Two generalizations of shininess

Guilherme V. Toledo and Yoni Zohar

Bar-Ilan University, Israel

Abstract. Shininess is one of the main theory combination methods in Satisfiability Modulo Theories. We present two infinite families of generalizations of it: one involving a property between gentleness and shininess, which helps us see gentleness itself as a variation of shininess; and another having no further computability requirements on the weaker theory, in addition dropping the need for smoothness. We prove these methods sharp, and compare them with existing methods.[1]

1 Introduction

There are three combination methods that have become widespread. First and foremost, the Nelson–Oppen method [7], requiring the two theories to be combined to be decidable and stably infinite. In a close second place, the polite method, which requires one theory to be decidable and strongly polite while the other is allowed to be simply decidable [9,5]; it is known, thanks to [1,2] that, for decidable theories, strong politeness is equivalent to shininess [10], a property that is more difficult to implement but is easier to understand. And, finally, the gentle method [4], which requires a theory to be gentle and the other to be decidable and have computable finite spectra [11].

These methods are all *sharp* [8], e.g.: a decidable theory that can be combined with all decidable theories must be strongly polite or, equivalently, shiny. The machinery necessary to prove sharpness of combination methods, involving the notions of test theories and a Galois connection, also provides a reasonably effective approach to finding new combination methods. We showcase this feature by producing two infinite families of methods closely resembling shiny combination.

The first is inspired by a method from [8] involving the notions of n-decidability and n-shininess. It offers a new way of understanding gentleness as a variation of shininess, helping contextualize that method. The second comes from dropping smoothness from the definition of shininess, a quite restrictive and often undesirable requirement, which imposes no computability requisites other than decidability on the weaker theory to be combined (much like the Nelson–Oppen and strongly polite combination methods).

We start with some basic notions necessary for the rest of the paper in Section 2. Sections 3 and 4 give two sets of generalizations of shininess. In Section 5 we prove that these new combination results are sharp, in the sense that they

[1] Funding: U.S-Israel Binational Science Foundation (BSF) grant 2024049, and Israel Science Foundation (ISF) grants 619/21 and 209/25.

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

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

Fig. 1. Cardinality formulas

cannot be theoretically strengthened. A brief conclusion and plans for future work can be found in Section 6.

2 Preliminaries

We denote $\mathbb{N}\setminus\{0\}$ by $\mathbb{N}^*$, for any $S\subseteq\mathbb{N}$ the set $\mathbb{N}\cup\{\aleph_0\}$ by S_ω, and $\{n\in\mathbb{N}_\omega : p\leq n\leq q\}$ by $[p,q]$.

2.1 First-order logic

We review standard definitions and notations regarding first-order logic. For more details, see, e.g., [3]. A **signature** Σ is an at most countable infinite set of function and predicate symbols with fixed arities and containing equality. Two signatures are **disjoint** if they share only the equality symbol; Fixed a countably infinite set of variables, we define literals, formulas and sentences (formulas with no free variables) as usual; let $QF(\Sigma)$ be the set of quantifier-free Σ-formulas.

A **Σ-interpretation** $\mathcal{A}$ is a non-empty set $dom(\mathbb{A})$ (its **domain**) equipped with functions $f^{\mathbb{A}}$ and predicates $P^{\mathbb{A}}$ over $dom(\mathcal{A})$ of the same arity as their corresponding symbols, and values $x^{\mathcal{A}}\in dom(\mathcal{A})$ assigned to all variables x. We define the truth-value of a formula φ in an interpretation $\mathcal{A}$ in the usual way, and if that value is true we write $\mathcal{A}\vDash\varphi$. In Figure 1 the formula $\neq(x_1,\dots,x_n)$ states the variables x_1 through x_n all take different values (for $n\geq 2$, when $n=1$ we can take it to be any tautology); the formulas $\psi_{\geq n}$, $\psi_{=n}$ and $\psi_{\leq n}$ state that there are at least, exactly, and at most n elements in the domain, respectively.

A **Σ-theory** is a set of Σ-sentences, called its **axioms**, and if an interpretation satisfies all of them it is said to be a **$\mathcal{T}$-interpretation**. A formula is **$\mathcal{T}$-satisfiable** if there is at least one $\mathcal{T}$-interpretation that satisfies it. A Σ-theory $\mathcal{T}$ is **decidable** when there is an algorithm that takes a quantifier-free Σ-formula and returns whether or not it is $\mathcal{T}$-satisfiable. Given signatures Σ_1 and Σ_2, $\Sigma_1\sqcup\Sigma_2$ is their disjoint union, which is the union of their symbols after renaming them to avoid overriding; and given a Σ_1-theory $\mathcal{T}_1$ and a Σ_2-theory $\mathcal{T}_2$, their **disjoint combination** $\mathcal{T}_1\sqcup\mathcal{T}_2$ is the $\Sigma_1\sqcup\Sigma_2$-theory with axiomatization equal to the union of $\mathcal{T}_1$ and $\mathcal{T}_2$ (after possible renaming).

The **spectrum** $Spec_{\mathcal{T}}(\phi)$ of a formula φ in a theory $\mathcal{T}$ is the subset of $\mathbb{N}_\omega$ of those elements n such that there exists a $\mathcal{T}$-interpretation $\mathcal{A}$ with $\mathcal{A}\vDash\varphi$ and $|dom(\mathcal{A})|=n$. We use the following ubiquitous result in first-order logic.

Theorem 1 ([6, Theorems 2.3.4 and 2.3.7]). *Let Δ be a set of formulas over a countable signature, and let $\kappa\geq\aleph_0$. Then, Δ is satisfied by an interpretation of size κ if and only if it is satisfied by an interpretation of size $\aleph_0$.*

The following is a classical result we make use of, obtained as an immediate corollary from compactness.

Corollary 1. *If ϕ is $\mathcal{T}$-satisfiable but $\aleph_0 \notin Spec_{\mathcal{T}}(\phi)$ then $Spec_{\mathcal{T}}(\phi)$ is finite.*

2.2 Theory combination and its Galois connection

$\mathcal{T}$ is said to have **computable finite spectra** [11] when there is an algorithm that takes a quantifier-free formula ϕ and an $n \in \mathbb{N}^*$, returning whether $n \in Spec_{\mathcal{T}}(\phi)$. A theory $\mathcal{T}$ is **stably infinite** if $Spec_{\mathcal{T}}(\phi) \neq \varnothing$, for a quantifier-free ϕ, implies $\aleph_0 \in Spec_{\mathcal{T}}(\phi)$. A theory $\mathcal{T}$ is **smooth** when, for every quantifier-free ϕ, $\mathcal{T}$-interpretation $\mathcal{A} \vDash \phi$ and $\kappa \geq |dom(\mathcal{A})|$, there is a $\mathcal{T}$-interpretation $\mathcal{B} \vDash \phi$ with $|dom(\mathcal{B})| = \kappa$. $\mathcal{T}$ has the **finite model property** when, for every $\mathcal{T}$-satisfiable quantifier-free formula ϕ, $Spec_{\mathcal{T}}(\phi) \cap \mathbb{N}^* \neq \varnothing$. And, if $\mathcal{T}$ is decidable, its **minimal model function** is the function $\mathrm{minmod}_{\mathcal{T}}$ that takes a quantifier-free formula ϕ and returns $\min Spec_{\mathcal{T}}(\phi)$, if that set is not empty (and any value in $\mathbb{N}_\omega$ otherwise). A theory is then said to be **shiny** [10] when it is smooth, has the finite model property, and its minimal model function is computable.

Proposition 1. *A theory is shiny if and only if there is an algorithm that takes a quantifier-free formula ϕ and returns an $n \in \mathbb{N}$ such that $Spec_{\mathcal{T}}(\phi) = \varnothing$ if $n = 0$, or $Spec_{\mathcal{T}}(\phi) = [n, \aleph_0]$ otherwise.*

Proof. Suppose $\mathcal{T}$ is shiny: if ϕ is not $\mathcal{T}$-satisfiable we return 0; if it is, we make $n = \mathrm{minmod}_{\mathcal{T}}(\phi)$. This value is in $\mathbb{N}$ because $\mathcal{T}$ has the finite model property; and $Spec_{\mathcal{T}}(\phi) = [n, \aleph_0]$ because $\mathcal{T}$ is smooth. For the reciprocal, we make $\mathrm{minmod}_{\mathcal{T}}(\phi) = n$; $\mathcal{T}$ has the finite model property because this value is always in $\mathbb{N}$; and it is smooth because $Spec_{\mathcal{T}}(\phi) = [n, \aleph_0]$ and due to Theorem 1. □

A theory $\mathcal{T}$ is **gentle** ([4], though we use the equivalent definition of [11]) when there is an algorithm that takes a quantifier-free formula ϕ and returns a pair (S, b), where $S \subset \mathbb{N}^*$ is finite and b is a Boolean, such that: if b is true, $Spec_{\mathcal{T}}(\phi) = S$; and if b is false, $Spec_{\mathcal{T}}(\phi) = \mathbb{N}_\omega \setminus S$. Using the notations for sets[2] of theories found in Table 1 we present the following proposition, whose proof can be found in Proposition 6.1 of [8]. It says that for decidable theories, shininess implies gentleness, that in turn implies having computable finite spectra.

Proposition 2. $\mathfrak{T}_{shiny} \subseteq \mathfrak{T}_{gentle} \subseteq \mathfrak{T}_{CFS} \subseteq \mathfrak{T}$.

One important tool in what is to come is the following result from [4,8] .

Lemma 1 (Fontaine's lemma). *$\mathcal{T}_1 \sqcup \mathcal{T}_2$ is decidable iff the following problem is decidable: given conjunctions of literals φ_1 and φ_2, determine whether $Spec_{\mathcal{T}_1}(\varphi_1) \cap Spec_{\mathcal{T}_2}(\varphi_2) = \varnothing$.*

[2] These, in principle, would not constitute sets, but proper classes. But, we can select all symbols for our (countable) signatures from a countable pool of symbols to ensure these are sets.

Symbol	Description (decidable and ...)	Definition
$\mathfrak{T}$	—	Section 2
$\mathfrak{T}_{\text{shiny}}$	Shiny	
$\mathfrak{T}_{gentle}$	Gentle	
$\mathfrak{T}_{\text{SI}}$	Stably infinite	
$\mathfrak{T}_{\text{CFS}}$	Computable finite spectra	
$\mathfrak{T}_{\text{S-d}}$	S-decidable	Section 3
$\mathfrak{T}_{\text{S-s}}$	S-shiny	
$\mathfrak{T}_{\text{S-r}}$	S-restricted	Section 4
$\mathfrak{T}_{\text{S-f}}$	S-final	

Table 1. Notations for sets of theories.

Two theories are **combinable** if their disjoint union is decidable, and two sets of theories are combinable if any theory in the first is combinable with any theory in the second; if one of these sets is a singleton, we may simply say the theory it contains is combinable with the other set. For each $X \subseteq \mathfrak{T}$, let $G(X)$ be the set of all decidable theories that are combinable with X, that is, the largest set of decidable theories that is combinable with X.

If $X \subseteq Y$, as every theory combinable with Y is also combinable with X, we get $G(Y) \subseteq G(X)$; and X is combinable with Y iff Y is combinable with X, meaning $X \subseteq G(Y)$ iff $Y \subseteq G(X)$. Both these facts make G an **antitone Galois connection** with respect to $\subseteq$. A **combination theorem** typically says two sets of theories are combinable, meaning $X \subseteq G(Y)$ and $Y \subseteq G(X)$. Such a theorem is *sharp* if, in fact, $X = G(Y)$ and $Y = G(X)$. In this case $G(G(X)) = X$ and $G(G(Y)) = Y$, which is the same as saying, in lattice theory, that X and Y are **closed**; for an arbitrary X, $G(G(X))$ is its **closure**, the minimum closed set containing X.

3 First Generalization

We now provide a new combination method, which generalizes shiny combination, and requires the following properties, that parameterize shininess and decidability by a computable set S.

Definition 1. *Take a computable set $S \subseteq \mathbb{N}^*$.*

1. *A decidable theory $\mathcal{T}$ is S-decidable when there is an algorithm that takes a quantifier-free formula ϕ and an $n \in \mathbb{N}^*$ and returns: whether $n \in Spec_{\mathcal{T}}(\phi)$ if $n \in S$; and an arbitrary Boolean value otherwise.*
2. *A decidable theory $\mathcal{T}$ is S-shiny if there is an algorithm that takes a quantifier-free formula ϕ and returns a pair (n, S_0), where $n \in \mathbb{N}$ and S_0 is a finite subset of S (possibly empty), such that: if $n = 0$, $Spec_{\mathcal{T}}(\phi) = S_0$; if $n \neq 0$, $Spec_{\mathcal{T}}(\phi) = S_0 \cup [n, \aleph_0]$.*

Algorithm 1 Pseudocode for showing $\mathcal{T}$ is S-decidable if it is S-shiny.

function SD(ϕ,m)
    $(n, S_0) \leftarrow$ computed for $\mathcal{T}$ and ϕ, according to Definition 1
    if $m \in S_0$ **then return** true
    else
        if $n > 0$ **then**
            if $m \geq n$ **then return** true
        return false

Remark 1. n-decidability and n-shininess were defined in [8] for any $n \in \mathbb{N}^*$. The former corresponds to $\{n\}$-decidability, the latter to $\{n\}$-shininess.

Now, it was shown in [11] that every shiny theory is decidable. This also holds for the new generalizations.

Proposition 3. *An S-shiny theory is S-decidable.*

Proof. Suppose $\mathcal{T}$ is S-shiny, and that we are given a quantifier-free ϕ and an $m \in S$: we want to return whether $m \in Spec_{\mathcal{T}}(\phi)$. As $\mathcal{T}$ is S-shiny, there is an algorithm that takes ϕ and returns a pair (n, S_0) such that $S_0 \subseteq S$ is finite, if $n = 0$ then $Spec_{\mathcal{T}}(\phi) = S_0$, and otherwise $Spec_{\mathcal{T}}(\phi) = S_0 \cup [n, \aleph_0]$. If $n = 0$, $m \in Spec_{\mathcal{T}}(\phi)$ iff $m \in S_0$; otherwise $m \in Spec_{\mathcal{T}}(\phi)$ iff $m \in S_0$ or $m \geq n$. This procedure may be seen in Algorithm 1. ◻

First, we show that these notions are indeed generalizations of decidable and shiny theories, in particular when S is taken to be the empty set. Additionally, we show that when S is taken to be the entire set $\mathbb{N}^*$, we get computable finite spectra and gentleness. This shows gentle and shiny theory combination are two ends of a spectrum: gentleness is the weakest form of S-shininess, and shininess the strongest; reciprocally, having computable finite spectra is the strongest form of S-decidability, and decidability itself the weakest.

Proposition 4. *$\mathfrak{T}_{\varnothing\text{-}d} = \mathfrak{T}$, $\mathfrak{T}_{\mathbb{N}^*\text{-}d} = \mathfrak{T}_{CFS}$, $\mathfrak{T}_{\mathbb{N}^*\text{-}s} = \mathfrak{T}_{gentle}$, and $\mathfrak{T}_{\varnothing\text{-}s} = \mathfrak{T}_{shiny}$.*

Proof. $\mathfrak{T}_{\varnothing\text{-d}} = \mathfrak{T}$ should be clear, as no condition is imposed on its theories other than being decidable. If a theory $\mathcal{T}$ is in $\mathfrak{T}_{\mathbb{N}^*\text{-d}}$ then there is an algorithm that takes a quantifier-free formula ϕ and an $n \in \mathbb{N}^*$ and returns whether $n \in Spec_{\mathcal{T}}(\phi)$, what corresponds precisely to the definition of having computable finite spectra.

That $\mathfrak{T}_{\varnothing\text{-s}} = \mathfrak{T}_{\text{shiny}}$ follows from the fact that in a $\varnothing$-shiny theory $\mathcal{T}$ there is an algorithm that takes a quantifier-free formula ϕ and returns a pair $(n, \varnothing)$ such that: either $n = 0$ and $Spec_{\mathcal{T}}(\phi) = \varnothing$ (meaning ϕ is not $\mathcal{T}$-satisfiable); or $n \in \mathbb{N}^*$ and $Spec_{\mathcal{T}}(\phi) = [n, \aleph_0]$, as expected of a shiny theory by Proposition 1.

Finally, it is clear $\mathbb{N}^*$-shiny theories are gentle. And if a theory $\mathcal{T}$ is gentle we either have $Spec_{\mathcal{T}}(\phi)$ finite, when we output $(0, Spec_{\mathcal{T}}(\phi))$, which is compatible with being an $\mathbb{N}^*$-shiny theory as $Spec_{\mathcal{T}}(\phi) \subseteq \mathbb{N}$; or $Spec_{\mathcal{T}}(\phi)$ is cofinite, and from gentleness we get $S = \mathbb{N}_\omega \setminus Spec_{\mathcal{T}}(\phi)$. In this case take the greatest element n not in $Spec_{\mathcal{T}}(\phi)$, and we have

$$Spec_{\mathcal{T}}(\phi) = ([1, n] \setminus S) \cup [n + 1, \aleph_0],$$

Algorithm 2 Pseudocode for determining non-emptyness of $Spec_{\mathcal{T}_1}(\phi_1) \cap Spec_{\mathcal{T}_2}(\phi_2)$ when $\mathcal{T}_1$ is S-decidable and $\mathcal{T}_2$ is S-shiny.

function SD+SS(ϕ_1, ϕ_2)
  if ϕ_2 is not $\mathcal{T}_2$-satisfiable **then return** false
  $(n, S_0) \leftarrow$ computed for ϕ_2 and $\mathcal{T}_2$, according to Definition 1.
  if $n \neq 0$ and $\phi_1 \wedge \neq (x_1, \dots, x_n)$ is $\mathcal{T}_1$-satisfiable **then** $\triangleright$ $x_1, \dots, x_n$ are fresh
    return true
  for $m \in S_0$ **do**
    if $m \in Spec_{\mathcal{T}_1}(\phi_1)$ **then**
      return true
  return false

also compatible with $\mathcal{T}$ being $\mathbb{N}$-shiny as certainly $[1, n] \setminus S \subseteq \mathbb{N}$, when we output $(n + 1, [1, n] \setminus S)$. □

The following theorem is the main result of this section. It generalizes the shiny combination theorem, according to which the disjoint combination of a shiny theory with a decidable theory is itself decidable [10].

Theorem 2. *If $\mathcal{T}_1$ is S-decidable and $\mathcal{T}_2$ is S-shiny for some S, and both $\mathcal{T}_1$ and $\mathcal{T}_2$ are decidable, then $\mathcal{T}_1 \sqcup \mathcal{T}_2$ is decidable.*

Proof (sketch). By Lemma 1, it suffices to show a decision procedure that decides, given conjunctions of literals ϕ_1 and ϕ_2 whether $Spec_{\mathcal{T}_1}(\phi_1) \cap Spec_{\mathcal{T}_2}(\phi_2)$ is empty. This decision procedure is shown in Algorithm 2. □

We now study how S-decidability and S-shininess change with S.

Proposition 5.

1. *If $S \subsetneq T$, then $\mathfrak{T}_{T\text{-}d} \subsetneq \mathfrak{T}_{S\text{-}d}$ and $\mathfrak{T}_{S\text{-}s} \subsetneq \mathfrak{T}_{T\text{-}s}$.*
2. *If S and T are incomparable (i.e. $S \nsubseteq T$ and $T \nsubseteq S$), so are $\mathfrak{T}_{S\text{-}d}$ and $\mathfrak{T}_{T\text{-}d}$, and also $\mathfrak{T}_{S\text{-}d}$ and $\mathfrak{T}_{T\text{-}d}$.*

Thus, the collections of S-decidable and S-shiny theories, for computable $S \subseteq \mathbb{N}^*$, have the structure of the computable subsets of $\mathbb{N}^*$, and its dual. More precisely, we have Figure 2: the Galois connection G is simply the reflection over the dashed line. Below it we find the many variations of S-shininess, having the structure of the computable subsets of $\mathbb{N}^*$, with shininess at the bottom and gentleness on top. Above the dashed line we find the variations of S-decidability, having the structure of the dual of the computable subsets of $\mathbb{N}^*$, with computable finite spectra at the bottom and decidability at the top.

We conclude the first generalization of shininess with an example:

Example 1. Consider a one-sorted variant of the SMT-LIB theory of bit-vectors, where the single domain includes all bit-vectors of the same width. The set X of all possible sizes of models is comprised of all the powers of two, a decidable set. Thus, this variant is X-decidable, and can be combined with any X-shiny theory, thanks to Theorem 2.

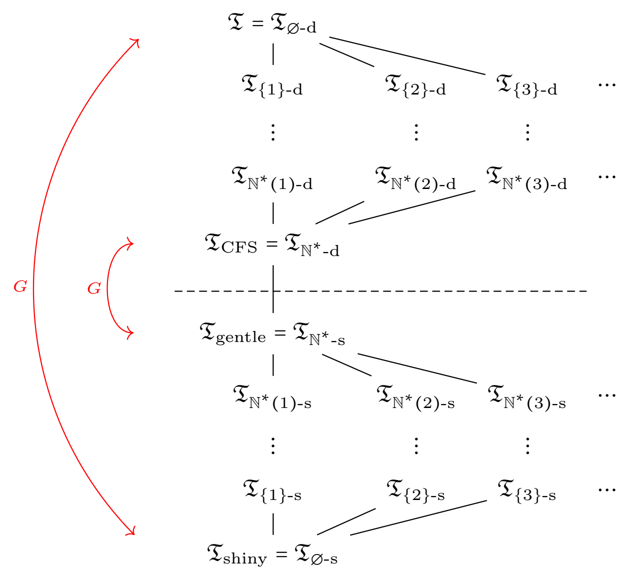

Fig. 2. The rich universe of S-decidable and S-shiny theory combination. For readability, $\mathbb{N}^*(n)$ denotes $\mathbb{N}^* \setminus \{n\}$.

4 Second Generalization

We now present our second generalization of the shiny combination theorem. It allows one to combine two sets of theories: the first, weaker set, has no requirement on the existence of algorithms for it (other than decidability, of course). The second requires an algorithm for computing a variant of a spectra.

Definition 2. *Take a computable, infinite set $S \subseteq \mathbb{N}^*$.*

1. *A decidable theory $\mathcal{T}$ is S-restricted if, for every quantifier-free formula ϕ, $Spec_{\mathcal{T}}(\phi) \setminus S \neq \varnothing$ implies $\aleph_0 \in Spec_{\mathcal{T}}(\phi)$.*
2. *A decidable theory is S-final if there is an algorithm that takes a quantifier-free formula ϕ and returns 0 if ϕ is not $\mathcal{T}$-satisfiable, and otherwise an $n \in \mathbb{N}^*$ such that*

$$Spec_{\mathcal{T}}(\phi) \cap S_\omega = [n, \aleph_0] \cap S_\omega.$$

Let us first see relations between these properties and others.

Proposition 6. *$\mathfrak{T}_{S\text{-}f} \subseteq \mathfrak{T}_{SI} \subseteq \mathfrak{T}_{S\text{-}r}$, for any S.*

Proof. Suppose $\mathcal{T}$ is S-final, and take a quantifier-free ϕ: if ϕ is not $\mathcal{T}$-satisfiable we have nothing to prove; otherwise our algorithm gives an $n \in \mathbb{N}^*$ such that $Spec_{\mathcal{T}}(\phi) \cap S_\omega = [n, \aleph_0] \cap S_\omega$, and since $\aleph_0 \in S$, $\mathcal{T}$ is stably infinite. If $\mathcal{T}$ is stably infinite we get that if $Spec_{\mathcal{T}}(\phi)$ is not empty then $\aleph_0 \in Spec_{\mathcal{T}}(\phi)$, meaning $\mathcal{T}$ is trivially S-restricted. $\square$

Algorithm 3 Pseudocode for determining non-emptiness of $Spec_{\mathcal{T}_1}(\phi_1) \cap Spec_{\mathcal{T}_2}(\phi_2)$ when $\mathcal{T}_1$ is S-restricted and $\mathcal{T}_2$ is S-final.

function SR+SF(ϕ_1, ϕ_2)
  if ϕ_2 is not $\mathcal{T}_2$-satisfiable **then return** false
  $n \leftarrow$ computed for ϕ_2 and $\mathcal{T}_2$, according to Definition 2.
  if $\phi_1 \wedge \neq(x_1, \dots, x_n)$ is $\mathcal{T}_1$-satisfiable **then** ▷ $x_1, \dots, x_n$ are fresh
    return true
  return false

Proposition 7. *$\mathfrak{T}_{\mathbb{N}^*\text{-}r} = \mathfrak{T}$, and $\mathfrak{T}_{\mathbb{N}^*\text{-}f} = \mathfrak{T}_{shiny}$.*

Proof. An $\mathbb{N}^*$-restricted theory imposes no further requirements on the theory other than being decidable, as we cannot have $Spec_{\mathcal{T}}(\phi) \setminus \mathbb{N}^*$ not empty unless $Spec_{\mathcal{T}}(\phi)$ already contains $\aleph_0$. $\mathbb{N}^*$-finality implies there is an algorithm that takes a quantifier-free ϕ and returns 0 if ϕ is not $\mathcal{T}$-satisfiable, and otherwise an $n \in \mathbb{N}^*$ such that $Spec_{\mathcal{T}}(\phi) = [n, \aleph_0]$, exactly the definition of shininess in Proposition 1. □

Now, we present our second new combination theorem.

Theorem 3. *Let $\mathcal{T}_1$ be S-restricted, and $\mathcal{T}_2$ be S-final. Then $\mathcal{T}_1 \sqcup \mathcal{T}_2$ is decidable.*

Proof (sketch). Much like Theorem 2, but this time using Algorithm 2. □

Example 2. The first example of a finitely witnessable (property related to politeness) theory that is not strongly finitely witnessable (property related to strong politeness, equivalent to shininess over decidable theories) was $\mathcal{T}^{\infty}_{even}$: this theory has only interpretations of even size, or infinite ones. Yet, in a sense, it is a surprisingly well-behaved theory: if S is the set of positive even numbers, then it should not be too hard to prove that $Spec_{\mathcal{T}^{\infty}_{even}}(\phi) = [n, \aleph_0] \cap S_{\omega}$, where n is the minimal size of any first-order model of ϕ, which we assume to be quantifier-free. This is almost the definition of shininess, except for the pesky intersection with S. Using the notions just defined, $\mathcal{T}^{\infty}_{even}$ is seen in fact to be S-shiny, and thus combinable with any S-restricted theory.

As we did for S-decidability and S-shininess, we analyze how the sets of S-restricted and S-final theories change as S changes.

Proposition 8. *1. If $S \subsetneq T$, $\mathfrak{T}_{S\text{-}r} \subsetneq \mathfrak{T}_{T\text{-}r}$, and $\mathfrak{T}_{T\text{-}f} \subsetneq \mathfrak{T}_{T\text{-}f}$.*
2. If S and T are incomparable, so are $\mathfrak{T}_{S\text{-}r}$ and $\mathfrak{T}_{T\text{-}r}$, and $\mathfrak{T}_{S\text{-}f}$ and $\mathfrak{T}_{T\text{-}f}$.

5 Sharpness of the new methods

In this section, we prove that the new combination theorems Theorems 2 and 3 are, in fact, sharp. For the former, we prove that every theory that can be combined with all S-decidable theories must be S-shiny, and that every theory that can be combined with all S-shiny theories must be S-decidable. Similar sharpness results are proven for the latter. We utilize the method of *test theories*, introduced in [8].

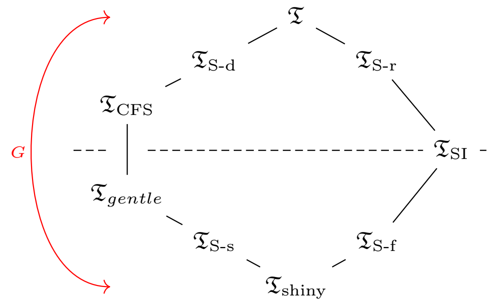

Fig. 3. The general picture involving the notions defined in this paper.

5.1 Test Theories

The idea behind test theories goes more or less as follows: we want to prove that if a theory $\mathcal{T}$ is combinable with every theory of X, then it admits some property; this could be done by looking at all $\mathcal{T} \sqcup \mathcal{T}'$, for $\mathcal{T}' \in X$. But considering all theories of X is rarely feasible, and is certainly not very effective. So, from X, we select special theories, *test theories*, such that the satisfiability problem of the combination of $\mathcal{T}$ and a test theory directly relates to properties from $\mathcal{T}$. In this sense, the set of test theories will be a very difficult set of theories to combine with $\mathcal{T}$, as being combinable with them forces $\mathcal{T}$ to have a given property.

We need 5 sets of test theories, three found in [8], and two new. Let $\Sigma_P^{\mathbb{N}}$ be the signature with no function symbols, and nullary predicate symbols $\{P_n : n \in \mathbb{N}^*\}$: all of our test theories will be $\Sigma_P^{\mathbb{N}}$-theories, and are described in Table 2, which relies on an undecidable set $U \subseteq \mathbb{N}^*$, and a function $F : \mathbb{N}^* \to \mathbb{N}^* \cup \{\aleph_0\}$ such that $\{(m,n) \in {\mathbb{N}^*}^2 : n \leq F(m)\}$ is decidable, but $\{n \in \mathbb{N}^* : F(n) = \aleph_0\}$ is not.[3]

The three test theories from [8] are $\mathcal{T}_n^m$, $\mathcal{T}_=^P$ and $\mathcal{T}_\leq^S$. $\mathcal{T}_n^m$ is a $\Sigma_P^{\mathbb{N}}$-theory with only interpretations of size m or n, and where P_k only has interpretations of size m whenever $k \in U$; so $Spec_{\mathcal{T}_n^m}(P_k) = \{n, m\}$ if $k \notin U$, and $\{m\}$ otherwise. A $\mathcal{T}_=^P$-interpretation where P_n is true must have exactly n elements; so $Spec_{\mathcal{T}_=^P}(P_n) = \{n\}$. The theory $\mathcal{T}_\leq^S$ has only interpretations of cardinality in S_ω, and if P_n is true and $F(n)$ is finite, then we have at most $F(n)$ elements; so $Spec_{\mathcal{T}}(P_n) = S \cap [1, F(n)]$ if $F(n)$ is finite, and S_ω otherwise.

The two new test theories are $\mathcal{T}_=^S$ and $\mathcal{T}_{\text{S-r}}^{m,n}$. The theory $\mathcal{T}_=^S$ is similar to $\mathcal{T}_=^P$, with a twist: if P_n is true (what can only happen for $n \in S$) then the interpretation has size n. So, $Spec_{\mathcal{T}_=^S}(P_n) = \varnothing$ if $n \notin S \subseteq \mathbb{N}^*$, and $\{n\}$ otherwise. The theory $\mathcal{T}_{\text{S-r}}^{m,n}$ is quite similar to $\mathcal{T}_n^m$, and will be used to test whether the spectrum of a formula is entirely contained in a computable, infinite set $S \subseteq \mathbb{N}^*$. It depends on two numbers m and n, with $m > n$ and $n \notin S$, and its interpretations must: have cardinality in S or equal to n; have cardinality n or at least m; and, if they have cardinality n, then P_k can only be true if $k \in U$. To summarize it, $Spec_{\mathcal{T}_{\text{S-r}}^{m,n}}(P_k) = S_\omega \cap (\{n\} \cup [m, \aleph_0])$ if $k \in U$, and $S_\omega \cap [m, \aleph_0]$ otherwise.

Theorem 4. *Table 2 is correct: each theory satisfies the properties listed for it.*

[3] An example for F assumes an enumeration of Turing machines, takes the index of a machine, and returns the number of steps it takes to halt, or $\aleph_0$ if it doesn't.

Theory	Details
$\mathcal{T}_n^m$	<u>Axiomatization</u> $\{\psi_{=n} \vee \psi_{=m}\} \cup \{P_k \rightarrow \psi_{=m} : k \in U\} \cup P_{\neq}$ <u>Properties</u> Decidable, S-decidable (if $n \notin S$), S-restricted (if $m, n \in S$) $m \notin Spec_{\mathcal{T}}(\phi), n \in Spec_{\mathcal{T}}(\phi) \Rightarrow [\phi \wedge P_k$ is $\mathcal{T} \sqcup \mathcal{T}_n^m$-SAT $\Leftrightarrow k \notin U]$
$\mathcal{T}_=^P$	<u>Axiomatization</u> $\{P_n \rightarrow \psi_{=n} : n \in \mathbb{N}^*\}$ <u>Properties</u> Decidable, S-decidable (for any S) $\phi \wedge P_n$ is $\mathcal{T} \sqcup \mathcal{T}_=^P$-SAT $\Leftrightarrow n \in Spec_{\mathcal{T}}(\phi)$
$\mathcal{T}_\leq^S$	<u>Axiomatization</u> $\{P_n \rightarrow \psi_{\leq F(n)} : F(n) \in \mathbb{N}\} \cup \{\neg\psi_{=n} : n \notin S\} \cup P_{\neq}$ <u>Properties</u> Decidable, $\overline{S}$-decidable, S-restricted $Spec_{\mathcal{T}}(\phi) \neq \varnothing, S \cap Spec_{\mathcal{T}}(\phi) = \varnothing \Rightarrow [\phi \wedge P_k$ is $\mathcal{T} \sqcup \mathcal{T}_\leq^S$-SAT $\Leftrightarrow F(k) = \aleph_0]$
$\mathcal{T}_=^S$	<u>Axiomatization</u> $\{P_n \rightarrow \psi_{=n} : n \in S\} \cup \{\neg P_n : n \notin S\} \cup P_{\neq}$ <u>Properties</u> Decidable, S-shiny, S-restricted $k \in S \Rightarrow [\phi \wedge P_k$ is $\mathcal{T} \sqcup \mathcal{T}_=^S$-SAT $\Leftrightarrow k \in Spec_{\mathcal{T}}(\phi)]$
$\mathcal{T}_{S-r}^{m,n}$	<u>Axiomatization</u> $\{\neg\psi_{=k} : k \notin S \cup \{n\}\} \cup \{\psi_{=n} \vee \psi_{\geq m}\} \cup \{\psi_{=n} \rightarrow \neg P_k : k \notin U\} \cup P_{\neq}$ <u>Properties</u> Decidable, S-final $n \in Spec_{\mathcal{T}}(\phi), m = M_{\mathcal{T}}(\phi) \Rightarrow [\phi \wedge P_k$ is $\mathcal{T} \sqcup \mathcal{T}_{S-r}^{m,n}$-SAT $\Leftrightarrow k \in U]$

Table 2. Test theories and their properties; $S \subseteq \mathbb{N}^*$; for the first and last theories, $m > n$; for the last, $n \notin S$. We define $P_{\neq} = \{P_i \rightarrow \neg P_j : i \neq j\}$ and $M_{\mathcal{T}}(\phi) = 1 + \max Spec_{\mathcal{T}}(\phi)$ to make the table more readable.

5.2 Sharpness for *S*-shininess

We now prove that Theorem 2 is sharp, that is: Given S, all decidable theories that can be combined with all S-shiny theories must be S-decidable, and all decidable theories that can be combined with all S-decidable theories must be S-shiny. Formally, we utilize operator G from Section 2.2.

Theorem 5. *$G(\mathfrak{T}_{S\text{-}s}) \subseteq \mathfrak{T}_{S\text{-}d}$ and $G(\mathfrak{T}_{S\text{-}d}) \subseteq \mathfrak{T}_{S\text{-}s}$.*

Proof. 1. Suppose that $\mathcal{T} \in G(\mathfrak{T}_{\text{S-s}})$, and take a quantifier-free formula ϕ: since $\mathcal{T}_=^S$ is decidable and S-shiny (see Theorem 4), for any $n \in S$ we have $n \in Spec_{\mathcal{T}}(\phi)$ if and only if $\phi \wedge P_n$ is $\mathcal{T} \sqcup \mathcal{T}_=^S$-satisfiable, a problem for which we have an algorithm since $\mathcal{T} \sqcup \mathcal{T}_=^S$ is decidable; we thus have $\mathcal{T}$ is S-decidable.
2. Suppose now that $\mathcal{T} \in G(\mathfrak{T}_{\text{S-d}})$. First we show that $\mathcal{T}$ must have computable finite spectra. $\mathcal{T}_=^P$ is decidable and S-decidable by Theorem 4; this means

Algorithm 4 Pseudocode for showing $\mathcal{T}$ is S-shiny if $\mathcal{T} \in G(\mathfrak{T}_{\text{S-d}})$.

1: **function** SS(ϕ)
2: $\quad k \leftarrow 1$
3: $\quad S_0 \leftarrow \varnothing$
4: $\quad$ **while** True **do**
5: $\quad\quad$ **if** $k \in \mathit{Spec}_{\mathcal{T}}(\phi)$ **then**
6: $\quad\quad\quad$ **if** $k \notin S$ **then return** (k, S_0)
7: $\quad\quad\quad S_0 \leftarrow S_0 \cup \{k\}$
8: $\quad\quad\quad k \leftarrow k + 1$
9: $\quad\quad$ **else**
10: $\quad\quad\quad$ **if** $\phi \wedge \neq (x_1, \ldots, x_k)$ is not $\mathcal{T}$-satisfiable **then return** $(0, S_0)$
11: $\quad\quad\quad k \leftarrow k + 1$

$\mathcal{T} \sqcup \mathcal{T}_{=}^{P}$ is decidable, and $k \in \mathit{Spec}_{\mathcal{T}}(\phi)$ (for a quantifier-free ϕ) if and only if $\phi \wedge P_k$ is $\mathcal{T} \sqcup \mathcal{T}_{=}^{P}$-satisfiable, proving $\mathcal{T}$ indeed has computable finite spectra. We now prove $\mathcal{T}$ is S-shiny by showing, in Algorithm 4, the algorithm that outputs (n, S_0). Checking (for $k \in \mathbb{N}^*$) whether $k \in S$ in line 5 is possible because S is computable, whether $k \in \mathit{Spec}_{\mathcal{T}}(\phi)$ in line 4 is possible because we just proved $\mathcal{T}$ has computable finite spectra, and whether $\phi \wedge \neq (x_1, \ldots, x_k)$ is $\mathcal{T}$-satisfiable in line 10 is possible because $\mathcal{T}$ is decidable. We notice that S_0 is nothing but $\mathit{Spec}_{\mathcal{T}}(\phi) \cap [1, k-1]$. Now, to see that our algorithm stops, there are two cases to consider: whether $\mathit{Spec}_{\mathcal{T}}(\phi) \setminus S$ is empty or not.

If it is empty, and $\mathit{Spec}_{\mathcal{T}}(\phi)$ is infinite, using Corollary 1 we get $\aleph_0 \in \mathit{Spec}_{\mathcal{T}}(\phi)$, and that leads to a contradiction: $\mathit{Spec}_{\mathcal{T}}(\phi) \setminus S = \varnothing$, $\aleph_0 \notin S$ and $\aleph_0 \in \mathit{Spec}_{\mathcal{T}}(\phi)$ cannot be simultaneously true. So, if $\mathit{Spec}_{\mathcal{T}}(\phi) \setminus S$ is empty, $\mathit{Spec}_{\mathcal{T}}(\phi)$ must be finite (and contained in S), and eventually we will find a k such that $k \notin \mathit{Spec}_{\mathcal{T}}(\phi)$ and $\phi \wedge \neq (x_1, \ldots, x_k)$ is not $\mathcal{T}$-satisfiable, when our algorithm must stop. In that case it returns $(0, S_0)$ for $S_0 = \mathit{Spec}_{\mathcal{T}}(\phi) \cap [1, k-1]$, which is a finite subset of S and equals $\mathit{Spec}_{\mathcal{T}}(\phi)$.

So now we assume $\mathit{Spec}_{\mathcal{T}}(\phi) \setminus S \neq \varnothing$: we will then prove $\mathit{Spec}_{\mathcal{T}}(\phi)$ has a finite value not in S, and our algorithm will stop when it finds the first such value, and will output (k, S_0) for $S_0 = \mathit{Spec}_{\mathcal{T}}(\phi) \cap [1, k-1]$ a finite subset of S; second, we will prove that the existence of a finite value k in $\mathit{Spec}_{\mathcal{T}}(\phi)$ that is not in S implies $[k, \aleph_0] \subseteq \mathit{Spec}_{\mathcal{T}}(\phi)$, so the spectrum indeed has the form suggested by the algorithm

$$\mathit{Spec}_{\mathcal{T}}(\phi) = (\mathit{Spec}_{\mathcal{T}}(\phi) \cap [1, k-1]) \cup [k, \aleph_0].$$

Suppose, then, that $\mathit{Spec}_{\mathcal{T}}(\phi) \setminus S \neq \varnothing$ but $(\mathit{Spec}_{\mathcal{T}}(\phi) \setminus S) \cap \mathbb{N} = \varnothing$. We then have that $\mathit{Spec}_{\mathcal{T}}(\phi) \setminus S = \{\aleph_0\}$, and therefore $\phi \wedge P_i$ is $\mathcal{T} \sqcup \mathcal{T}_{\leq}^{\overline{S}}$-satisfiable if and only if $F(i) = \aleph_0$, an undecidable problem. This leads to a contradiction, as $\mathcal{T}_{\leq}^{\overline{S}}$, for $\overline{S} = \mathbb{N}^* \setminus S$, is decidable and has no interpretations with sizes in S, thus being S-decidable (see Theorem 4). Next, suppose $n \in \mathit{Spec}_{\mathcal{T}}(\phi) \setminus S$ is finite and there is an $n < m < \aleph_0$ such that $m \notin \mathit{Spec}_{\mathcal{T}}(\phi)$. Then $\phi \wedge P_i$ is $\mathcal{T} \sqcup \mathcal{T}_{n}^{m}$-satisfiable if and only if $i \notin U$, an undecidable problem. This again leads to a contradiction, as $\mathcal{T}_{n}^{m}$ is decidable and S-decidable (Theorem 4).

□

5.3 Sharpness for S-finality

We now prove that Theorem 3 is sharp, namely: given S, every decidable theory that can be combined with all S-final theories must be S-restricted, and every decidable theory that can be combined with all S-restricted theories must be S-final. Again, we utilize operator G from Section 2.

Theorem 6. *$G(\mathfrak{T}_{S\text{-}f}) \subseteq \mathfrak{T}_{S\text{-}r}$ and $G(\mathfrak{T}_{S\text{-}r}) \subseteq \mathfrak{T}_{S\text{-}f}$.*

Proof. 1. Suppose that $\mathcal{T} \in G(\mathfrak{T}_{\text{S-f}})$ is not S-restricted, that is, there are a quantifier-free formula ϕ, an $n \in Spec_{\mathcal{T}}(\phi) \setminus S$ and a maximum element of $Spec_{\mathcal{T}}(\phi)$; define $m = 1 + \max Spec_{\mathcal{T}}(\phi)$. We state that $\mathcal{T}^{m,n}_{\text{S-r}}$ is decidable and S-final, so $\mathcal{T} \sqcup \mathcal{T}^{m,n}_{\text{S-r}}$ should be decidable; yet $\phi \wedge P_k$ is $\mathcal{T} \sqcup \mathcal{T}^{m,n}_{\text{S-r}}$-satisfiable if and only if $k \in U$, leading to a contradiction.

2. Suppose that $\mathcal{T} \in G(\mathfrak{T}_{\text{S-r}})$: first we prove that, if ϕ is $\mathcal{T}$-satisfiable (otherwise we just return 0), there is a finite number in $Spec_{\mathcal{T}}(\phi) \cap S_{\omega}$, for any quantifier-free formula ϕ. Were that not the case we would get $\phi \wedge P_k$ is $\mathcal{T} \sqcup \mathcal{T}^{S}_{\leq}$ if and only if $F(k) = \aleph_0$, what contradicts the fact that $\mathcal{T}^{S}_{\leq}$ is S-restricted.
Next we show how the least element of $Spec_{\mathcal{T}}(\phi) \cap S_{\omega}$ can be calculated: using that $\mathcal{T}^{S}_{=}$ is decidable and S-restricted, we simply search for the first $k \in S$ such that $\phi \wedge P_k$ is $\mathcal{T} \sqcup \mathcal{T}^{S}_{=}$-satisfiable, which must be finite by what we have seen above. Finally, suppose $n \in Spec_{\mathcal{T}}(\phi) \cap S$ but, for some $n < m < \aleph_0$ in S (which we can always find as S is infinite), $m \notin Spec_{\mathcal{T}}(\phi)$. Then $\phi \wedge P_k$ is $\mathcal{T} \sqcup \mathcal{T}^{m}_{n}$-satisfiable if and only if $k \notin U$, contradicting the fact that $\mathcal{T}^{m}_{n}$ is decidable and S-restricted.

□

6 Conclusion and Future Work

We have presented two infinite families of sharp combination methods resembling shininess. One motivates gentleness as simply one extreme of a range of properties, the other extreme being shininess itself. The other requires very little from the weaker theory to be combined, not unlike Nelson–Oppen and polite theory combination. This is in line with [8], which suggests that new combination methods should also be proven sharp given the new techniques available to researchers in the area.

Directions for future work are many: the Galois connection of theory combination makes finding new combination methods much easier; but perhaps more interesting than finding new methods would be trying to generalize those we have at hand to the many-sorted case, a task currently being carried out.

References

1. 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. doi:10.1007/978-3-642-45221-5_15.
2. Filipe Casal and João Rasga. Many-sorted equivalence of shiny and strongly polite theories. *Journal of Automated Reasoning*, 60(2):221–236, February 2018. doi:10.1007/s10817-017-9411-y.
3. Herbert Bruce Enderton. *A Mathematical Introduction to Logic*. New York, Academic Press, 1972.
4. 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.
5. Dejan Jovanović and Clark Barrett. Polite theories revisited. Technical Report TR2010-922, Depatrment of Computer Science, New York University, January 2010. URL: http://www.cs.stanford.edu/~barrett/pubs/JB10-TR.pdf.
6. David Marker. *Model theory: An introduction*, volume 217 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.
7. Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Trans. Program. Lang. Syst.*, 1(2):245–257, October 1979. doi:10.1145/357073.357079.
8. Benjamin Przybocki, Guilherme V. Toledo, and Yoni Zohar. Characterizing sets of theories that can be disjointly combined. In *53rd ACM SIGPLAN Symposium on Principles of Programming Languages (POPL 2026)*, 2026.
9. 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. URL: https://hal.inria.fr/inria-00000570, doi:10.1007/11559306.
10. 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. URL: ftp://ftp.cs.uiowa.edu/pub/tinelli/papers/TinZar-RR-04.pdf.
11. Guilherme V. Toledo, Benjamin Przybocki, and Yoni Zohar. Being polite is not enough (and other limits of theory combination). In *Automated Deduction – CADE 30: 30th International Conference on Automated Deduction, Stuttgart, Germany, July 28-31, 2025, Proceedings*, page 17–34, Berlin, Heidelberg, 2025. Springer-Verlag. doi:10.1007/978-3-031-99984-0_2.

A Proof of Theorem 4

In what follows, $\mathcal{T}_{Eq}$ is the theory having no axioms, over the empty signature (with only equality).

Theorem 4. *Table 2 is correct: each theory satisfies the properties listed for it.*

Proof. 1. $\mathcal{T}_n^m$ is decidable according to [8]. So we move on to prove that, if $n \notin S$, $\mathcal{T}_n^m$ is S-decidable, and if $m, n \in S$, then $\mathcal{T}_n^m$ is S-restricted. $\mathcal{T}_n^m$ only has interpretations of cardinality n or m, so the S-decidability of $\mathcal{T}_n^m$ will follow from the fact every $\mathcal{T}_n^m$-satisfiable quantifier-free formula ϕ has a $\mathcal{T}_n^m$-interpretation of size m: indeed, it is enough to simply add $m - n$ more elements to the domain of a $\mathcal{T}_n^m$-interpretation of size n that satisfies ϕ. It is obvious that $\mathcal{T}_n^m$ is S-restricted as the spectra of all its formulas are contained in $\{n, m\}$, itself contained in S.

2. $\mathcal{T}_{=}^{P}$ is decidable by [8]. Let ϕ be a conjunction of literals, without loss of generality with at most one positive P-literal, and let ϕ' be the result of removing from ϕ any P-literals. If the positive P-literal in ϕ is P_n, then $Spec_{\mathcal{T}_{=}^{P}}(\phi) = \{n\}$ if $\mathrm{minmod}_{\mathcal{T}_{Eq}}(\phi') \leq n$, and is empty otherwise. If ϕ contains no positive P-literals then $Spec_{\mathcal{T}_{=}^{P}}(\phi) = [\mathrm{minmod}_{\mathcal{T}_{Eq}}(\phi'), \aleph_0]$. In both cases we can decide whether an element m of S is in these sets, so $\mathcal{T}_{=}^{P}$ is S-decidable.

3. Again we find the proof that $\mathcal{T}_{\leq}^{S}$ is decidable in [8]. Now we provide a proof that $\mathcal{T}_{\leq}^{S}$ is $\overline{S}$-decidable and S-restricted. It is obvious this theory is $\overline{S}$-decidable, as it contains no interpretations of cardinality in $\overline{S}$; it is also obvious this theory is S-restricted, as it only contains interpretations of cardinality in S_ω.

4. Let ϕ be a conjunction of literals, with at most one positive P-literal (otherwise the formula is immediately not $\mathcal{T}_{=}^{S}$-satisfiable), and let ϕ' be the result of removing from ϕ all of its P-literals. If the positive P-literal in ϕ is P_n, for $n \notin S$, then ϕ is not $\mathcal{T}_{=}^{S}$-satisfiable; if $n \in S$, then $Spec_{\mathcal{T}_{=}^{S}}(\phi) = \{n\}$ if $\mathrm{minmod}_{\mathcal{T}_{Eq}}(\phi') \leq n$, and is empty otherwise. For the non-trivial part of this statement, when $\mathrm{minmod}_{\mathcal{T}_{Eq}}(\phi') \leq n$, we take a $\mathcal{T}_{Eq}$-interpretation satisfying ϕ' with n elements (what we can do as $\mathcal{T}_{Eq}$ is smooth), make P_n true and all P_m false.
If ϕ contains no positive P-literals then $Spec_{\mathcal{T}_{=}^{S}}(\phi) = [\mathrm{minmod}_{\mathcal{T}_{Eq}}(\phi'), \aleph_0]$. This proves not only that this theory is decidable, but also that it is S-shiny: in the case that the spectrum equals $\{n\}$ we must have $n \in S$.
Suppose now that $k \in Spec_{\mathcal{T}}(\phi)$, so there is a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies ϕ with k elements in its domain; we make it into a $\mathcal{T} \sqcup \mathcal{T}_{=}^{S}$-interpretation by making P_k true and all other P_i false, and of course it satisfies $\phi \wedge P_k$. Reciprocally, suppose $\phi \wedge P_k$ is satisfied by the $\mathcal{T} \sqcup \mathcal{T}_{=}^{S}$-interpretation $\mathcal{A}$: because $\mathcal{A}$ is a $\mathcal{T}_{=}^{S}$-interpretation satisfying P_k it contains k elements; because $\mathcal{A}$ is also a $\mathcal{T}$-interpretation satisfying ϕ, we get $k \in Spec_{\mathcal{T}}(\phi)$.
It is obvious $\mathcal{T}_{=}^{S}$ is S-restricted, as it only has interpretations of cardinality in S.

5. We proceed to prove $\mathcal{T}_{\mathsf{S-r}}^{m,n}$ is decidable and, if $n \notin S$ and $m \geq n$, S-final. Take a conjunction of literals ϕ: we may assume it contains at most one positive P-literal, otherwise it will not be $\mathcal{T}_{\mathsf{S-r}}^{m,n}$-satisfiable. Let ϕ' be the result of taking from ϕ any P-literals, and we state ϕ is $\mathcal{T}_{\mathsf{S-r}}^{m,n}$-satisfiable if and only if ϕ' is $\mathcal{T}_{Eq}$-satisfiable.
   For the non-trivial direction take a $\mathcal{T}_{Eq}$-interpretation $\mathcal{A}$ that satisfies ϕ' with any cardinality in $S_\omega \cap [m_0, \aleph_0]$, for $m_0 = \max\{m, \mathrm{minmod}_{\mathcal{T}_{Eq}}(\phi')\}$ (so the set does not contain n and is not empty), make P_k true only if it occurs positively in ϕ, and the resulting interpretation is a $\mathcal{T}_{\mathsf{S-r}}^{m,n}$-interpretation that satisfies ϕ. Moreover, no element in $S_\omega \cap [1, m_0 - 1]$ can be in the spectrum of ϕ, as either: ϕ' is not satisfiable, if $m_0 = \mathrm{minmod}_{\mathcal{T}_{Eq}}(\phi')$; or simply $S \cap [1, m_0 - 1]$ is empty, if $m_0 = m$.
   Now, let $n \in \mathit{Spec}_{\mathcal{T}}(\phi) \setminus S$ and $m = 1 + \max(\mathit{Spec}_{\mathcal{T}}(\phi) \cup \{n\})$, and suppose that $k \in U$: because $n \in \mathit{Spec}_{\mathcal{T}}(\phi)$ there is a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies ϕ with n elements in its domain; we make P_k true and all other P_i false in $\mathcal{A}$, and the result is a $\mathcal{T} \sqcup \mathcal{T}_{S-r}^{m,n}$-interpretation as $k \in U$, that in addition satisfies $\phi \wedge P_k$. Reciprocally, assume $\mathcal{A}$ is a $\mathcal{T} \sqcup \mathcal{T}_{S-r}^{m,n}$-interpretation that satisfies $\phi \wedge P_k$: because $\mathcal{A}$ is a $\mathcal{T}$-interpretation that satisfies ϕ, the cardinality of its domain is bounded by $\max(\mathit{Spec}_{\mathcal{T}}(\phi))$; since $\mathcal{A}$ is also a $\mathcal{T}_{S-r}^{m,n}$-interpretation this implies the domain of $\mathcal{A}$ has exactly n elements, what in turn implies $k \in U$.

$\square$

B Proof of Theorem 2

Theorem 2. *If $\mathcal{T}_1$ is S-decidable and $\mathcal{T}_2$ is S-shiny for some S, and both $\mathcal{T}_1$ and $\mathcal{T}_2$ are decidable, then $\mathcal{T}_1 \sqcup \mathcal{T}_2$ is decidable.*

Proof. Using Lemma 1, it suffices to show that it is then decidable whether $\mathit{Spec}_{\mathcal{T}_1}(\phi_1) \cap \mathit{Spec}_{\mathcal{T}_2}(\phi_2) \neq \varnothing$, where ϕ_1 and ϕ_2 are conjunctions of literals over the signatures Σ_1 and Σ_2, respectively. We state Algorithm 2 presents an adequate decision procedure. We can compute a pair (n, S_0), where $S_0 \subseteq S$ is finite, such that $\mathit{Spec}_{\mathcal{T}_2}(\phi_2)$ equals $S_0 \cup [n, \aleph_0]$ if $n \neq 0$, and S_0 otherwise, from the fact that $\mathcal{T}_2$ is S-shiny. We can test whether $\phi_1 \wedge \neq(x_1, \ldots, x_n)$ is $\mathcal{T}_1$-satisfiable from the fact that $\mathcal{T}_1$ is decidable. Finally, we can test if an $m \in S_0$ is also in $\mathit{Spec}_{\mathcal{T}_1}(\phi_1)$ since $\mathcal{T}_1$ is S-decidable, and $S_0 \subseteq S$.

Next we prove correctness.

($\Rightarrow$): Suppose $\mathit{Spec}_{\mathcal{T}_1}(\phi_1) \cap \mathit{Spec}_{\mathcal{T}_2}(\phi_2)$ is not empty. If $n \neq 0$ and they share an element in $[n, \aleph_0]$ then line 4 of our algorithm will return true. If not they must share an element in S_0 and the loop starting on line 5 will return true.

($\Leftarrow$): Suppose now $\mathit{Spec}_{\mathcal{T}_1}(\phi_1) \cap \mathit{Spec}_{\mathcal{T}_2}(\phi_2)$ is empty. The conditional in line 3 won't be satisfied, as otherwise the intersection would not be empty; and neither will be any conditional in the loop starting at line 5, so eventually, we will return false.

$\square$

C Proof of Proposition 5

Proposition 5.

1. *If $S \subsetneq T$, then $\mathfrak{T}_{T\text{-}d} \subsetneq \mathfrak{T}_{S\text{-}d}$ and $\mathfrak{T}_{S\text{-}s} \subsetneq \mathfrak{T}_{T\text{-}s}$.*
2. *If S and T are incomparable (i.e. $S \nsubseteq T$ and $T \nsubseteq S$), so are $\mathfrak{T}_{S\text{-}d}$ and $\mathfrak{T}_{T\text{-}d}$, and also $\mathfrak{T}_{S\text{-}d}$ and $\mathfrak{T}_{T\text{-}d}$.*

Proof. Suppose $\mathcal{T}$ is T-decidable: this means there is an algorithm that takes a quantifier-free ϕ and an n and, if $n \in T$, it returns whether $n \in Spec_{\mathcal{T}}(\phi)$; since $S \subseteq T$, this same algorithm shows that $\mathcal{T}$ is S-decidable. For S and T-shininess it is enough to notice that if S_0 is a finite subset of S, and $S \subseteq T$, then S_0 is also a finite subset of T.

Let $n_s \in S \setminus T$, and $n_t \in T \setminus S$. Define the $\Sigma_P^{\mathbb{N}}$-theory $\mathcal{T}_s$ with axiomatization

$$\{\psi_{=n_s} \vee \psi_{\geq k} : k \in \mathbb{N}^*\} \cup \{\psi_{=n_s} \to \neg P_k : k \notin U\} \cup \{P_i \to \neg P_j : i \neq j\};$$

we define $\mathcal{T}_t$ analogously. $\mathcal{T}_s$ is decidable: given a quantifier-free ϕ, we may assume it contains at most one positive P-literal; if ϕ' is obtained from ϕ by removing any P-literals, then ϕ is $\mathcal{T}_s$-satisfiable if and only if ϕ' is $\mathcal{T}_{Eq}$-satisfiable. One direction of this is trivial; for the other we can use $\mathcal{T}_{Eq}$ is smooth to find an infinite interpretation that satisfies ϕ', and we make it into a $\mathcal{T}_s$-interpretation that satisfies ϕ by making P_k true if and only if it occurs positively in ϕ. $\mathcal{T}_s$ is also T-decidable, as it contains no interpretations with cardinality in T; yet it is not S-decidable, as $n_s \in Spec_{\mathcal{T}_s}(P_k)$ if and only if $k \in U$. The results for $\mathcal{T}_t$ are completely analogous.

Now consider the theories $\mathcal{T}_{=n_s}$ and $\mathcal{T}_{=n_t}$, axiomatized by respectively $\{\psi_{=n_s}\}$ and $\{\psi_{=n_t}\}$, which are decidable ([8]). $\mathcal{T}_{=n_s}$ is S-shiny, since $Spec_{\mathcal{T}_{n_s}}(\phi) = \{n_s\}$ if $\text{minmod}_{\mathcal{T}_{Eq}}(\phi) \leq n_s$, and otherwise it is empty: for the non-trivial direction, if $\text{minmod}_{\mathcal{T}_{Eq}}(\phi) \leq n_s$ we can use the fact $\mathcal{T}_{Eq}$ is smooth to get a $\mathcal{T}_{Eq}$-interpretation that satisfies ϕ with n_s elements, and of course that will also be a $\mathcal{T}_{=n_s}$-interpretation. But $\mathcal{T}_{=n_s}$ is obviously not T-shiny: the spectrum of a tautology such as $x = x$ in $\mathcal{T}_{=n_s}$ is $\{n_s\}$, and this set is not of the form $T_0 \cup [n, \aleph_0]$, for any n and $\mathcal{T}_0 \subseteq T$, since it does not contain $\aleph_0$; it is also not of the form T_0, for a finite $T_0 \subseteq T$, as $n_s \notin T$. Analogous results hold for $\mathcal{T}_{=n_t}$.

If $S \subsetneq T$ and $n_t \in T \setminus S$, the theories $\mathcal{T}_t$ and $\mathcal{T}_{=n_t}$ above also show $\mathfrak{T}_{\text{T-d}} \subsetneq \mathfrak{T}_{\text{S-d}}$ and $\mathfrak{T}_{\text{S-s}} \subsetneq \mathfrak{T}_{\text{T-s}}$. □

D Proof of Theorem 3

Theorem 3. *Let $\mathcal{T}_1$ be S-restricted, and $\mathcal{T}_2$ be S-final. Then $\mathcal{T}_1 \sqcup \mathcal{T}_2$ is decidable.*

Proof. Let ϕ_1 and ϕ_2 be conjunctions of literals over the signatures Σ_1 and Σ_2, respectively. We state Algorithm 3 presents an adequate decision procedure. That we can compute an $n \in \mathbb{N}^*$ such that $Spec_{\mathcal{T}_2}(\phi_2) \cap S_\omega = [n, \aleph_0] \cap S_\omega$ follows from the fact that $\mathcal{T}_2$ is S-final. We can compute if $\phi_1 \wedge \neq (x_1, \ldots, x_n)$ is $\mathcal{T}_1$-satisfiable since $\mathcal{T}_1$ is decidable. Next we prove correctness.

($\Rightarrow$): Suppose $Spec_{\mathcal{T}_1}(\phi_1) \cap Spec_{\mathcal{T}_2}(\phi_2)$ is not empty. If $\aleph_0$ is an element they share then $\phi_1 \wedge \neq (x_1, \dots, x_n)$ is $\mathcal{T}_1$-satisfiable, so we will get true. If $\aleph_0$ is not in the intersection of the spectra, then $\aleph_0 \notin Spec_{\mathcal{T}_1}(\phi_1)$ and so $Spec_{\mathcal{T}_1}(\phi_1) \subseteq S$, and the fact that $\phi_1 \wedge \neq (x_1, \dots, x_n)$ is $\mathcal{T}_1$-satisfiable implies $Spec_{\mathcal{T}_1}(\phi_1) \cap [n, \aleph_0] \subseteq S_\omega \cap [n, \aleph_0]$ is not empty, thus again we get true.

($\Leftarrow$) Suppose this time $Spec_{\mathcal{T}_1}(\phi_1) \cap Spec_{\mathcal{T}_2}(\phi_2)$ is empty. Then we have that $\phi_1 \wedge \neq (x_1, \dots, x_n)$ will not be $\mathcal{T}_1$-satisfiable, since otherwise the intersection would not be empty. As $\aleph_0 \notin Spec_{\mathcal{T}_1}(\phi_1)$ we get $Spec_{\mathcal{T}_1}(\phi_1) \subseteq S$, and since $Spec_{\mathcal{T}_2}(\phi_2) \cap S_\omega = [n, \aleph_0] \cap S_\omega$ we get false.

□

E Proof of Proposition 8

Proposition 8. *1. If $S \subsetneq T$, $\mathfrak{T}_{S\text{-}r} \subsetneq \mathfrak{T}_{T\text{-}r}$, and $\mathfrak{T}_{T\text{-}f} \subsetneq \mathfrak{T}_{T\text{-}f}$.*
2. If S and T are incomparable, so are $\mathfrak{T}_{S\text{-}r}$ and $\mathfrak{T}_{T\text{-}r}$, and $\mathfrak{T}_{S\text{-}f}$ and $\mathfrak{T}_{T\text{-}f}$.

Proof. Suppose $\mathcal{T}$ is S-restricted, so if $Spec_{\mathcal{T}}(\phi) \setminus S \neq \varnothing$ then $\aleph_0 \in Spec_{\mathcal{T}}(\phi)$ (for quantifier-free ϕ). If $Spec_{\mathcal{T}}(\phi) \setminus T \neq \varnothing$, as $S \subseteq T$ we obtain $Spec_{\mathcal{T}}(\phi) \setminus S \neq \varnothing$ and thus $\aleph_0 \in Spec_{\mathcal{T}}(\phi)$, meaning $\mathcal{T}$ is T-restricted.

Now suppose $\mathcal{T}$ is T-final: there is an algorithm that takes a (quantifier-free) ϕ and returns 0 if it is not $\mathcal{T}$-satisfiable, and otherwise an $n \in \mathbb{N}^*$ such that $Spec_{\mathcal{T}}(\phi) \cap T_\omega = [n, \aleph_0] \cap T_\omega$. In this last case, since, again, $S \subseteq T$, $T \cap S = S$, and so $Spec_{\mathcal{T}}(\phi) \cap S_\omega = (Spec_{\mathcal{T}}(\phi) \cap S_\omega) \cap T_\omega = ([n, \aleph_0] \cap T_\omega) \cap S_\omega = [n, \aleph_0] \cap S_\omega$, we get the same algorithm shows $\mathcal{T}$ is S-final.

Take $n_s \in S \setminus T$, and $n_t \in T \setminus S$. The theory $\mathcal{T}_{=n_s}$ is decidable and S-restricted, but not T-restricted, and analogous statements can be made about $\mathcal{T}_{=n_t}$.

Define the $\Sigma_P^{\mathbb{N}}$-theory $\mathcal{T}_s$ axiomatized by

$$\{\psi_{=n_t} \to \neg P_k : k \notin U\} \cup \{P_i \to \neg P_j : i \neq j\}.$$

It is decidable: given a conjunction of literals $\phi$with at most one positive P-literal P_k, where ϕ' is obtained from ϕ by removing all P-literals, we state ϕ is $\mathcal{T}_s$-satisfiable if and only if ϕ' is $\mathcal{T}_{Eq}$-satisfiable; furthermore, its spectrum equals $[\mathrm{minmod}_{\mathcal{T}_s}(\phi'), \aleph_0]$ if $k \in U$, and $[\mathrm{minmod}_{\mathcal{T}_s}(\phi'), \aleph_0] \setminus \{n_t\}$ otherwise. Indeed, given a $\mathcal{T}_{Eq}$-interpretation satisfying ϕ' with any cardinality greater than or equal to $\mathrm{minmod}_{\mathcal{T}_{Eq}}(\phi')$, we simply make P_n true if and only if $n = k$; this construction only fails to work when $k \notin U$ and the desired cardinality is n_t.

Because n_t is not in S, $Spec_{\mathcal{T}_s}(\phi) \cap S_\omega = [\mathrm{minmod}_{\mathcal{T}_{Eq}}(\phi'), \aleph_0] \cap S_\omega$ (assuming ϕ' is $\mathcal{T}_{Eq}$-satisfiable, otherwise we just return 0), and thus $\mathcal{T}_s$ is S-final. But either n_t is not the first element of T, and $Spec_{\mathcal{T}_s}(P_k) \cap T_\omega = T_\omega \setminus \{n_t\}$ for a $k \notin U$ (which is not a set of the form $[n, \aleph_0] \cap T_\omega$); or it is, and then $Spec_{\mathcal{T}_s}(P_k) \cap T_\omega = [n_t, \aleph_0] \cap T_\omega$ if $k \in U$, and $[n, \aleph_0] \cap T_\omega$ otherwise, for any $n > n_t$, meaning no algorithm could output the n_t, respectively n, given k. $\mathcal{T}_s$ is therefore not T-final, a similar construction giving a T-final theory which is not S-final.

Finally, if $S \subsetneq T$ and $n_t \in T \setminus S$, we have that the theories $\mathcal{T}_{=n_t}$ and $\mathcal{T}_t$ also show $\mathfrak{T}_{\text{S-r}} \subsetneq \mathfrak{T}_{\text{T-r}}$ and $\mathfrak{T}_{\text{T-f}} \subsetneq \mathfrak{T}_{\text{T-f}}$. □