary, as there are three properties to exclude. Thus, the number of pairs of theories that we present is optimal.

The signatures for the theories are described in Table 1. Σ_1 is simply the empty signature. Atomic formulas are therefore only equalities between variables. Σ_s has a unary function symbol s. And Σ_P^n has infinitely many 0-ary predicate symbols $P_1, P_2, \ldots$.

The 4 theories are described in Table 2. The first two are taken from [25], which introduced and studied a wide collection of theories. $\mathcal{T}_\infty$ is the theory over the empty signature whose models have infinitely many elements.

$\mathcal{T}_f$ is more involved. Its axiomatization as a Σ_s-theory assumes the existence of a non-computable function $f : \mathbb{N}^* \to \{0,1\}$, such that $f(1) = 1$, and for every $k \geq 1$, f maps half of the numbers between 1 and 2^k to 1, and the other half to 0. Such a function was proven to exist in [25, Lemma 6]. The axiomatization utilizes two derived functions: $f_0(k)$ returns the number of numbers between 1 and k that f maps to 0, while $f_1(k)$ returns the number of numbers between 1 and k that f maps to 1. Obviously, when k is a power of 2, then $f_0(k) = f_1(k)$. Now, f itself is not a part of the signature Σ_s of $\mathcal{T}_f$. Instead, the axiomatization relies on the formulas from Figure 3, that involve counting elements for which the function symbol s acts as the identity. Intuitively, a finite $\mathcal{T}_f$-interpretation $\mathcal{A}$ with n elements has $f_0(n)$ of them satisfying $s^{\mathcal{A}}(e) \neq e$, and $f_1(n)$ satisfying $s^{\mathcal{A}}(e) = e$; an infinite such interpretation has infinitely many elements of each kind.

The definition of $\mathcal{T}_{\leq}$ assumes an arbitrary non-computable function $F : \mathbb{N}^* \to \mathbb{N}^* \cup \{\aleph_0\}$ such that the set $\{(m,n) \in \mathbb{N}^* \times \mathbb{N}^* \mid F(m) \geq n\}$ is decidable. Such

$$\psi^{=}_{\geq n} = \exists \overrightarrow{x}. [\neq (x_1, \dots, x_n) \wedge \bigwedge_{i=1}^{n} p(x_i)], \qquad \psi^{\neq}_{\geq n} = \exists \overrightarrow{x}. [\neq (x_1, \dots, x_n) \wedge \bigwedge_{i=1}^{n} \neg p(x_i)],$$

$$\psi^{=}_{=n} = \exists \overrightarrow{x}. [\neq (x_1, \dots, x_n) \wedge \bigwedge_{i=1}^{n} p(x_i) \wedge \forall x. [p(x) \rightarrow \bigvee_{i=1}^{n} x = x_i]],$$

$$\psi^{\neq}_{=n} = \exists \overrightarrow{x}. [\neq (x_1, \dots, x_n) \wedge \bigwedge_{i=1}^{n} \neg p(x_i) \wedge \forall x. [\neg p(x) \rightarrow \bigvee_{i=1}^{n} x = x_i]].$$

Fig. 3: Formulas for the axiomatization of $\mathcal{T}_f$. $\overrightarrow{x}$ stands for $x_1, \dots, x_n$, and $p(x)$ for $s(x) = x$.

Name	Signature	Axiomatization	Source
$\mathcal{T}_\infty$	Σ_1	$\{\psi_{\geq n} : n \in \mathbb{N}^*\}$	[25]
$\mathcal{T}_f$	Σ_s	$\{[\psi^{=}_{\geq f_1(k)} \wedge \psi^{\neq}_{\geq f_0(k)}] \vee \bigvee_{i=1}^{k} [\psi^{=}_{=f_1(i)} \wedge \psi^{\neq}_{=f_0(i)}] : k \in \mathbb{N}^*\}$	[25]
$\mathcal{T}_\leq$	Σ^n_P	$\{P_n \rightarrow \psi_{\leq F(n)} : n \in \mathbb{N}^*, F(n) \in \mathbb{N}^*\}$	new
$\mathcal{T}_=$	Σ^n_P	$\{P_n \rightarrow \psi_{=n} : n \in \mathbb{N}^*\}$	new

Table 2: Theories. $f : \mathbb{N}^* \rightarrow \{0, 1\}$ is assumed to be a non-computable function, such that $f(1) = 1$ and, for every $k \geq 0$, f maps half of the numbers between 1 and 2^k to 1, and the other half to 0. $f_i(k)$ is the number of numbers between 1 and k that are mapped by f to i. $F : \mathbb{N}^* \rightarrow \mathbb{N}^* \cup \{\aleph_0\}$ is non-computable, but the set $\{(m, n) \mid F(m) \geq n\}$ is decidable. Formulas from Figure 3 are used.

a function F exists: for example, suppose F maps every $n \in \mathbb{N}^*$ to the number of steps the nth Turing machine (under some encoding) takes to halt, returning $\aleph_0$ if it does not halt. This function is clearly not computable. But, given m and n, we can decide whether $F(m) \geq n$ by executing the mth Turing machine for n steps. If a $\mathcal{T}_\leq$-interpretation $\mathcal{A}$ satisfies P_n, then it has at most $F(n)$ elements.[8]

Finally, $\mathcal{T}_=$ consists of all Σ^n_P-interpretations $\mathcal{A}$ in which for all $n \in \mathbb{N}^*$, either P_n is interpreted as false, or $|\mathcal{A}| = n$. It therefore allows quantifier-free formulas to enforce finite sizes of models, as P_n being true implies the model has n elements.

3.2 Nelson–Oppen and gentle combination

We begin by proving the sharpness of Theorem 1 in the following sense: although two theories can be combined if both are stably-infinite, this is no longer the case if only one has that property. This result was previously proven in [5, Theorem 4.1], but with a different proof.

Theorem 5. *There are decidable theories $\mathcal{T}_1$ and $\mathcal{T}_2$ over disjoint signatures such that $\mathcal{T}_1$ is stably infinite but $\mathcal{T}_1 \oplus \mathcal{T}_2$ is undecidable.*

[8] $\mathcal{T}_\leq$ generalizes the theory TM_∞ from [5].

Proof (sketch). [9] Take $\mathcal{T}_1$ and $\mathcal{T}_2$ to be $\mathcal{T}_f$ and $\mathcal{T}_=$, respectively. Clearly, their signatures (namely Σ_s and Σ_P^n) are disjoint. Further, $\mathcal{T}_f$ is shown in [9, Lemma 54] to be stably infinite; it was also proven to have the same set of quantifier-free satisfiable formulas as the theory of an uninterpreted unary function, which makes it decidable. Finally, although both $\mathcal{T}_f$ and $\mathcal{T}_=$ are decidable, $\mathcal{T}_f \oplus \mathcal{T}_=$ is not. Indeed, the formulas $P_{n+1} \wedge \varphi^{=}_{\geq f_1(n)+1}$, where

$$\varphi^{=}_{\geq n} \coloneqq \bigwedge_{1 \leq i < j \leq n} \neg(x_i = x_j) \wedge \bigwedge_{i=1}^{n} s(x_i) = x_i,$$

are $\mathcal{T}_= \oplus \mathcal{T}_f$-satisfiable if and only if $f(n+1) = 1$, whereas f is a non-computable function. □

Clearly, Theorem 1 and the proof of Theorem 5 imply that $\mathcal{T}_=$ is not stably infinite. And indeed, for every n, the formula P_n is $\mathcal{T}_=$-satisfiable, but only by a finite model.

As it turns out, the same theories can be used to show a similar result for gentleness.

Theorem 6. *There are decidable theories $\mathcal{T}_1$ and $\mathcal{T}_2$ over disjoint signatures such that $\mathcal{T}_1$ is gentle, but $\mathcal{T}_1 \oplus \mathcal{T}_2$ is undecidable.*

Proof (sketch). We reuse the proof of Theorem 5, but flip the roles of the theories. Now, we set $\mathcal{T}_1$ to be $\mathcal{T}_=$ and $\mathcal{T}_2$ to be $\mathcal{T}_f$. $\mathcal{T}_f$ and $\mathcal{T}_=$ are both decidable, are over disjoint signatures, but $\mathcal{T}_= \oplus \mathcal{T}_f$ is undecidable. The only thing left to show is that $\mathcal{T}_=$ is gentle, which indeed can be shown. □

Theorem 4 and the proof of Theorem 6 tell us that $\mathcal{T}_f$ is not gentle. And indeed, were $\mathcal{T}_f$ gentle, one would be able to calculate f. Similarly, $\mathcal{T}_f$ does not satisfy any of the other two requirements from Theorem 4.

3.3 Polite combination

Theorem 2 demands two properties from $\mathcal{T}_2$ in order for it to be combinable with any decidable theory $\mathcal{T}_1$ over a disjoint signature: strong finite witnessability and smoothness. We start by showing that if smoothness is removed from the requirements, the theorem fails.

Theorem 7. *There are decidable theories $\mathcal{T}_1$ and $\mathcal{T}_2$ over disjoint signatures such that $\mathcal{T}_2$ is strongly finitely witnessable but $\mathcal{T}_1 \oplus \mathcal{T}_2$ is undecidable.*

Proof (sketch). Take $\mathcal{T}_1$ to be $\mathcal{T}_f$ and $\mathcal{T}_2$ to be $\mathcal{T}_=$, as was done in the proof of Theorem 5, where both theories were shown to be decidable while their combination was shown to be undecidable. The only thing that is left to be shown, and indeed can be shown by providing an appropriate strong witness, is that $\mathcal{T}_=$ is strongly finitely witnessable. □

[9] Due to lack of space, some proofs are omitted, and can be found in the appendix.

As before, Theorem 2 and the proof of Theorem 7 imply that $\mathcal{T}_{=}$ is not smooth. And indeed, it is not, as it is not even stably infinite.

Next, we show that dropping the strong finite witnessability requirement also leads to a failure in the polite combination method.

Theorem 8. *There are decidable theories $\mathcal{T}_1$ and $\mathcal{T}_2$ over disjoint signatures such that $\mathcal{T}_2$ is smooth but $\mathcal{T}_1 \oplus \mathcal{T}_2$ is undecidable.*

Proof. Take $\mathcal{T}_1$ to be $\mathcal{T}_{=}$ and $\mathcal{T}_2$ to be $\mathcal{T}_f$, again as in the proof of Theorem 6, only now we rely on the fact that, proven in [9, Lemma 54], that $\mathcal{T}_f$ is smooth, and as we already know, $\mathcal{T}_{=} \oplus \mathcal{T}_f$ is not decidable. □

Clearly, Theorem 2 and the proof of Corollary 1 imply that $\mathcal{T}_f$ is not strongly finitely witnessable. This was also proven in [9, Lemma 56].

Now, $\mathcal{T}_f$ was proven in [9, Lemmas 55] to not be smooth, but it is also finitely witnessable (without being strongly finitely witnessable), which makes it polite. Thus, the proof of Theorem 8 also gives us the following corollary, by again taking $\mathcal{T}_1$ to be $\mathcal{T}_{=}$ and $\mathcal{T}_2$ to be $\mathcal{T}_f$.

Corollary 1. *There are decidable theories $\mathcal{T}_1$ and $\mathcal{T}_2$ over disjoint signatures such that $\mathcal{T}_2$ is polite but $\mathcal{T}_1 \oplus \mathcal{T}_2$ is undecidable.*

Recall that [17] claimed that politeness is enough for theory combination, but a problem in the proof was later discovered and corrected in [11] by strengthening the politeness assumption to strong politeness. But was the problem of [17] in the proof or in the statement itself? In other words: does Conjecture 1 hold? What we immediately get from Corollary 1 is that it does not.

Corollary 2. *Conjecture 1 does not hold.*

Hence, politeness is not enough for theory combination, which justifies the title of this paper.

3.4 Shiny combination

In this section we consider the three requirements Theorem 3 makes on one of the combined theories, namely: computability of the minimal model function, the finite model property, and smoothness.

We start with the computability of the minimal model function.

Theorem 9. *There are decidable theories $\mathcal{T}_1$ and $\mathcal{T}_2$ over disjoint signatures such that $\mathcal{T}_2$ is smooth and has the finite model property, but $\mathcal{T}_1 \oplus \mathcal{T}_2$ is undecidable.*

Proof. By taking $\mathcal{T}_1$ to be $\mathcal{T}_{=}$ and $\mathcal{T}_2$ to be $\mathcal{T}_f$, we can use proofs of previous theorems in order to show most properties that are needed. Further, it was proven in [26, Theorem 2] that $\mathcal{T}_f$ admits the finite model property. □

From Theorem 3 and the proof of Theorem 9, $\mathcal{T}_f$ does not have a computable minimal model function, which was also proven in [8, Lemma 126].

For the next sharpness theorem we need the following lemma, according to which for decidable theories, strong finite witnessability implies computability of the minimal model function. This was essentially proven in [7], but was never explicitly stated there; indeed, as they were focused on strong politeness and shininess, they have assumed smoothness, even if that assumption was never actually used in the part of the proof that concerned the computability of the minimal model function.

Lemma 1. *If $\mathcal{T}$ is decidable and strongly finitely witnessable, then it has a computable minimal model function.*

Remark 1. Notice that the reciprocal of Lemma 1 is not true: decidability and computability of the minimal model function do not entail strong finite witnessability. For example, $\mathcal{T}^{\infty}_{even}$, defined in [21] by the axiomatization $\{\neg\psi_{=2\cdot n+1} \mid n \in \mathbb{N}\}$, is proven in [8, Lemma 126] to have a computable minimal model function; furthermore it is decidable, as it satisfies all and only the quantifier-free formulas that are satisfiable in first-order logic, but it is not strongly finitely witnessable (as proven in [21]).

Now, using Lemma 1, we show that shiny combination (Theorem 3) fails without the smoothness requirement. We once again essentially reuse Theorem 5 to obtain the following:

Theorem 10. *There are decidable theories $\mathcal{T}_1$ and $\mathcal{T}_2$ over disjoint signatures such that $\mathcal{T}_2$ has the finite model property and a computable minimal model function, but $\mathcal{T}_1 \oplus \mathcal{T}_2$ is undecidable.*

Proof. Take $\mathcal{T}_1$ and $\mathcal{T}_2$ to be $\mathcal{T}_f$ and $\mathcal{T}_=$, respectively: we have already shown that they are both decidable even though $\mathcal{T}_f \oplus \mathcal{T}_=$ is not. From Lemma 1 and the fact that $\mathcal{T}_=$ is strongly finitely witnessable (which was established in the proof of Theorem 7), we get $\mathcal{T}_=$ has a computable minimal model function. Using then [26, Theorem 2], according to which finite witnessability implies the finite model property, $\mathcal{T}_=$ has the finite model property. □

Next, we show that the requirement of the finite model property cannot be removed. Unlike the previous results, we are unable to reuse $\mathcal{T}_=$ and $\mathcal{T}_f$. Therefore, we use the theory $\mathcal{T}_\infty$ from [25] and the theory $\mathcal{T}_\leq$.

Theorem 11. *There are decidable theories $\mathcal{T}_1$ and $\mathcal{T}_2$ over disjoint signatures such that $\mathcal{T}_2$ is smooth and has a computable minimal model function, but $\mathcal{T}_1 \oplus \mathcal{T}_2$ is undecidable.*

Proof (sketch). Take $\mathcal{T}_1$ to be $\mathcal{T}_\leq$ and $\mathcal{T}_2$ to be $\mathcal{T}_\infty$, axiomatized in Table 2. Clearly, they are defined over disjoint signatures. $\mathcal{T}_\infty$ is smooth and has a computable minimal model function. The proofs for these facts are simple, and are given in [9, Lemma 22] and [8, Lemma 130]. It is also decidable, as it satisfies all quantifier-free formulas in its signature that are satisfiable in first-order logic (and only them). Perhaps surprisingly, it is possible to show that $\mathcal{T}_\leq$ is decidable. However, it can also be shown that $\mathcal{T}_\leq \oplus \mathcal{T}_\infty$ is not. □

Name	0-ary Functions	1-ary Functions	Predicates
Σ_t^a	$\{a\}$	$\{t\}$	$\emptyset$

Table 3: A finite signature.

$$\psi^{orb}_{\geq n}(x) = \bigwedge_{0\leq i<j\leq n-1} \neg(t^i(x) = t^j(x)) \text{ for } n \in \mathbb{N}^* \setminus \{1\}$$

$$\psi^{orb}_{=1}(x) = t(x) = x$$

$$\psi^{orb}_{=n}(x) = \psi^{orb}_{\geq n}(x) \wedge \neg\psi^{orb}_{\geq n+1}(x) \text{ for } n \geq 2$$

Fig. 4: Formulas in Σ_t^a.

4 Finite signatures

Every proof in Section 3 uses a pair of theories, one of them always over the infinite signature Σ_P^n (in all cases the used theory is $\mathcal{T}_=$, except for in Theorem 11, where the Σ_P^n-theory $\mathcal{T}_\leq$ is used instead). And indeed, both $\mathcal{T}_=$ and $\mathcal{T}_\leq$ are theories that are relatively easy to understand. This is, among other things, thanks to the availability of infinitely many predicates.

In this section, we aim to provide finitistic proofs of the limitation theorems from Section 3, in the sense that all theories that are used are over finite signatures. Doing so provides a more succinct set of examples, over more minimal signatures. The cost, however, is that the theories that we use in this section are more complex.

4.1 New theories over finite signatures

The two theories over the infinite signature Σ_P^n from Section 3 are $\mathcal{T}_=$ and $\mathcal{T}_\leq$. They will be replaced by theories over the finite signature Σ_t^a: this signature has a unary function t, a constant a, and no predicates, as described in Table 3.

In order to introduce the new theories, we define the formulas $\psi^{orb}_{\geq n}(x)$ and $\psi^{orb}_{=n}(x)$ in Figure 4.

The *orbit* (see e.g., [20,12]) of an element e in a Σ_t^a-interpretation $\mathcal{A}$ is the set $\{(t^{\mathcal{A}})^n(e) \mid n \in \mathbb{N}\}$. Since e itself is always an element of this set, the orbit is always non-empty. We sometimes view this set as the following sequence indexed by n: $e, t^{\mathcal{A}}(e), (t^{\mathcal{A}})^2(e), \ldots$ In this context, in an interpretation $\mathcal{A}$ that satisfies $\psi^{orb}_{\geq n}(x)$, we have that there are at least n elements that can be obtained by recursively applying $t^{\mathcal{A}}$ to $x^{\mathcal{A}}$, meaning its orbit has at least n elements; similar, if $\mathcal{A}$ satisfies $\psi^{orb}_{=n}(x)$, the orbit of $x^{\mathcal{A}}$ has precisely n elements.

With these formulas, we can now define the new theories over Σ_t^a. These are specified in Table 4. In the finite $\mathcal{T}^2_{orb}$-interpretations, the orbit of the interpretation of the constant symbol a consists of at least half of the elements of the

Name	Signature	Axiomatization	Source
$\mathcal{T}^2_{orb}$	Σ^a_t	$\{\psi^{orb}_{=n}(a) \rightarrow \psi_{\leq 2n} : n \in \mathbb{N}^*\}$	new
$\mathcal{T}^{orb}_{\leq}$	Σ^a_t	$\{\psi^{orb}_{=n}(a) \rightarrow \psi_{\leq F(n)+n} : n \in \mathbb{N}^*,\ F(n) \in \mathbb{N}^*\}$	new

Table 4: Theories over the finite signature Σ^a_t. In the definition of $\mathcal{T}^{orb}_{\leq}$, the function F admits the same assumptions as in Table 2. The axiomatizations utilize formulas that are defined in Figure 4.

interpretation; meanwhile in the infinite $\mathcal{T}^2_{orb}$-interpretations this orbit is infinite. $\mathcal{T}^{orb}_{\leq}$ is very similar to $\mathcal{T}_{\leq}$, replacing P_n by the assumption that the orbit has size n, and also concluding that the number of elements in the domain is at most $F(n) + n$ (and not $F(n)$ as in $\mathcal{T}_{\leq}$).

With these new theories, we can now turn to making the proofs of the theorems from Section 3 rely solely on finite signatures.

4.2 Finitizing the proofs of Theorems 5 to 10 and Corollary 1

The proof of Theorem 5 sets $\mathcal{T}_1$ to be $\mathcal{T}_f$ and $\mathcal{T}_2$ to be $\mathcal{T}_=$. In order to only use finite signatures, we set $\mathcal{T}_2$ to be $\mathcal{T}^2_{orb}$ instead.[10]

As for $\mathcal{T}_1$, we can still use $\mathcal{T}_f$, but we need to restrict the possible functions f it relies on. To make it clear that the functions f are now required to satisfy some extra properties we denote them by g, so that $\mathcal{T}_f$ becomes $\mathcal{T}_g$. We then require $g : \mathbb{N}^* \rightarrow \{0, 1\}$ to be any non-computable function such that: $g(1) = 1$ and g is zero as often as it is 1 in each interval from 1 to 2^k (as required for f in the definition of $\mathcal{T}_f$); and in addition to the requirements in $\mathcal{T}_f$, we now also require that $g(2n + 1) = g(2n + 2)$ for all $n \geq 2$. Such functions exist: for an example, take the function $f : \mathbb{N}^* \rightarrow \{0, 1\}$ defined in [25], make $g(1) = g(3) = 1$, $g(2) = g(4) = 0$, and $g(2n + 1) = g(2n + 2) = f(n + 1)$ for $n \geq 2$. Since $\mathcal{T}_f$ is decidable and stably infinite regardless of the specific f, we have $\mathcal{T}_g$ is decidable and stably infinite.

Although it can be shown that $\mathcal{T}^2_{orb}$ is decidable, $\mathcal{T}_g \oplus \mathcal{T}^2_{orb}$ is undecidable. Indeed, were it decidable, one would be able to calculate the function g by using the fact that $\psi^{orb}_{=n+1}(a) \wedge \varphi^{=}_{\geq g_1(2n)+2}$ is $\mathcal{T}_g \oplus \mathcal{T}^2_{orb}$-satisfiable if and only if $g(2n + 1) = g(2n + 2) = 1$ for $n \geq 2$: if we know g up to $2n$ we can calculate $g_1(2n)$, obtain the formula $\psi^{orb}_{=n+1}(a) \wedge \varphi^{=}_{\geq g_1(2n)+2}$, and by testing whether it is $\mathcal{T}_g \oplus \mathcal{T}^2_{orb}$-satisfiable we find the value for $g(2n + 1) = g(2n + 2)$. We know $g(1) = g(3) = 1$ and $g(2) = g(4) = 0$, and then we proceed from there on forward inductively.

We can mimic the same process for other results from Section 3, by replacing $\mathcal{T}_=$ by $\mathcal{T}^2_{orb}$, and instantiating $\mathcal{T}_f$ by $\mathcal{T}_g$ with f satisfying the aforementioned condition. In particular, we can do so in: Theorem 6, by proving that $\mathcal{T}^2_{orb}$ is

[10] Notice that [5] has also produced a finitary proof of Theorem 5, using a theory named $TM_{\forall\omega}$ over a finite signature.

also gentle; Theorem 7, by proving that $\mathcal{T}^2_{orb}$ is strongly finitely witnessable; Theorem 8, and Corollary 1, by remembering $\mathcal{T}_g$ is both smooth and polite; Theorem 9, as $\mathcal{T}_g$ has the finite model property; and Theorem 10, by proving that $\mathcal{T}^2_{orb}$ has a computable minimal model function, and the finite model property.

4.3 Finitizing the proof of Theorem 11

Finally, notice that the proof of Theorem 11 sets $\mathcal{T}_1$ to be $\mathcal{T}_{\leq}$, which is defined over an infinite signature, and $\mathcal{T}_2$ to be $\mathcal{T}_{\infty}$. While we can leave $\mathcal{T}_2$ as $\mathcal{T}_{\infty}$, we replace $\mathcal{T}_{\leq}$ by its Σ^a_t-variant $\mathcal{T}^{orb}_{\leq}$, in order to get two theories over a finite signature.

The proof that $\mathcal{T}^{orb}_{\leq}$ is decidable follows the proof that $\mathcal{T}^2_{orb}$ is decidable. Yet the combination $\mathcal{T}^{orb}_{\leq} \oplus \mathcal{T}_{\infty}$ is not decidable, $\psi^{orb}_{=n}(a)$ being satisfiable in it if, and only if, $F(n) = \aleph_0$.

5 New combination theorems

In this section we prove new combination theorems, that strengthen Theorems 3 and 4. In Section 5.1, we show that the conditions from Theorem 4 can be weakened. In Section 5.2, we show that the finite model property *can* be dismissed from shiny combination, as long as we compensate it by requiring another property from the second theory being combined. This does not contradict Theorem 11, as the example there does not meet the additional criterion.

In what follows, we assume that Σ_1 and Σ_2 are disjoint signatures, and that $\mathcal{T}_1$ is a Σ_1-theory, and $\mathcal{T}_2$ a Σ_2-theory.

5.1 Recovering gentle combination

Assuming $\mathcal{T}_1$ is gentle, Theorem 4 provided three conditions on $\mathcal{T}_2$, any one of which suffices for theory combination. We prove a strengthening of Theorem 4.

Definition 1. *We say that a theory $\mathcal{T}$ has* computable finite spectra *if there is an algorithm that, given a quantifier-free formula φ and $k \in \mathbb{N}^*$, decides whether $k \in Spec(\mathcal{T}, \varphi)$.*

Intuitively, having computable finite spectra means that we can query the set $Spec(\mathcal{T}, \varphi)$ to check whether it contains a given *finite* cardinality. In contrast to gentleness, it does not imply that we can compute any concrete set S, nor does it require the ability to check whether $\aleph_0$ is in the spectra.

Theorem 12. *Suppose that $\mathcal{T}_1$ is gentle and $\mathcal{T}_2$ has computable finite spectra. Then, $\mathcal{T}_1 \oplus \mathcal{T}_2$ is decidable.*

Each of the three properties in Theorem 4 imply that $\mathcal{T}_2$ has computable finite spectra, so Theorem 12 is indeed a strengthening. We now present two theories that can be combined by Theorem 12 but not by any other combination method discussed in this paper.

Example 1. Fix any $n \in \mathbb{N}^*$ and let $\mathcal{T}_1 = \mathcal{T}_{\leq n}$ be the Σ_1-theory axiomatized by $\{\psi_{\leq n}\}$. Now, let $\mathcal{T}_2$ be $\mathcal{T}_{\leq}$ from Table 2. Then, $\mathcal{T}_{\leq n}$ and $\mathcal{T}_{\leq}$ are decidable, $\mathcal{T}_{\leq n}$ is gentle, and $\mathcal{T}_{\leq}$ has computable finite spectra. By Theorem 12, $\mathcal{T}_{\leq n} \oplus \mathcal{T}_{\leq}$ is decidable. On the other hand, $\mathcal{T}_{\leq}$ does not satisfy any of the three properties in Theorem 4. Furthermore, neither theory is strongly polite, shiny, or stably infinite, and so none of the other combination theorems can be used to decide this combination of theories.

5.2 Recovering shiny combination without finite models

In Theorem 11 we have seen that the shiny combination theorem fails if the finite model property is dropped from the definition of shininess. However, we now show that we can do without the finite model property if we impose another condition on the other theory being combined.

Definition 2. *We say that a theory $\mathcal{T}$ is* infinitely decidable *if it is decidable whether a quantifier-free formula is satisfied by an infinite $\mathcal{T}$-interpretation.*

A very similar notion to infinite decidability, that also requires the theory to be decidable, was defined in [5] and called $\exists_\infty$-decidability, but not considered along gentleness.

Theorem 13. *Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be decidable theories over disjoint signatures. Suppose that $\mathcal{T}_1$ is smooth and has a computable minimal model function and that $\mathcal{T}_2$ is infinitely decidable. Then, $\mathcal{T}_1 \oplus \mathcal{T}_2$ is decidable.*

In the next example, we present 2 theories that can be combined using Theorem 13, but not with any other combination theorem studied in this paper.

Example 2. Let $\mathcal{T}_\infty^h$ be the Σ_P^n-theory axiomatized by

$$\{P_1 \rightarrow \psi_{=1}\} \cup \{P_1 \rightarrow \neg P_n : n \geq 2\} \cup \{P_n \rightarrow \psi_{\geq m} : m, n \geq 2, h(n) = 1\},$$

for $h : \mathbb{N}^* \rightarrow \{0, 1\}$ a non-computable function. Also, consider the theory $\mathcal{T}_\infty$ from Table 2. Both theories are decidable. It can be shown that $\mathcal{T}_\infty^h$ is neither stably infinite nor has computable finite spectra, and so it cannot be combined with $\mathcal{T}_\infty$ using the Nelson–Oppen method, the gentle method, or the new method we propose in Theorem 12. It can also be shown that neither theory is strongly polite or shiny, and so they cannot be combined using the polite or shiny methods. But, $\mathcal{T}_\infty$ has a computable minimal model function and is smooth. Further, $\mathcal{T}_\infty^h$ is infinitely decidable. By Theorem 13, $\mathcal{T}_\infty \oplus \mathcal{T}_\infty^h$ is decidable. Thus, Theorem 13 is able to combine two theories that none of the other methods can.[11]

[11] We use $\mathcal{T}_\infty$ in the example to keep things simple, reusing the theories that are already defined in the paper. However, any decidable theory $\mathcal{T}$ on a countable signature (disjoint from Σ_P^n) with only infinite models could replace $\mathcal{T}_\infty$, such as the theory of dense linear orders without endpoints [13].

Approach	Property	Theorem	Infinite $\mathcal{T}_1$	Infinite $\mathcal{T}_2$	Finite $\mathcal{T}_1$	Finite $\mathcal{T}_2$
Nelson–Oppen (Thm. 1)	Stable Infiniteness	Thm. 5	$\mathcal{T}_f$	$\mathcal{T}_=$	$\mathcal{T}_g$	$\mathcal{T}^2_{orb}$
Gentle (Thm. 4)	Gentleness*	Thm. 6	$\mathcal{T}_=$	$\mathcal{T}_f$	$\mathcal{T}^2_{orb}$	$\mathcal{T}_g$
Polite (Thm. 2)	Smoothness	Thm. 7	$\mathcal{T}_f$	$\mathcal{T}_=$	$\mathcal{T}_g$	$\mathcal{T}^2_{orb}$
	Strong Finite Witnessability**	Thm. 8	$\mathcal{T}_=$	$\mathcal{T}_f$	$\mathcal{T}^2_{orb}$	$\mathcal{T}_g$
Shiny (Thm. 3)	Comp. Min. Mod.	Thm. 9	$\mathcal{T}_=$	$\mathcal{T}_f$	$\mathcal{T}^2_{orb}$	$\mathcal{T}_g$
	Smoothness	Thm. 10	$\mathcal{T}_f$	$\mathcal{T}_=$	$\mathcal{T}_g$	$\mathcal{T}^2_{orb}$
	Finite Model Property*	Thm. 11	$\mathcal{T}_\leq$	$\mathcal{T}_\infty$	$\mathcal{T}^{orb}_\leq$	$\mathcal{T}_\infty$

Table 5: Summary of the main results. Only 3 tuples of theories were used, and these are assigned different colors in the table.

* Theorems 12 and 13, in a sense, remedy Theorems 6 and 11.
** In particular, Corollary 1 is a consequence of the proof of Thm. 8.

6 Conclusion

For each combination method and each of its associated properties, we have proven in Section 3 that the corresponding combination theorem fails if the property is not assumed. The proofs always involve producing two theories that are decidable while their combination is not. The proofs of these results were improved in Section 4, where only finite signatures were used.

Table 5 lists the theories used in Sections 3 and 4. It also lists the original combination theorem whose limits are identified. Notice that for each theorem we produced 2 pairs of theories: one pair for its original proof, and another pair for its improved proof. In total, we were able to prove all theorems, with finite and infinite signatures, using only three quadruples of theories, built from only six theories, by reusing the introduced theories as much as possible.

We have also proven that politeness is not enough for theory combination. Further, we have introduced two new combination theorems, based on shiny and gentle combinations (Theorems 12 and 13).

The main direction for further work is to find more theorems like Theorems 12 and 13, with the purpose of varying the set of requirements for theory combination. We hope that such theorems will make it to introduce algorithms for new combination of theories. In addition, we are working on stronger limitation theorems: while the classical combination methods provide *sufficient* conditions for combinability, we plan to study *necessary* conditions.

Acknowledgments We thank Christophe Ringeissen, Pascal Fontaine, and Cesare Tinelli for fruitful discussions that led to and helped writing this paper.

References

1. Haniel Barbosa, Clark W. Barrett, Martin Brain, Gereon Kremer, Hanna Lachnitt, Makai Mann, Abdalrhman Mohamed, Mudathir Mohamed, Aina Niemetz, Andres Nötzli, Alex Ozdemir, Mathias Preiner, Andrew Reynolds, Ying Sheng, Cesare Tinelli, and Yoni Zohar. cvc5: A versatile and industrial-strength SMT solver. In *TACAS (1)*, volume 13243 of *Lecture Notes in Computer Science*, pages 415–442. Springer, 2022.
2. Clark W. Barrett, Christopher L. Conway, Morgan Deters, Liana Hadarean, Dejan Jovanovic, Tim King, Andrew Reynolds, and Cesare Tinelli. CVC4. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings*, volume 6806 of *Lecture Notes in Computer Science*, pages 171–177. Springer, 2011.
3. Clark W. Barrett and Cesare Tinelli. CVC3. In *CAV*, volume 4590 of *Lecture Notes in Computer Science*, pages 298–302. Springer, 2007.
4. Maria Paola Bonacina, Pascal Fontaine, Christophe Ringeissen, and Cesare Tinelli. Theory combination: Beyond equality sharing. In Carsten Lutz, Uli Sattler, Cesare Tinelli, Anni-Yasmin Turhan, and Frank Wolter, editors, *Description Logic, Theory Combination, and All That - Essays Dedicated to Franz Baader on the Occasion of His 60th Birthday*, volume 11560 of *Lecture Notes in Computer Science*, pages 57–89. Springer, 2019.
5. Maria Paola Bonacina, Silvio Ghilardi, Enrica Nicolini, Silvio Ranise, and Daniele Zucchelli. Decidability and undecidability results for nelson-oppen and rewrite-based decision procedures. In Ulrich Furbach and Natarajan Shankar, editors, *Automated Reasoning*, pages 513–527, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
6. Filipe Casal and João Rasga. Revisiting the equivalence of shininess and politeness. In *LPAR*, volume 8312 of *Lecture Notes in Computer Science*, pages 198–212. Springer, 2013.
7. Filipe Casal and João Rasga. Many-sorted equivalence of shiny and strongly polite theories. *Journal of Automated Reasoning*, 60(2):221–236, February 2018.
8. Guilherme Vicentin de Toledo and Yoni Zohar. Combining combination properties: Minimal models, 2024.
9. Guilherme Vicentin de Toledo, Yoni Zohar, and Clark Barrett. Combining combination properties: An analysis of stable infiniteness, convexity, and politeness, 2023.
10. 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.
11. 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.
12. Mati Kilp, Ulrich Knauer, and Alexander V. Mikhalev. *Monoids, Acts and Categories: With Applications to Wreath Products and Graphs. A Handbook for Students and Researchers.* DE GRUYTER, December 2000.
13. Dexter Kozen. *Theory of Computation.* Texts in Computer Science. Springer, London, England, 2006 edition, March 2006.
14. Sava Krstić, Amit Goel, Jim Grundy, and Cesare Tinelli. Combined satisfiability modulo parametric theories. In O. Grumberg and M. Huth, editors, *Proceedings*

of the 13th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (Braga, Portugal), volume 4424 of *Lecture Notes in Computer Science*, pages 618–631. Springer, 2007.

15. Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Trans. Program. Lang. Syst.*, 1(2):245–257, October 1979.
16. Benjamin Przybocki, Guilherme Vicentin de Toledo, Yoni Zohar, and Clark W. Barrett. The nonexistence of unicorns and many-sorted löwenheim-skolem theorems. In André Platzer, Kristin Yvonne Rozier, Matteo Pradella, and Matteo Rossi, editors, *Formal Methods - 26th International Symposium, FM 2024, Milan, Italy, September 9-13, 2024, Proceedings, Part I*, volume 14933 of *Lecture Notes in Computer Science*, pages 658–675. Springer, 2024.
17. 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.
18. 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.
19. Yann-Joachim Ringard. Mustard watches: An integrated approach to time and food. *Equipe de logicue mathématique, prépublications*, 1990.
20. Joseph J. Rotman. *An Introduction to the Theory of Groups.* Springer New York, 1995.
21. 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.
22. Cesare Tinelli and Calogero Zarba. Combining decision procedures for theories in sorted logics. Technical Report 04-01, Department of Computer Science, The University of Iowa, February 2004.
23. Cesare Tinelli and Calogero G. Zarba. Combining nonstably infinite theories. *J. Autom. Reason.*, 34(3):209–238, 2005.
24. Guilherme V. Toledo and Yoni Zohar. Combining combination properties: Minimal models. In Nikolaj Bjorner, 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.

A Proof of Theorem 5

Proposition 1. *The theory $\mathcal{T}_=$ is gentle.*

Proof. Let φ be a conjunction of literals. Write $\varphi = \varphi_1 \wedge \varphi_2$, where φ_1 contains the equalities and disequalities in φ and φ_2 contains the literals of the form P_n and $\neg P_n$ in φ.

If φ_1 is unsatisfiable in equational logic, then $Spec(\mathcal{T}_=, \varphi) = \emptyset$; otherwise, let m be the size of the smallest interpretation that satisfies φ_1, which is possible since the theory of equality is shiny (see [23]).

If φ_2 contains no positive literals, then $Spec(\mathcal{T}_=, \varphi) = \{n \in \mathbb{N}^* \mid n \geq m\} \cup \{\aleph_0\}$. If φ_2 contains exactly one positive literal P_n, then $Spec(\mathcal{T}_=, \varphi) = \{n\}$. If φ_2 contains two positive literals P_n and $P_{n'}$ where $n \neq n'$, then $Spec(\mathcal{T}_=, \varphi) = \emptyset$.

Thus, we have shown that $Spec(\mathcal{T}_=, \varphi)$ is either a finite set of finite cardinalities that can be computed or a cofinite set whose complement can be computed; that is, $\mathcal{T}_=$ is gentle. $\square$

Proposition 2. *The theory $\mathcal{T}_= \oplus \mathcal{T}_f$ is undecidable.*

Proof. Let

$$\varphi^{=}_{\geq n} \coloneqq \bigwedge_{1 \leq i < j \leq n} \neg(x_i = x_j) \wedge \bigwedge_{i=1}^{n} s(x_i) = x_i.$$

Given that f is non-computable, it suffices to show that for each $n \in \mathbb{N}^*$, the sentence $P_{n+1} \wedge \varphi^{=}_{\geq f_1(n)+1}$ is $\mathcal{T}_= \oplus \mathcal{T}_f$-satisfiable if and only if $f(n+1) = 1$; if $\mathcal{T}^h_P \oplus \mathcal{T}_f$ were decidable, this would allow us to compute $f(n+1)$ recursively in terms of $f(1), \ldots, f(n)$.

First, suppose $f(n+1) = 1$. Then, $f_1(n+1) = f_1(n) + 1$. Since $\mathcal{T}_f$ is smooth [9, Lemma 54] and has an interpretation of size 1, there is a $\mathcal{T}_f$-interpretation of every size in $\mathbb{N}^*$. Further, any $\mathcal{T}_f$-interpretation of size $m \in \mathbb{N}^*$ satisfies $\varphi^{=}_{\geq f_1(m)}$. Thus, there is a $\mathcal{T}_f$-interpretation $\mathcal{A}$ of size $n+1$ satisfying $\varphi^{=}_{\geq f_1(n)+1}$. We can extend $\mathcal{A}$ to a $\mathcal{T}_= \oplus \mathcal{T}_f$-interpretation $\mathcal{B}$ satisfying $P_{n+1} \wedge \varphi^{=}_{\geq f_1(n)+1}$ by letting $P^{\mathcal{B}}_{n+1}$ be true and $P^{\mathcal{B}}_{n'}$ be false for all $n' \neq n+1$.

Second, suppose $f(n+1) = 0$. Then, $f_1(n+1) = f_1(n)$, so any $\mathcal{T}_f$-interpretation $\mathcal{A}$ of size $n+1$ satisfies $\psi^{=}_{=f_1(n)}$ and therefore does not satisfy $\varphi^{=}_{\geq f_1(n)+1}$. Hence, $P_{n+1} \wedge \varphi^{=}_{\geq f_1(n)+1}$ is $\mathcal{T}_= \oplus \mathcal{T}_f$-unsatisfiable. $\square$

B Proof of Theorem 6

It suffices to show that $\mathcal{T}_=$ is gentle, and this was done in Proposition 1.

C Proof of Theorem 7

Proposition 3. *The theory $\mathcal{T}_=$ is strongly finitely witnessable.*

Proof. It suffices to define a strong witness for Σ_P^n-formulas that are conjunctions of literals, so let φ be a conjunction of literals. If φ does not contain any literals of the form P_n, then let $wit(\varphi) := \varphi \wedge w = w$ (where w is fresh). Otherwise, let n be the largest natural number such that the literal P_n is in φ. Then, let

$$wit(\varphi) := \varphi \wedge \bigwedge_{1 \leq i < j \leq n} w_i \neq w_j,$$

where each w_i is a fresh variable.

First, we show that φ and $\exists \overrightarrow{w}.wit(\varphi)$ are $\mathcal{T}_=$-equivalent, where $\overrightarrow{w} = vars(wit(\varphi)) \setminus vars(\varphi)$. This is clear if φ does not contain any literals of the form P_n. Otherwise, let n be the largest natural number such that the literal P_n is in φ. Since $\mathcal{T}_=$ has the axiom $P_n \rightarrow \psi_{=n}$, any $\mathcal{T}_=$-interpretation that satisfies φ has n elements. In particular, any $\mathcal{T}_=$-interpretation that satisfies φ satisfies

$$\exists \overrightarrow{w}. \bigwedge_{1 \leq i < j \leq n} w_i \neq w_j.$$

It follows that φ and $\exists \overrightarrow{w}.\, wit(\varphi)$ are $\mathcal{T}_=$-equivalent.

Now, let δ be an arrangement on a finite set of variables V such that $wit(\varphi) \wedge \delta$ has a $\mathcal{T}_=$-interpretation $\mathcal{A}'$ satisfying it. We need to show that there is a $\mathcal{T}_=$-interpretation $\mathcal{A}$ satisfying $wit(\varphi) \wedge \delta$ such that $dom(\mathcal{A}) = vars(wit(\varphi) \wedge \delta)^{\mathcal{A}}$. If φ does not contain any literals of the form P_n, then we get our desired interpretation $\mathcal{A}$ by letting $dom(\mathcal{A}) = vars(wit(\varphi) \wedge \delta)^{\mathcal{A}'}$, letting $x^{\mathcal{A}} = x^{\mathcal{A}'}$ for each variable $x \in vars(wit(\varphi) \wedge \delta)$ (and letting $x^{\mathcal{A}}$ be arbitrary for $x \notin vars(wit(\varphi) \wedge \delta)$), and letting $P_n^{\mathcal{A}}$ be false for all $n \in \mathbb{N}^*$. Otherwise, let n be the (necessarily unique) natural number such that the literal P_n is in φ. Then, $\mathcal{A}'$ has exactly n elements, so $dom(\mathcal{A}') = \{w_1^{\mathcal{A}'}, \ldots, w_n^{\mathcal{A}'}\}$. Thus, we can simply take $\mathcal{A} = \mathcal{A}'$ in this case. □

D Proof of Lemma 1

We actually prove this in the many-sorted setting,[12] and then Lemma 1 follows as a particular instance, where the number of sorts is 1. As mentioned, this was already proven in [7] (also for the many-sorted case), but not explicitly stated. More precisely, there, they proved shininess from strong politeness, and so relied on smoothness. However, a careful look at the proof reveals the fact that smoothness was not relied on in the specific part of the proof that showed the computability of the minimal model function. To be safe, we provide a full proof here.

Let Σ be a many-sorted, first-order signature, S a finite set of its sorts, and $vars_\sigma(\varphi)$ be the set of variables of sort σ in φ.

[12] See [25] for the definitions of a many-sorted signature, and what an interpretation is in that case.

Finite witnessability: A Σ-theory $\mathcal{T}$ is finitely witnessable *w.r.t.* S if there exists a function *wit* (called a witness) from the quantifier-free formulas of Σ into themselves such that, for every quantifier-free formula φ, one has that: (I) φ and $\exists\,\overrightarrow{x}.\,wit(\varphi)$ are $\mathcal{T}$-equivalent, where $\overrightarrow{x} = vars(wit(\varphi)) \setminus vars(\varphi)$; (II) if $wit(\varphi)$ is $\mathcal{T}$-satisfiable there exists a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies $wit(\varphi)$ with $\sigma^{\mathcal{A}} = vars_{\sigma}(wit(\varphi))^{\mathcal{A}}$.

Strong finite witnessability: A Σ-theory $\mathcal{T}$ is strongly finitely witnessable *w.r.t.* S if it is finitely witnessable *w.r.t.* S, with witness *wit*, which in addition satisfies: (II') for every quantifier-free formula φ, finite set of variables V, and arrangement δ_V on V, if $wit(\varphi) \wedge \delta_V$ is $\mathcal{T}$-satisfiable then there exists a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies $wit(\varphi) \wedge \delta_V$ with $\sigma^{\mathcal{A}} = vars_{\sigma}(wit(\varphi) \wedge \delta_V)^{\mathcal{A}}$. In that case *wit* is called a strong witness.

Minimal model function The minimal model function **minmod** *w.r.t.* S of a Σ-theory $\mathcal{T}$ is a function from the quantifier-free formulas of Σ to the power set of $\mathbb{N}_{\omega}^{S}$ (that is, the set of functions from S to $\mathbb{N}_{\omega}$) such that: (I) if φ is a quantifier-free $\mathcal{T}$-satisfiable formula and $\mathbf{n} \in \mathbf{minmod}(\varphi)$, there exists a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies φ with $|\sigma^{\mathcal{A}}| = \mathbf{n}(\sigma)$ for every $\sigma \in S$; (II) if φ is a quantifier-free $\mathcal{T}$-satisfiable formula, $\mathbf{n} \in \mathbf{minmod}(\varphi)$, and $\mathcal{B}$ is a $\mathcal{T}$-interpretation that satisfies φ with $|\sigma^{\mathcal{B}}| \neq \mathbf{n}(\sigma)$ for some $\sigma \in S$, there exists $\sigma_* \in S$ such that $\mathbf{n}(\sigma_*) < |\sigma_*^{\mathcal{B}}|$.

Theorem 14. *If $\mathcal{T}$ is decidable and strongly finitely witnessable with respect to a finite set of sorts S, then it has a computable minimal model function with respect to S.*

Proof. Assume, without loss of generality, that $S = \{\sigma_1, \dots, \sigma_n\}$ so that we may write an element of the minimal model function as $(n_1, \dots, n_n)$. Let *wit* be the strong witness for $\mathcal{T}$, V be the set of variables in $wit(\varphi)$, V_i the set of variables in $wit(\varphi)$ of sort σ_i, $Eq_i(V)$ the set of equivalence relations on V_i (finite and easily algorithmically found), and $Eq(V)$ the product of $Eq_i(V)$. We then define $\mathbf{minmod}(\varphi)$ as the set of minimal elements of the set

$$T(\varphi) = \{(|V_1/E_1|, \dots, |V_n/E_n|) : E \in Eq(V) \text{ and } wit(\varphi) \wedge \delta_V^E \text{ is } \mathcal{T}\text{-satisfiable}\},$$

under the order such that $(n_1, \dots, n_n) \leq (m_1, \dots, m_n)$ if and only if $n_i \leq m_i$ for each $1 \leq i \leq n$, and where: $E = (E_1, \dots, E_n)$, and δ_V^E is the arrangement on V inducing the equivalence E_i on each V_i. This is computable as *wit* is computable, $\mathcal{T}$ is decidable, and the set whose minimal elements we must find is finite.

Take an element $(|V_1/E_1|, \dots, |V_n/E_n|)$ of $\mathbf{minmod}(\varphi)$, meaning $wit(\varphi) \wedge \delta_V^E$ is $\mathcal{T}$-satisfiable; as *wit* is a strong witness, there is a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies $wit(\varphi) \wedge \delta_V^E$ with $\sigma_i^{\mathcal{A}} = vars_{\sigma_i}(wit(\varphi) \wedge \delta_V^E)^{\mathcal{A}} = V_i^{\mathcal{A}}$ for every $\sigma_i \in S$. Since $\mathcal{A}$ satisfies δ_V^E, $V_i^{\mathcal{A}}$ has as many elements as V_i/E_i, and so $(|\sigma_1^{\mathcal{A}}|, \dots, |\sigma_n^{\mathcal{A}}|) = (|V_1/E_1|, \dots, |V_n/E_n|)$, meaning the first property of a minimal model function is satisfied.

Now, suppose for the sake of contradiction that there is a tuple $(|V_1/E_1|, \dots, |V_n/E_n|)$ in $\mathbf{minmod}(\varphi)$ and a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies φ such that

$|\sigma_i^{\mathcal{A}}| \leq |V_i/E_i|$, for all $\sigma_i \in S$, and for at least one of them $|\sigma_i^{\mathcal{A}}| < |V_i/E_i|$. As *wit* is a strong witness we have that $\mathcal{A}$ satisfies $\exists \overrightarrow{x}.\, wit(\varphi)$, for $\overrightarrow{x} = vars(wit(\varphi)) \setminus vars(\varphi)$, and therefore there is a $\mathcal{T}$-interpretation $\mathcal{B}$, differing from $\mathcal{A}$ at most on the values assigned to $\overrightarrow{x}$, that satisfies $wit(\varphi)$. Let F_i be the equivalence relation induced by $\mathcal{B}$ over V_i, and $F = (F_1, \dots, F_n)$, so that $\mathcal{B}$ satisfies $wit(\varphi) \wedge \delta_V^F$. Again, by using the fact that *wit* is a strong witness there must exist a third $\mathcal{T}$-interpretation $\mathcal{C}$ that satisfies $wit(\varphi) \wedge \delta_V^F$ with $\sigma_i^{\mathcal{C}} = vars_{\sigma_i}(wit(\varphi) \wedge \delta_V^F)^{\mathcal{C}} = V_i^{\mathcal{C}}$ for each $1 \leq i \leq n$. Using $\mathcal{C}$ satisfies δ_V^F, $|\sigma_i^{\mathcal{C}}| = |V_i/F_i|$, and since $\mathcal{B}$ also satisfies δ_V^F, $|V_i/F_i| \leq |\sigma_i^{\mathcal{B}}| = |\sigma_i^{\mathcal{A}}|$. This means that $(|\sigma_1^{\mathcal{C}}|, \dots, |\sigma_n^{\mathcal{C}}|)$, although being in $T(\varphi)$, is strictly less than $(|V_1/E_1|, \dots, |V_n/E_n|)$, a minimal element of $T(\varphi)$, leading to a contradiction and finishing the proof. □

E Proof of Theorem 11

Proposition 4. *The theory $\mathcal{T}_{\leq}$ is decidable.*

Proof. It suffices to show that it is decidable whether a conjunction of literals is $\mathcal{T}_{\leq}$-satisfiable, so let φ be a conjunction of literals. Write $\varphi = \varphi_1 \wedge \varphi_2$, where φ_1 contains the equalities and disequalities in φ and φ_2 contains the literals of the form P_n and $\neg P_n$ in φ.

We describe our decision procedure as follows. If φ_1 is unsatisfiable in equational logic, then φ is $\mathcal{T}_{\leq}$-unsatisfiable. Otherwise, let m be the size of the smallest interpretation that satisfies φ_1.

We claim that, in this case, φ is $\mathcal{T}_{\leq}$-satisfiable if and only if for every n such that the literal P_n is in φ_2, we have $F(n) \geq m$. This is because if the latter condition holds, we can extend an interpretation that satisfies φ_1 to a $\mathcal{T}_{\leq}$-interpretation $\mathcal{A}$ satisfying φ by setting $P_n^{\mathcal{A}}$ to true for every n such that the literal P_n is in φ_2 and setting $P_n^{\mathcal{A}}$ to false otherwise. Otherwise, there is some n such that the literal P_n is in φ_2 and $F(n) < m$. In this case, there is no interpretation that satisfies φ_1 of size at most $F(n)$, so φ is $\mathcal{T}_{\leq}$-unsatisfiable (since $P_n \rightarrow \psi_{\leq F(n)}$ is an axiom of $\mathcal{T}_{\leq}$). □

Proposition 5. *The theory $\mathcal{T}_{\infty}$ is decidable.*

Proof. We prove $\mathcal{T}_{\infty}$ and equational logic satisfy the same quantifier-free formulas, and since the latter is decidable so will be $\mathcal{T}_{\infty}$. Of course equational logic satisfies all quantifier-free formulas that $\mathcal{T}_{\infty}$ satisfies, given that it has more models than $\mathcal{T}_{\infty}$. Reciprocally, suppose the quantifier-free formula φ is satisfied by equational logic, and let $\mathcal{A}$ be an interpretation that satisfies φ. We consider the interpretation $\mathcal{B}$ with $dom(\mathcal{B}) = dom(\mathcal{A}) \cup B$, for a set $B = \{b_n : n \in \mathbb{N}\}$ disjoint from $dom(\mathcal{A})$, and $x^{\mathcal{B}} = x^{\mathcal{A}}$ for all variables x. $\mathcal{B}$ is a $\mathcal{T}_{\infty}$-interpretation, of course, but it also satisfies φ, which can be proven by a simple induction on the subformulas of φ. □

Proposition 6. *The theory $\mathcal{T}_{\infty} \oplus \mathcal{T}_{\leq}$ is undecidable.*

Proof. It suffices to show that P_n is $\mathcal{T}_\infty \oplus \mathcal{T}_\leq$-satisfiable if and only if $F(n) = \aleph_0$. If $F(n) = \aleph_0$, then P_n is satisfied by the $\mathcal{T}_\infty \oplus \mathcal{T}_\leq$-interpretation $\mathcal{A}$ of size $\aleph_0$ where $P_n^{\mathcal{A}}$ is true and $P_{n'}^{\mathcal{A}}$ is false for all $n' \neq n$. If $F(n) < \aleph_0$, then any $\mathcal{T}_\leq$-interpretation satisfying P_n must be finite (since $\mathcal{T}_\leq$ has the axiom $P_n \rightarrow \psi_{\leq F(n)}$). Hence, P_n is $\mathcal{T}_\infty \oplus \mathcal{T}_\leq$-unsatisfiable. □

F Proofs concerning $\mathcal{T}^2_{orb}$

Proposition 7. *The theory $\mathcal{T}^2_{orb}$ is gentle.*

Proof. Let: $vars(\varphi)$ equal $\{x_1, \dots, x_n\}$; M_i be the maximum of j such that $t^j(x_i)$ shows up in φ; M_0' be the maximum of j such that $t^j(a)$ appears in φ, and if it doesn't we set M_0' to 0; $M_0 = M_0' + \sum_{i=1}^{n}(M_i + 1)$; and take fresh variables $x_{i,j}$, for $0 \leq i \leq n$ and $0 \leq j \leq M_i$. We then flatten and Ackermannize φ, meaning: we replace any term $t^j(x_i)$ by $x_{i,j}$, and any term $t^j(a)$ by $x_{0,j}$, in order to obtain the formula of equational logic φ'; and define the formula φ_* as $\varphi' \wedge Fun(V)$, where $V = \{x_{i,j} : 0 \leq i \leq n, 0 \leq j \leq M_i\}$ and

$$Fun(V) = \bigwedge_{0 \leq i,p \leq n} \bigwedge_{0 \leq j < M_i} \bigwedge_{0 \leq q < M_p} (x_{i,j} = x_{p,q}) \rightarrow (x_{i,j+1} = x_{p,q+1}).$$

Now, consider the set $Eq(V)$ of equivalence relations on V, which is easily computable; $[x_{i,j}]_E$ shall denote in what follows the equivalence class under E with representative $x_{i,j}$. We define a subset $Eq'(V)$ of $Eq(V)$ such that E is in $Eq'(V)$ if, and only if, when defining the interpretation of equational logic with domain V/E and where $x_{i,j}$ is assigned the value $[x_{i,j}]_E$ (an interpretation we shall denote by $\mathcal{V}/\mathcal{E}$), φ_* is true in this interpretation (this can be decided algorithmically given the finiteness of V/E).

For $E \in Eq'(V)$ we define a partial function t_E on V/E by making $t_E([x_{i,j}]_E) = [x_{i,j+1}]_E$ for all $i \in \{0, \dots, n\}$ and $j \in \{0, \dots, M_i - 1\}$ (notice that $t_E([x_{i,M_i}]_E)$ may still be defined if $[x_{i,M_i}]_E = [x_{p,q}]_E$ for a $q \in \{0, \dots, M_p - 1\}$). This is well-defined: if $[x_{i,j}]_E = [x_{p,q}]_E$ for $0 \leq i, p \leq n$, $0 \leq j < M_i$ and $0 \leq q < M_p$, we have that $\mathcal{V}/\mathcal{E}$ satisfies $x_{i,j} = x_{p,q}$; since it also satisfies $Fun(V)$, we have that it satisfies $x_{i,j+1} = x_{p,q+1}$, meaning that $t_E([x_{i,j}]_E) = [x_{i,j+1}]_E = [x_{p,q+1}]_E = t_E([x_{p,q}]_E)$. The partial function t_E can be computed by an exhaustive search, as V, and thus V/E, is finite.

We then let B_0^E be the orbit of $[x_{0,0}]_E$ under t_E: for a partial function, this means either the list $\{t_E^j([x_{0,0}]_E) : j \in \mathbb{N}\}$ if t_E is always defined on $t_E^j([x_{0,0}]_E)$; or the list $\{t_E^0([x_{0,0}]_E), \dots, t_E^J([x_{0,0}]_E)\}$, if t_E is defined on all $t_E^j([x_{0,0}]_E)$ for $0 \leq j \leq J-1$, but not on $t_E^J([x_{0,0}]_E)$; this can be easily found algorithmically.[13] Define $Eq''(V)$ as the subset of $Eq'(V)$ where $2|B_0^E| \geq |V/E|$. For every $E \in Eq''(V)$, define the interval $I(E)$ as

$$I(E) = \begin{cases} [|V/E|, 2|B_0^E|] & \text{if } t_E \text{ is defined for all of } B_0^E, \\ \{n \in \mathbb{N}^* : n \geq |V/E|\} \cup \{\aleph_0\} & \text{otherwise,} \end{cases}$$

[13] The proof of Proposition 13 continues from here.

and we state that

$$Spec(\varphi) = \bigcup_{E \in Eq''(V)} I(E)$$

if $Eq''(V)$ is not empty, and $Spec(\varphi) = \emptyset$ otherwise. Given that the sets $I(E)$ are computable and either finite or cofinite, so is $Spec(\varphi)$ if the identity truly holds, meaning $\mathcal{T}^2_{orb}$ is gentle. We prove the identity in three cases.

1. If $Eq''(V)$ is not empty and t_E is defined for all of B^E_0, for each $0 \leq j \leq 2|B^E_0| - |V/E|$ take a set B with cardinality j disjoint from V/E, and we define a $\mathcal{T}^2_{orb}$-interpretation $\mathcal{A}_j$ as follows.
We make $dom(\mathcal{A}_j) = (V/E) \cup B$, which then has $|V/E| + j \leq 2|B^E_0|$ elements. Of course $a^{\mathcal{A}_j} = [x_{0,0}]_E$. $t^{\mathcal{A}_j}(b) = t_E(b)$ for all b where t_E is defined, and $t^{\mathcal{A}_j}(b) = b$ otherwise: this way the orbit of $a^{\mathcal{A}_j}$ under $t^{\mathcal{A}}$ has $|B^E_0|$ elements, making $\mathcal{A}_j$ a $\mathcal{T}^2_{orb}$-interpretation. And, finally, $x^{\mathcal{A}_j} = [x]_E$ for all variables $x \in V$, $x_i^{\mathcal{A}_j} = [x_{i,0}]_E$ for all variables x_i in φ, and arbitrarily otherwise, so $\mathcal{A}_j$ satisfies φ as $E \in Eq'(V)$.
There cannot exist a $\mathcal{T}^2_{orb}$-interpretation $\mathcal{A}$ that induces the equivalence E on V with fewer than $|V/E|$ elements, obviously. And there cannot exist a $\mathcal{T}^2_{orb}$-interpretation $\mathcal{A}$ that induces the equivalence E on V with more than $2|B^E_0|$ elements as the orbit of $a^{\mathcal{A}}$ under $t^{\mathcal{A}}$ has necessarily $|B^E_0|$ elements.
2. If $Eq''(V)$ is not empty and t_E is not defined over all of B^E_0, take any $j \in \mathbb{N}$ and a set $B = \{b_1, \dots, b_j\}$ with cardinality j disjoint from V/E, and we define a $\mathcal{T}^2_{orb}$-interpretation $\mathcal{A}_j$ as follows.
First $dom(\mathcal{A}_j) = (V/E) \cup B$, so $|dom(\mathcal{A}_j)| = |V/E| + j$. Second, of course $a^{\mathcal{A}_j} = [x_{0,0}]_E$. Third: $t^{\mathcal{A}_j}(b) = t_E(b)$ for all b where t_E is defined; $t^{\mathcal{A}_j}(b) = b_1$ for the one element $b \in B^E_0$ where t_E is not defined; $t^{\mathcal{A}_j}(b_i) = b_{i+1}$ for $1 \leq i \leq j-1$; and, for all elements b where $t^{\mathcal{A}_j}$ hasn't been defined yet, including b_j, $t^{\mathcal{A}}(b) = b$. Notice that, this way, the orbit of $a^{\mathcal{A}_j}$ under $t^{\mathcal{A}_j}$ has size $|B^E_0| + j$, and since $2(|B^E_0| + j) \geq |V/E| + 2j \geq |V/E| + j = |dom(\mathcal{A}_j)|$ we get $\mathcal{A}_j$ is a $\mathcal{T}^2_{orb}$-interpretation. And, finally, $x^{\mathcal{A}_j} = [x]_E$ for all variables $x \in V$, $x_i^{\mathcal{A}_j} = [x_{i,0}]_E$ for all variables x_i in φ, and arbitrarily otherwise, so $\mathcal{A}_j$ satisfies φ as $E \in Eq'(V)$.
Of course an interpretation that satisfies φ and induces the equivalence E must have at least $|V/E|$ elements, so we are done.
3. Suppose a $\mathcal{T}^2_{orb}$-interpretation $\mathcal{A}$ satisfies φ: we can change the values assigned to the variables $x_{i,j}$ while keeping φ satisfied, as they are not in φ, so that $x^{\mathcal{A}}_{i,j} = (t^{\mathcal{A}})^j(x^{\mathcal{A}}_i)$; take then the equivalence E on V such that $x_{i,j} E x_{p,q}$ if $x^{\mathcal{A}}_{i,j} = x^{\mathcal{A}}_{p,q}$. Of course φ_* is satisfied by $\mathcal{V}/\mathcal{E}$, so $E \in Eq'(V)$. If t_E is defined for all elements of the orbit of $[x_{0,0}]_E$ under t_E, we have that $|dom(\mathcal{A})| = 2|B^E_0|$ and, since $|dom(\mathcal{A})| \geq |V/E|$, we get $2|B^E_0| \geq |V/E|$; if it is not, then it contains M_0 elements, and since $|V/E| \leq M_0$ we again get $2|B^E_0| \geq |V/E|$, proving that $Eq''(V)$. Therefore, for φ to be satisfiable we must have some $E \in Eq''(V)$.

□

Proposition 8. *The theory $\mathcal{T}_{orb}^2$ is decidable.*

Proof. Follows from Proposition 7: a quantifier-free formula φ is $\mathcal{T}_{orb}^2$-satisfiable if and only if $\mathit{Spec}(\varphi)$ is not empty, something that is decidable. $\square$

Proposition 9. *The theory $\mathcal{T}_{orb}^2$ is strongly finitely witnessable.*

Proof. Let x_1 through x_n be the variables in a quantifier-free formula φ, M_i be the maximum of j such that $t^j(x_i)$ occurs in φ, M_0' be the maximum of j such that $t^j(x_0)$ occurs in φ, M_0 be $M_0'+\sum_{i=1}^{n}(M_i+1)$, and take fresh variables $x_{i,j}$, for $0 \leq i \leq n$ and $0 \leq j \leq M_i$. We state

$$\mathit{wit}(\varphi) = \varphi \wedge \bigwedge_{i=0}^{n} \bigwedge_{j=0}^{M_i} x_{i,j} = t^j(x_i)$$

is a strong witness for $\mathcal{T}_{orb}^2$. Of course it maps quantifier-free formulas into other quantifier-free formulas, and is computable. Furthermore, for $\overrightarrow{x} = \mathit{vars}(\mathit{wit}(\varphi)) \setminus \mathit{vars}(\varphi)$, it is obvious that $\exists\, \overrightarrow{x}.\mathit{wit}(\varphi)$ implies φ, since $\mathit{wit}(\varphi)$ itself already implies φ. Reciprocally, if the $\mathcal{T}_{orb}^2$-interpretation $\mathcal{A}$ satisfies φ, we produce a new $\mathcal{T}_{orb}^2$-interpretation $\mathcal{B}$ by changing the values assigned by $\mathcal{A}$ to those variables in $\overrightarrow{x}$ so that $x_{i,j}^{\mathcal{B}} = (t^{\mathcal{A}})^j(x_i^{\mathcal{A}})$; this way $\mathcal{B}$ satisfies $\mathit{wit}(\varphi)$, and therefore $\mathcal{A}$ satisfies $\exists\, \overrightarrow{x}.\, \mathit{wit}(\varphi)$.

Now, take a finite set of variables V (not to be confused with the V used in the proof above that $\mathcal{T}_{orb}^2$ is decidable), an arrangement δ_V over V, and a $\mathcal{T}_{orb}^2$-interpretation $\mathcal{A}$ that satisfies $\mathit{wit}(\varphi) \wedge \delta_V$: there are two cases we consider; for simplicity, let U denote $\mathit{vars}(\mathit{wit}(\varphi))$.

1. Suppose that the orbit of $a^{\mathcal{A}}$ under $t^{\mathcal{A}}$ is a subset of $U^{\mathcal{A}} \cup V^{\mathcal{A}}$, and we then define an interpretation $\mathcal{B}$ by making: $\mathit{dom}(\mathcal{B}) = U^{\mathcal{A}} \cup V^{\mathcal{A}}$; $a^{\mathcal{B}} = a^{\mathcal{A}}$; $t^{\mathcal{B}}(b) = t^{\mathcal{A}}(b)$ whenever the latter value is in $\mathit{dom}(\mathcal{B})$, and otherwise $t^{\mathcal{B}}(b) = b$ (this way, the orbit of $a^{\mathcal{B}}$ under $t^{\mathcal{B}}$ is the same as the orbit of $a^{\mathcal{A}}$ under $t^{\mathcal{A}}$, and since $|\mathit{dom}(\mathcal{B})| \leq |\mathit{dom}(\mathcal{A})|$ we get $\mathcal{B}$ is a $\mathcal{T}_{orb}^2$-interpretation); and $x^{\mathcal{B}} = x^{\mathcal{A}}$ for every variable $x \in U \cup V$, and arbitrarily otherwise (so $\mathit{dom}(\mathcal{B}) = \mathit{vars}(\mathit{wit}(\varphi) \wedge \delta_V)^{\mathcal{B}})$).
   It is clear that $\mathcal{B}$ satisfies δ_V, given the definition of how it assigns values to variables; since $x_{i,j}^{\mathcal{B}} = x_{i,j}^{\mathcal{A}}$ and $x_{i,j}^{\mathcal{A}} = (t^{\mathcal{A}})^j(x_i^{\mathcal{A}})$, and thus $(t^{\mathcal{B}})^j(x_i^{\mathcal{B}}) = (t^{\mathcal{A}})^j(x_i^{\mathcal{A}})$ (for $0 \leq j \leq M_i$), we get $\mathcal{B}$ satisfies $\bigwedge_{i=0}^{n} \bigwedge_{j=0}^{M_i} x_{i,j} = t^j(x_i)$.
   Finally, for any atomic subformula of φ of the form $t^j(x_i) = t^q(x_p)$ (the cases $t^j(x_i) = t^q(a)$ and $t^j(a) = t^q(a)$ being analogous), since $(t^{\mathcal{B}})^j(x_i^{\mathcal{B}}) = (t^{\mathcal{A}})^j(x_i^{\mathcal{A}})$ and $(t^{\mathcal{B}})^q(x_p^{\mathcal{B}}) = (t^{\mathcal{A}})^q(x_p^{\mathcal{A}})$, we get that the subformula is true in $\mathcal{B}$ if and only if it is true in $\mathcal{A}$; since φ is quantifier-free, $\mathcal{B}$ satisfies φ, and thus $\mathit{wit}(\varphi)$, and we are done.
2. Suppose then that the orbit of $a^{\mathcal{A}}$ under $t^{\mathcal{A}}$ is not entirely contained in $U^{\mathcal{A}} \cup V^{\mathcal{A}}$: there are, however, at least $M_0 + 1$ (distinct) elements of it in $U^{\mathcal{A}} \cup V^{\mathcal{A}}$. For simplicity, assume $\{b_1, \dots, b_R\}$ is an enumeration of $V^{\mathcal{A}} \setminus U^{\mathcal{A}}$. We then define an interpretation $\mathcal{B}$ as follows: $\mathit{dom}(\mathcal{B}) = U^{\mathcal{A}} \cup V^{\mathcal{A}}$; $a^{\mathcal{B}} = a^{\mathcal{A}}$;

$t^{\mathcal{B}}(c) = t^{\mathcal{A}}(c)$ whenever the latter is in $U^{\mathcal{A}}$, $(t^{\mathcal{B}})^{M_0+1}(a^{\mathcal{B}}) = b_1$, $t^{\mathcal{B}}(b_r) = b_{r+1}$ for $1 \leq r < R$, and if $t^{\mathcal{B}}(c)$ hasn't been defined yet we make it equal to c (the orbit of $a^{\mathcal{B}}$ under $t^{\mathcal{B}}$ has then $M_0 + R + 1$ elements, and since $|U^{\mathcal{A}}| \leq 2M_0 + 1$ and $|V^{\mathcal{A}} \setminus U^{\mathcal{A}}| = R$ we conclude $\mathcal{B}$ is a $\mathcal{T}^2_{orb}$-interpretation); and $x^{\mathcal{B}} = x^{\mathcal{A}}$ for every variable $x \in U \cup V$, and arbitrarily otherwise (so $dom(\mathcal{B}) = vars(wit(\varphi) \wedge \delta_V)^{\mathcal{B}}))$.
We finally prove $\mathcal{B}$ satisfies $wit(\varphi) \wedge \delta_V$ just as in the item above.

□

Proposition 10. *The theory $\mathcal{T}^2_{orb}$ has the finite model property.*

Proof. Follows from Proposition 9 and the fact that strong finite witnessability implies the finite model property, as shown in [26]. □

Proposition 11. *The theory $\mathcal{T}^2_{orb}$ has a computable minimal model function.*

Proof. Follows from Lemma 1 and Propositions 8 and 9. □

Proposition 12. *The theory $\mathcal{T}_g \oplus \mathcal{T}^2_{orb}$ is undecidable.*

Proof. Assume instead $\mathcal{T}_g \oplus \mathcal{T}^2_{orb}$ is decidable, and let us define a function $G : \mathbb{N}^* \to \{0, 1\}$ and formulas φ_n by making: $G(1) = G(3) = 1$ and $G(2) = G(4) = 0$; assuming G defined up to $2n$, φ_n equal to $\psi^{orb}_{=n+1}(a) \wedge \varphi^{=}_{\geq G_1(2n)+2}$ (where, as before, $G_1(n) = |\{1 \leq i \leq n : G(i) = 1\}|$), and

$$G(2n+1) = G(2n+2) = \begin{cases} 1 & \text{if } \varphi_n \text{ is } \mathcal{T}_g \oplus \mathcal{T}^2_{orb}\text{-satisfiable,} \\ 0 & \text{otherwise;} \end{cases}$$

Of course G is computable, but we shall show that $G = g$ and reach a contradiction; assume this is true for all values up to $2n$ for $n \geq 2$, meaning in particular that $G_1(n) = g_1(n)$, and we shall show $g(2n+1) = g(2n+2) = 1$ if and only if φ_n is $\mathcal{T}_g \oplus \mathcal{T}^2_{orb}$-satisfiable.

If it is satisfiable, there is a $\mathcal{T}_g \oplus \mathcal{T}^2_{orb}$-interpretation $\mathcal{A}$ that satisfies φ_n, and thus $\psi^{orb}_{=n+1}(a)$ as well as $\varphi^{=}_{\geq g_1(2n)+2}$; from the axiom $\psi^{orb}_{=n+1}(a) \to \psi_{\leq 2n+2}$ of $\mathcal{T}^2_{orb}$ we get $\mathcal{A}$ has at most $2n+2$ elements, and from the axiomatization of $\mathcal{T}_g$ we get that it has at most $g_1(2n+2)$ elements satisfying $s(x) = x$. The fact that $\mathcal{A}$ satisfies $\varphi^{=}_{\geq g_1(2n)+2}$ implies it has at least $g_1(2n)+2$ elements satisfying $s(x) = x$, and these two last facts are only compatible if $g(2n+1) = g(2n+2) = 1$.

Reciprocally, suppose $g(2n+1) = g(2n+2) = 1$, and so there exists a $\mathcal{T}_g$-interpretation $\mathcal{A}$ with exactly $g_1(2n+2) = g_1(2n) + 2$ elements satisfying $s(x) = x$, and $g_0(2n+2) = g_0(2n)$ satisfying $s(x) \neq x$ (and thus $2n+2$ elements in total, which we name a_1 through a_{2n+2}). Extend $\mathcal{A}$ to a $\Sigma^a_t \oplus \Sigma_s$-interpretation $\mathcal{B}$ by making $a^{\mathcal{B}} = a_1$, $t^{\mathcal{B}}(a_i) = a_{i+1}$ for $1 \leq i \leq n$, $t^{\mathcal{B}}(a_{n+1}) = a_{n+1}$, and $t^{\mathcal{B}}(a_i) = a_i$ for $n+2 \leq i \leq 2n+2$: $\psi^{orb}_{=n+1}(a)$ is then true in $\mathcal{B}$, while all $\psi^{orb}_{=m}(a)$, for $m \neq n+1$, are obviously false. We prove that $\mathcal{B}$ is then a $\mathcal{T}_g \oplus \mathcal{T}^2_{orb}$-interpretation that satisfies φ_n, which shall finish the proof. $\mathcal{B}$ is certainly a $\mathcal{T}_g$-interpretation, vacuously satisfies all axioms $\psi^{orb}_{=m}(a) \to \psi_{\leq 2m}$ for $m \neq n+1$, and satisfies

$\psi^{orb}_{=n+1}(a) \rightarrow \psi_{\leq 2n+2}$ given that it satisfies $\psi^{orb}_{=n+1}(a)$ and has $2n+2$ elements, making of it a $\mathcal{T}_g \oplus \mathcal{T}^2_{orb}$-interpretation. Furthermore, as mentioned before it satisfies $\psi^{orb}_{=n+1}(a)$, and satisfies $\varphi^{=}_{\geq g_1(2n)+2}$ since it has $g_1(2n+2) = g_1(2n)+2$ elements that satisfy $s(x) = x$. □

G Proofs concerning $\mathcal{T}^{orb}_{\leq}$

Proposition 13. *The theory $\mathcal{T}^{orb}_{\leq}$ is decidable.*

Proof. This is essentially the same as the proof of Proposition 8, the proofs being exactly the same up to Footnote 13. The difference is that φ is $\mathcal{T}^{orb}_{\leq}$-satisfiable if and only if there is an $E \in Eq'(V)$ such that

$$|B_0| + F(|B_0|) \geq |V/E|.$$

Now, we may not be able to calculate $F(|B_0|)$, but we can equivalently write this condition as $F(|B_0|) \geq |V/E| - |B_0|$, and such tests being computable are a prerequisite for F. □

Proposition 14. *The theory $\mathcal{T}_\infty \oplus \mathcal{T}^{orb}_{\leq}$ is undecidable.*

Proof. Consider the formulas $\psi^{orb}_{=n}$: we state that they are $\mathcal{T}_\infty \oplus \mathcal{T}^{orb}_{\leq}$-satisfiable if and only if $F(n) = \aleph_0$, what we know cannot be tested algorithmically.

Take first an n such that $F(n) = \aleph_0$, and define an interpretation $\mathcal{A}$ as follows: $dom(\mathcal{A}) = \mathbb{N}$ (so $\mathcal{A}$ is a $\mathcal{T}_\infty$-interpretation) and $a^{\mathcal{A}} = 0$; $t^{\mathcal{A}}(i) = i+1$ for all $i \neq n-1$, and $t^{\mathcal{A}}(n-1) = 0$, so the orbit of 0 is $\{0, \dots, n-1\}$, meaning $\mathcal{A}$ satisfies $\psi^{orb}_{=n}$ and, since $F(n) = \aleph_0$ and $\mathcal{A}$ is infinite we have that it is a $\mathcal{T}^{orb}_{\leq}$-interpretation; and $x^{\mathcal{A}}$ can be defined arbitrarily for all variables x. This means $\mathcal{A}$ is a $\mathcal{T}_\infty \oplus \mathcal{T}^{orb}_{\leq}$-interpretation that satisfies φ.

Reciprocally, suppose $\mathcal{A}$ is a $\mathcal{T}_\infty \oplus \mathcal{T}^{orb}_{\leq}$-interpretation that satisfies $\psi^{orb}_{=n}$: if $F(n) \in \mathbb{N}^*$ we have $|dom(\mathcal{A})| \leq F(n) + n$, which is finite and thus contradicts the fact that $\mathcal{A}$ is supposed to be a $\mathcal{T}_\infty$-interpretation. Thus $F(n) = \aleph_0$. □

H Proof of Theorem 12

The key to the proof is the following result due to Fontaine.

Lemma 2 ([10, Corollary 1]). *Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be theories over disjoint signatures Σ_1 and Σ_2, respectively. Suppose that it is decidable whether $Spec(\mathcal{T}_1, \varphi_1) \cap Spec(\mathcal{T}_2, \varphi_2) = \emptyset$, where φ_1 and φ_2 are conjunctions of literals over the signatures Σ_1 and Σ_2, respectively. Then, $\mathcal{T}_1 \oplus \mathcal{T}_2$ is decidable.*

In light of the lemma, the following implies Theorem 12.

Lemma 3. *Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be decidable theories over disjoint signatures Σ_1 and Σ_2 respectively. Suppose that $\mathcal{T}_1$ is gentle and $\mathcal{T}_2$ has computable finite spectra. Then, it is decidable whether $Spec(\mathcal{T}_1, \varphi_1) \cap Spec(\mathcal{T}_2, \varphi_2) = \emptyset$, where φ_1 and φ_2 are conjunctions of literals over the signatures Σ_1 and Σ_2 respectively.*

Proof. Let φ_1 and φ_2 be conjunctions of literals over the signatures Σ_1 and Σ_2 respectively. We describe our decision procedure as follows. Since $\mathcal{T}_1$ is gentle, $Spec(\mathcal{T}_1, \varphi_1)$ is either of the form S or $S \cup \{n \in \mathbb{N}_\omega \mid n \geq k\}$ for some $k \in \mathbb{N}^*$, where $S \subset \mathbb{N}^*$ is a finite set. We have $S \cap Spec(\mathcal{T}_2, \varphi_2) = \emptyset$ if and only if $n \notin Spec(\mathcal{T}_2, \varphi_2)$ for each $n \in S$, which we can check algorithmically since $\mathcal{T}_2$ has computable finite spectra. We also have $\{n \in \mathbb{N}_\omega \mid n \geq k\} \cap Spec(\mathcal{T}_2, \varphi_2) = \emptyset$ if and only if $\varphi_2 \wedge \neq (x_1, \ldots, x_k)$ is $\mathcal{T}_2$-unsatisfiable (where the variables x_i are fresh). These computations allow us to determine whether $Spec(\mathcal{T}_1, \varphi_1) \cap Spec(\mathcal{T}_2, \varphi_2) = \emptyset$. □

We also prove here that Theorem 12 is a strengthening of Theorem 4.

Proposition 15. *If a theory $\mathcal{T}$ is gentle, then $\mathcal{T}$ has computable finite spectra.*

Proof. Let φ be a quantifier-free formula, and let $k \in \mathbb{N}^*$. If $\mathcal{T}$ is gentle, then we can compute an explicit representation of the set $Spec(\mathcal{T}, \varphi)$, from which we can decide whether $k \in Spec(\mathcal{T}, \varphi)$. □

Proposition 16. *If a theory $\mathcal{T}$ is finitely axiomatizable, then $\mathcal{T}$ has computable finite spectra.*

Proof. Let φ be a quantifier-free formula, and let $k \in \mathbb{N}^*$. Let Σ be the signature over which $\mathcal{T}$ is defined. We may assume that Σ only contains the symbols appearing in $Ax(\mathcal{T}) \cup \{\varphi\}$ so that, in particular, Σ is finite. We can enumerate the Σ-interpretations of size k, checking whether any of them satisfy all of the formulas in $Ax(\mathcal{T}) \cup \{\varphi\}$. If we find a $\mathcal{T}$-interpretation of size k satisfying φ, then $k \in Spec(\mathcal{T}, \varphi)$; otherwise, $k \notin Spec(\mathcal{T}, \varphi)$. □

Proposition 17. *Suppose that for a theory $\mathcal{T}$, there is an algorithm that, given a conjunction φ of literals, outputs a finite set $S \subset \mathbb{N}_\omega$ such that $Spec(\mathcal{T}, \varphi) = S$. Then, $\mathcal{T}$ has computable finite spectra.*

Proof. Let φ be a quantifier-free formula, and let $k \in \mathbb{N}^*$. Without loss of generality, φ is a conjunction of literals (the general case follows by putting φ in disjunctive normal form). We can compute a finite set $S \subset \mathbb{N}_\omega$ such that $Spec(\mathcal{T}, \varphi) = S$, from which we can decide whether $k \in Spec(\mathcal{T}, \varphi)$. □

I Proofs concerning Example 1

Proposition 18. *For any $n \in \mathbb{N}^*$, the theory $\mathcal{T}_{\leq n}$ is gentle (and therefore decidable).*

Proof. Let φ be a conjunction of literals. If φ is unsatisfiable in equational logic, then $Spec(\mathcal{T}_{\leq n}, \varphi) = \emptyset$. Otherwise, let m be the size of the size of the smallest interpretation that satisfies φ. Then, $Spec(\mathcal{T}_{\leq n}, \varphi) = [m, n]$. □

That $\mathcal{T}_{\leq}$ is decidable is proved in Proposition 4.

Proposition 19. *The theory $\mathcal{T}_{\leq}$ has computable finite spectra.*

Proof. Let φ be a conjunction of literals, and let $k \in \mathbb{N}^*$. Write $\varphi = \varphi_1 \wedge \varphi_2$, where φ_1 contains the equalities and disequalities in φ and φ_2 contains the literals of the form P_n and $\neg P_n$ in φ.

We describe our decision procedure as follows. If φ_1 is unsatisfiable in equational logic, then φ is $\mathcal{T}_{\leq}$-unsatisfiable, so $k \notin Spec(\mathcal{T}_{\leq}, \varphi)$. Otherwise, let m be the size of the smallest interpretation that satisfies φ_1. If $k < m$, then $k \notin Spec(\mathcal{T}_{\leq}, \varphi)$. So assume that $k \geq m$.

We claim that, in this case, $k \in Spec(\mathcal{T}_{\leq}, \varphi)$ if and only if for every n such that the literal P_n is in φ_2, we have $F(n) \geq k$. This is because if the latter condition holds, we can extend an interpretation of size k that satisfies φ_1 to a $\mathcal{T}_{\leq}$-interpretation $\mathcal{A}$ of size k satisfying φ by setting $P_n^{\mathcal{A}}$ to true for every n such that the literal P_n is in φ_2 and setting $P_n^{\mathcal{A}}$ to false otherwise. Otherwise, there is some n such that the literal P_n is in φ_2 and $F(n) < k$. Then, there is no interpretation of φ_2 of size k (since $P_n \rightarrow \psi_{\leq F(n)}$ is an axiom of $\mathcal{T}_{\leq}$), so $k \notin Spec(\mathcal{T}_{\leq}, \varphi)$. □

Proposition 20. *The theory $\mathcal{T}_{\leq}$ is not gentle.*

Proof. It suffices to show that it is undecidable whether $Spec(\mathcal{T}_{\leq}, \varphi)$ is infinite. And indeed, $Spec(\mathcal{T}_{\leq}, P_n)$ is infinite if and only if $F(n) = \aleph_0$, which is undecidable by our assumptions on F. □

Proposition 21. *The theory $\mathcal{T}_{\leq}$ is not finitely axiomatizable.*

Proof. If $\mathcal{T}$ is a finitely axiomatizable Σ_P^n-theory, then the predicate P_n does not appear in the axiomatization of $\mathcal{T}$ for some n such that $F(n) < \aleph_0$. Then, $P_n \rightarrow \psi_{\leq F(n)}$ is $\mathcal{T}_{\leq}$-valid but not $\mathcal{T}$-valid. □

Proposition 22. *There is no algorithm that, given a conjunction of literals φ in the language of $\mathcal{T}_{\leq}$, outputs a finite set $S \subset \mathbb{N}_\omega$ such that $Spec(\mathcal{T}_{\leq}, \varphi) = S$*

Proof. Indeed, $Spec(\mathcal{T}_{\leq}, \top)$ is not even a finite set. □

Proposition 23. *For any $n \in \mathbb{N}^*$, the theory $\mathcal{T}_{\leq n}$ is not stably infinite (and therefore neither strongly polite nor shiny).*

Proof. The theory has no models of size greater than n and hence no infinite models. □

Proposition 24. *The theory $\mathcal{T}_{\leq}$ is not stably infinite (and therefore neither strongly polite nor shiny).*

Proof. Let $n \in \mathbb{N}^*$ be such that $F(n) < \aleph_0$. Then, P_n has no infinite $\mathcal{T}_{\leq}$-models, since $P_n \rightarrow \psi_{\leq F(n)}$ is an axiom of $\mathcal{T}_{\leq}$. □

J Proof of Theorem 13

In light of Lemma 2, the following implies Theorem 13.

Lemma 4. *Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be decidable theories over disjoint signatures Σ_1 and Σ_2 respectively. Suppose that $\mathcal{T}_1$ is smooth and has a computable minimal model function and that $\mathcal{T}_2$ is infinitely decidable. Then, it is decidable whether $Spec(\mathcal{T}_1, \varphi_1) \cap Spec(\mathcal{T}_2, \varphi_2) = \emptyset$, where φ_1 and φ_2 are conjunctions of literals over the signatures Σ_1 and Σ_2 respectively.*

Proof. Let φ_1 and φ_2 be conjunctions of literals over the signatures Σ_1 and Σ_2 respectively. We describe our decision procedure as follows. If φ_1 is $\mathcal{T}_1$-unsatisfiable, then $Spec(\mathcal{T}_1, \varphi_1) \cap Spec(\mathcal{T}_2, \varphi_2) = \emptyset$, so suppose φ_1 is $\mathcal{T}_1$-satisfiable.

If $\mathbf{minmod}_{\mathcal{T}_1}(\varphi_1) = n$ for some $n \in \mathbb{N}^*$, then $Spec(\mathcal{T}_1, \varphi_1) = \{m \in \mathbb{N}^* \mid m \geq n\} \cup \{\aleph_0\}$, since $\mathcal{T}_1$ is smooth. In this case, we have $Spec(\mathcal{T}_1, \varphi_1) \cap Spec(\mathcal{T}_2, \varphi_2) = \emptyset$ if and only if φ_2 has no $\mathcal{T}_2$-interpretations of size at least n, which happens exactly when $\varphi_2 \wedge \neq (x_1, \ldots, x_n)$ is $\mathcal{T}_2$-unsatisfiable (where the variables x_i are fresh).

If $\mathbf{minmod}_{\mathcal{T}_1}(\varphi_1) = \aleph_0$, then $Spec(\mathcal{T}_1, \varphi_1) = \{\aleph_0\}$. In this case, we have $Spec(\mathcal{T}_1, \varphi_1) \cap Spec(\mathcal{T}_2, \varphi_2) = \emptyset$ if and only if φ_2 is not satisfied by any infinite $\mathcal{T}_2$-interpretation, which we can check since $\mathcal{T}_2$ is infinitely decidable. □

K Proofs concerning Example 2

Proposition 25. *The theory $\mathcal{T}^h_\infty$ is not stably-infinite.*

Proof. P_1 is only satisfied by a $\mathcal{T}^h_\infty$-interpretation with one element. □

Proposition 26. *The theory $\mathcal{T}^h_\infty$ is not strongly polite.*

Proof. Follows from Proposition 25, which implies $\mathcal{T}^h_\infty$ is not smooth. □

Proposition 27. *The theory $\mathcal{T}_\infty$ is not strongly polite.*

Proof. $\mathcal{T}_\infty$ was proven in [25] not to be strongly finitely witnessable, so the result follows. □

Proposition 28. *The theory $\mathcal{T}^h_\infty$ is not shiny.*

Proof. Follows from Proposition 25, which implies $\mathcal{T}^h_\infty$ is not smooth. □

Proposition 29. *The theory $\mathcal{T}_\infty$ is not shiny.*

Proof. $\mathcal{T}_\infty$ does not have the finite model property, and therefore is not shiny, since it is not trivial but has no finite models. □

Proposition 30. *The theory $\mathcal{T}^h_\infty$ does not have computable finite spectra.*

Proof. For all $n \geq 2$, we have $1 \in Spec(\mathcal{T}^h_\infty, P_n)$ if and only if $h(n) = 0$, which is undecidable by assumption. □

Proposition 31. *The theory $\mathcal{T}_\infty$ has a computable minimal model function.*

Proof. Proven in [24]. □

Proposition 32. *The theory $\mathcal{T}_\infty$ is smooth.*

Proof. Obvious, but proven in [25]. □

Proposition 33. *The theory $\mathcal{T}_\infty^h$ is decidable.*

Proof. Given a quantifier-free formula φ, assume without loss of generality φ is a cube and write $\varphi = \varphi_1 \wedge \varphi_2$, where φ_1 contains only equalities and disequalities, and φ_2 contains only the predicates P_n and their negations. We state φ is $\mathcal{T}_\infty^h$-satisfiable if, and only if, φ_1 is satisfiable in equational logic, φ_2 does not contain a predicate and its negation, and if φ_2 contains the literal P_1, then φ_1 has a model in equational logic of size 1 and φ_2 contains no other positive literals.

If φ is $\mathcal{T}_\infty^h$-satisfiable so is φ_1, and by forgetting the predicates on a $\mathcal{T}_\infty^h$-interpretation that satisfies φ_1 we get an interpretation in equational logic that satisfies that formula. Of course φ_2 is also satisfiable, and then we cannot have both P_n and $\neg P_n$ in φ_2. Finally, if φ_2 contains the literal P_1, any $\mathcal{T}_\infty^h$-interpretation $\mathcal{A}$ that satisfies φ must satisfy P_1 as well, and by the axiom $P_1 \rightarrow \psi_{=1}$ of $\mathcal{T}_\infty^h$ we get $\mathcal{A}$ has only one element; by forgetting its predicates, we get an interpretation in equational logic that satisfies φ_1 and has only one element. Furthermore, by the axioms $P_1 \rightarrow \neg P_n$ we get $\mathcal{A}$ satisfies $\neg P_n$ for $n \geq 2$, and thus φ_2 cannot contain the positive literals P_n.

For the reciprocal, start by assuming φ_2 does not contain the literal P_1, and since φ_1 is satisfiable in equational logic it has an infinite model $\mathcal{A}$, as equational logic is stably-infinite; turn $\mathcal{A}$ into a Σ_P^n-interpretation by setting $P_n^{\mathcal{A}}$ to true if the literal P_n occurs in φ_2, and otherwise to false, meaning it satisfies φ_2 as well (this is possible because φ_2 cannot contain both a literal and its negation). $\mathcal{A}$ then satisfies the axioms $P_1 \rightarrow \psi_{=1}$ and $P_1 \rightarrow \neg P_n$ vacuously, as it does not satisfy P_1; and it satisfies the axioms $P_n \rightarrow \psi_{\geq m}$ for $h(n) = 1$ since it is infinite and therefore satisfies all $\psi_{\geq m}$. If φ_2 contains P_1, we take an interpretation $\mathcal{A}$ in equational logic that satisfies φ_1 with only one element, and make it into a Σ_P^n-interpretation by setting $P_1^{\mathcal{A}}$ to true, and all other $P_n^{\mathcal{A}}$ to false (so $\mathcal{A}$ satisfies φ_2 as well, since that formula cannot contain any positive literals other than P_1). It satisfies $P_1 \rightarrow \psi_{=1}$ as it contains only one element, $P_1 \rightarrow \neg P_n$ as all P_n different from P_1 are false in $\mathcal{A}$, and $P_n \rightarrow \psi_{\geq m}$ (for $n \geq 2$ with $h(n) = 1$) vacuously. □

Proposition 34. *The theory $\mathcal{T}_\infty^h$ is infinitely decidable.*

Proof. Given a quantifier-free formula φ, we assume, without loss of generality, that φ is a cube; write then $\varphi = \varphi_1 \wedge \varphi_2$, where φ_1 contains only equalities and disequalities, and φ_2 contains only the predicates P_n and their negations. As proven in Proposition 33, $\mathcal{T}_\infty^h$ is decidable, so we only need to worry whether $\mathcal{T}_\infty^h$-satisfiable φ are satisfied by an infinite $\mathcal{T}_\infty^h$-interpretation: we state that such φ have infinite $\mathcal{T}_\infty^h$-models if and only if P_1 does not appear as a literal in φ_2.

One direction is obvious: if φ_1 contains P_1 and the $\mathcal{T}^h_\infty$-interpretation $\mathcal{A}$ satisfies φ, by the axiom $P_1 \rightarrow \psi_{=1}$ we get $\mathcal{A}$ can only have one element. If P_1 is not in φ_2 we proceed as in the proof of Proposition 33: as equational logic is stably-infinite, take an infinite interpretation $\mathcal{A}$ that satisfies φ_1; and set $P_n^{\mathcal{A}}$ to true if and only if the positive literal P_n occurs in φ_2, so $\mathcal{A}$, as a Σ_P^n-interpretation, satisfies φ_2 as well. □