



Maslov’s class K with Equivalence^{*}

Oskar Fiuk

University of Wrocław
contact.oskarfiuk@gmail.com

Abstract

Maslov’s class K is a decidable fragment of equality-free first-order logic. It subsumes numerous classical decidable fragments, including the equality-free monadic fragment, the equality-free two-variable fragment, and the Gödel class. Introduced over 50 years ago, the class K has been studied in automated deduction, primarily through resolution-based methods. Despite this long history, the precise complexity of the satisfiability problem for K remained open until recently, when NExpTime-completeness was finally established.

The two main contributions of this work are as follows.

- The recent proofs of NExpTime-completeness and the finite model property for K are highly non-trivial and technically sophisticated. We present considerably simpler proofs of both results. Our approach combines a novel reduction from K to the class of solvable Skolem sentences with a randomised construction of finite models.
- We introduce the class K+E by allowing sentences in K to contain one distinguished binary predicate E whose interpretation is constrained to be an equivalence relation. By leveraging our approach to this extended class, we prove that K+E has the finite model property and that its satisfiability problem is 2-NExpTime-complete.

1 Introduction

Hilbert’s *Entscheidungsproblem* asked for a general algorithm to determine the satisfiability of first-order statements. Church and Turing famously showed that no such algorithm exists by establishing the undecidability of first-order logic. This negative result gave rise to a broad research programme aimed at identifying fragments of first-order logic for which satisfiability remains decidable. Over the decades, this effort has led to the discovery of numerous decidable formalisms, including the Bernays–Schönfinkel class, the monadic fragment, the Ackermann class, the Gödel class, the two-variable fragment, and the guarded fragment, among others.

One motivation for studying decidable fragments comes from automated deduction, a branch of computer science that seeks effective methods for solving the satisfiability problem, not in full generality, but for broad and practically relevant classes of instances. One of the pioneers of automated deduction was S. J. Maslov, who introduced the class K [8, 26]. In the 1960s, Maslov developed a proof procedure, called the *inverse method* [24], for deciding validity and applied it to the class K. Subsequent work shifted attention to the dual class, denoted $\bar{K}$, and

^{*}Supported by Polish National Science Center, grant No. 2021/41/B/ST6/00996.

to its satisfiability problem. In the 1980s, Zamov proposed a resolution procedure for $\bar{K}$; see Chapter 6 of the monograph *Resolution Methods for the Decision Problem* [5]. In 1999, Hustadt and Schmidt [17] continued Zamov’s work and developed a refined variant of his procedure.

The importance of Maslov’s class $\bar{K}$ stems from the fact that it is one of the broadest fragments of (equality-free) first-order logic. In particular, $\bar{K}$ subsumes many classical fragments (all considered without equality), including: the monadic fragment [25], the Ackermann class [1] and its generalised variant due to Voigt [33], the Gödel class [10], and the solvable Skolem class from Dreben and Goldfarb’s monograph [3] (there denoted as Class 2.4). Moreover, the even broader class of conjunctions of $\bar{K}$ sentences, denoted $\overline{DK}$, is decidable [17]. The absence of equality in $\bar{K}$ is unavoidable: already the Gödel class with equality is undecidable [11].

The class $\bar{K}$ also subsumes the equality-free two-variable fragment, FO^2 , in the sense that sentences in Scott normal form belong to $\bar{K}$ [32]. As a consequence, many standard modal and description logics can be translated into $\bar{K}$. This connection is discussed in the survey [18], where $\bar{K}$ is considered alongside other well-known fragments such as the guarded fragment and the fluted fragment, all of which are incomparable in expressive power with $\bar{K}$.

After the early 2000s, $\bar{K}$ has received relatively little attention. A renewed interest in $\bar{K}$ is motivated by the study of the uniform one-dimensional fragment, denoted UF_1 . Introduced by Hella and Kuusisto in 2014 [16], see also [19, 23], UF_1 was designed as a canonical generalisation of FO^2 to predicates of arity greater than 2. The satisfiability problem for UF_1 (with equality) is NEXPTIME-complete [19], matching the complexity of FO^2 . To increase the expressive power of (equality-free) UF_1 , Kieroński [6] proposed a generalisation denoted AUF_1^- . He observed that AUF_1^- , when put into a suitable normal form, falls within $\overline{DK}$. In [7], an even broader extension of (equality-free) UF_1 , denoted $\forall\text{-}UF_1$, is introduced. It, too, is subsumed by $\overline{DK}$.

Despite this broad applicability, the class $\bar{K}$ remained poorly understood from a theoretical perspective, as the two central questions for logical fragments—the finite model property and the precise complexity of the satisfiability problem—remained open for over 50 years.¹

These questions were resolved only recently: in 2024, Fiuk, Kieroński, and Michielini [7] established that $\bar{K}$ has the finite model property and that the satisfiability problem for $\bar{K}$ is NEXPTIME-complete. In light of this result, the significance of $\bar{K}$ becomes even more apparent: it can express many non-trivial properties, not expressible in other fragments, while being no harder than FO^2 , UF_1 , or the Gödel class. This contrasts with the guarded fragment, which is 2-EXPTIME-complete [14], and with the fluted fragment, which is TOWER-complete [30].

Our contribution. The proof of the finite model property and the NEXPTIME upper bound for $\bar{K}$ given in [7] is technically involved and difficult to follow; we discuss it in more detail in Section 5. Our first contribution is a new and substantially simpler proof of these results.

Our approach is based on a novel Skolemisation procedure that transforms sentences in $\bar{K}$ into sentences in the solvable Skolem class [3, 13]. Notably, the precise complexity of the solvable Skolem class was itself open prior to [7]. We then apply Goldfarb’s probabilistic construction of finite models for sentences in the solvable Skolem class [13], which generalises the earlier approach of Gurevich and Shelah for the Gödel class [15]. A direct application of this randomised model construction, however, does not yield the desired model-size bound. To obtain a tight singly exponential bound, we refine the construction using ideas from [7], which, thanks to our Skolemisation procedure, can be applied in a considerably simpler setting.

Our second contribution extends the decidability result beyond the class $\bar{K}$. A common way of increasing the expressive power of decidable fragments is to restrict the class of admissible

¹In particular, it appears that [5, 17, 4] refrain from making explicit claims regarding the complexity of $\bar{K}$, or even the running time of their decision procedures. Nevertheless, elementary upper bounds can be derived from Theorem 6.3 in [5] and Theorem 6 in [17]; however, these bounds are at best doubly or triply exponential.

structures. For instance, FO^2 has been studied over structures equipped with linear orders [27, 31], equivalence relations [21, 20, 22], and transitive relations [9, 28]. We follow this line of research and prove that $\bar{K}$ remains decidable when extended with one equivalence relation. More precisely, we introduce the class $\bar{K} + E$, obtained by adding to $\bar{K}$ one distinguished binary predicate E whose interpretation is required to be an equivalence relation. We show that $\bar{K} + E$ has the finite model property and that its satisfiability problem is 2-NEXPTIME-complete.

Organisation. Section 2 introduces the notation and basic definitions. Section 3 presents the class $\bar{K}$ and related fragments. Section 4 develops the Skolemisation procedure. Section 5 gives a simplified proof of the NEXPTIME upper bound for $\bar{K}$. Section 6 establishes that $\bar{K}$ extended with one distinguished equivalence relation is 2-NEXPTIME-complete.

2 Preliminaries

Notation. Let $\mathbb{N} = \{0, 1, 2, \dots\}$. For $k \in \mathbb{N}$, the notation $[k]$ stands for the set $\{1, \dots, k\}$, with the convention that $[0] = \emptyset$. More generally, we use interval notation $[a, b] \subseteq \mathbb{N}$ to denote the set $\{a, a+1, \dots, b\}$ for $a \leq b$, and the empty set $\emptyset$ for $a > b$. The powerset of a set S is denoted by 2^S . Tuples are denoted by bars and, by abuse of notation, sometimes regarded as sets of their elements. Concatenation of tuples is denoted by juxtaposition.

First-order logic. We assume general familiarity with first-order logic (FO). The logical symbols are $=, \perp, \top, \vee, \wedge, \neg, \rightarrow, \leftrightarrow$, and the quantifiers $\forall, \exists$. The non-logical symbols are relation symbols (from a countably infinite set), constant symbols (also from a countably infinite set), and variables (again countably many). We do not allow function symbols of positive arity.

A *signature* is a pair $\Sigma = \langle \bar{R}, \bar{c} \rangle$, where $\bar{R}$ is a non-empty tuple of relation symbols and $\bar{c}$ is a tuple of constant symbols. When convenient, we identify Σ with the set $\bar{R} \cup \bar{c}$. Every relation symbol $R \in \bar{R}$ comes with associated arity, denoted $\text{arity}(R)$.

The *length* of a formula φ , denoted $|\varphi|$, is the total number of symbols it contains, where each occurrence of a symbol—be it a variable, relation symbol, or constant—contributes 1. The signature of φ consists of relation and constant symbols that appear in the formula. The set of atomic subformulas of φ is denoted by $\text{atoms}(\varphi)$. The set of *free* variables of φ is denoted by $\text{vars}(\varphi)$; we often view this set as a tuple $\bar{x}$ and indicate this by $\varphi[\bar{x}]$. Given a tuple of elements $\bar{a}$ of the same length as $\bar{x}$, we write $\varphi[\bar{a}]$ for the result of componentwise substitution of the variables in $\bar{x}$ by the corresponding elements of $\bar{a}$. More generally, for any tuple of variables $\bar{y}$, not necessarily coinciding with the free variables of φ , and any tuple of elements $\bar{b}$ of the same length as $\bar{y}$, we write $\varphi[\bar{b}/\bar{y}]$ for the result of componentwise substitution of the variables in $\bar{y}$ by the corresponding elements of $\bar{b}$. We write $\pm\varphi$ to denote either φ or $\neg\varphi$.

We use Fraktur letters such as $\mathfrak{A}, \mathfrak{B}, \dots$ to denote structures, and the corresponding Roman letters $A, B, \dots$ for their domains. Let $\Sigma = \langle \bar{R}, \bar{c} \rangle$ be a signature. A Σ -*structure* $\mathfrak{A}$ interprets every relation symbol $R \in \bar{R}$ as a relation $R^{\mathfrak{A}} \subseteq A^{\text{arity}(R)}$ and every constant symbol $c \in \bar{c}$ as an element $c^{\mathfrak{A}} \in A$. An element $a \in A$ is *unnamed* if $a \neq c^{\mathfrak{A}}$ for every constant $c \in \bar{c}$.

For a sub-signature $\Sigma_0 \subseteq \Sigma$, the Σ_0 -*reduct* of $\mathfrak{A}$ is a Σ_0 -structure $\mathfrak{A}_0$ obtained from $\mathfrak{A}$ by forgetting the interpretations of the symbols in $\Sigma \setminus \Sigma_0$. Equivalently, $\mathfrak{A}$ is a Σ -*expansion* of $\mathfrak{A}_0$.

Types and outer-types. Let $\Sigma = \langle \bar{R}, \bar{c} \rangle$ be a signature. For $k \in \mathbb{N}$, let $\text{lit}^k(\Sigma)$ denote the set of all literals over Σ and the variables $x_1, \dots, x_k$; literals involving $=$ or $\neq$ are not included.

For $k \in \mathbb{N}$, a k -*type* over Σ is a maximal consistent subset of $\text{lit}^k(\Sigma)$; that is, for every pair of dual literals $\lambda, \neg\lambda \in \text{lit}^k(\Sigma)$, exactly one of λ and $\neg\lambda$ belongs to the k -type. A 0-type contains literals involving only constants. A $(k+1)$ -type σ *extends* a k -type τ if $\tau \subseteq \sigma$. A *type* is simply a

k -type for some $k \geq 0$. When convenient, we identify types with the conjunctions of the literals they contain. Let $\tau^k(\Sigma)$ denote the set of all k -types over Σ , and write $\tau(\Sigma) := \bigcup_{k \in \mathbb{N}} \tau^k(\Sigma)$.

Let $\mathfrak{A}$ be a Σ -structure, and let $\bar{a} = \langle a_1, \dots, a_k \rangle$ be a k -tuple of pairwise distinct unnamed elements of $\mathfrak{A}$. We write $\text{tp}^{\mathfrak{A}}[a_1, \dots, a_k] := \{\lambda \in \text{lit}^k(\Sigma) \mid \mathfrak{A} \models \lambda[\bar{a}/\bar{x}]\}$ for the k -type realised by $\bar{a}$ in $\mathfrak{A}$, where $[\bar{a}/\bar{x}]$ denotes the substitution $x_i \mapsto a_i$.

In our proofs, we also need a refined variant of types. For $k \in \mathbb{N}$, let

$$\text{lit}_o^k(\Sigma) := \{\lambda \in \text{lit}^k(\Sigma) \mid \text{vars}(\lambda) = \{x_1, \dots, x_k\} \text{ or } |\text{vars}(\lambda)| \leq 1\}.$$

A k -outer-type over Σ is a maximal consistent subset of $\text{lit}_o^k(\Sigma)$. Thus, a k -outer-type is akin to a k -type, but contains only literals involving either all variables $x_1, \dots, x_k$ or at most one of them. In particular, k -types and k -outer-types coincide for $k \leq 2$.

We use for outer-types the terminology and conventions introduced above for types. In particular, an *outer-type* is a k -outer-type for some $k \geq 0$, and we identify outer-types with the conjunctions of their literals when convenient. We write $\tau_o^k(\Sigma)$ and $\tau_o(\Sigma)$ for the sets of all k -outer-types and all outer-types over Σ , respectively.

For a k -tuple $\bar{a} = \langle a_1, \dots, a_k \rangle$ of pairwise distinct unnamed elements of $\mathfrak{A}$, the k -outer-type realised by $\bar{a}$ in $\mathfrak{A}$ is $\text{tp}_o^{\mathfrak{A}}[a_1, \dots, a_k] := \text{tp}^{\mathfrak{A}}[a_1, \dots, a_k] \cap \text{lit}_o^k(\Sigma)$.

For a k -outer-type τ_o , we write $\tau_o \langle x_i \rangle$ for the 1-type induced on the variable x_i . It is obtained by taking all literals $\lambda \in \tau_o$ such that $\text{vars}(\lambda) \subseteq \{x_i\}$ and renaming x_i with x_1 . Intuitively, if a tuple $\bar{a}$ realises τ_o in a structure, then $\tau_o \langle x_i \rangle$ is the 1-type of the i th element of $\bar{a}$.

For a k -type τ and pairwise distinct indices $i_1, \dots, i_\ell \in [k]$, we write $\tau \langle x_{i_1}, \dots, x_{i_\ell} \rangle$ for the ℓ -outer-type induced by τ on $\langle x_{i_1}, \dots, x_{i_\ell} \rangle$. It is obtained by taking all literals $\lambda \in \tau$ such that either $\text{vars}(\lambda) = \{x_{i_1}, \dots, x_{i_\ell}\}$ or $\text{vars}(\lambda) \subseteq \{x_{i_j}\}$ for some $j \in [\ell]$, and then renaming the variables according to the mapping $x_{i_j} \mapsto x_j$ for every $j \in [\ell]$. Intuitively, if a tuple $\langle a_1, \dots, a_k \rangle$ realises τ in a structure, then $\tau \langle x_{i_1}, \dots, x_{i_\ell} \rangle$ is the ℓ -outer-type of the subtuple $\langle a_{i_1}, \dots, a_{i_\ell} \rangle$.

For the set of all outer-types induced by the k -type τ , we write

$$\tau_o(\tau) := \{\tau \langle x_{i_1}, \dots, x_{i_\ell} \rangle \mid \ell \in [0, k] \text{ and } i_1, \dots, i_\ell \in [k] \text{ are pairwise distinct}\}.$$

3 The Class $\bar{K}$ and Related Fragments

Definition. Let φ be an equality-free sentence in negation normal form. The sentence φ may contain constant symbols. Without loss of generality, we assume that distinct quantifier occurrences in φ bind distinct variables.

For an atom $\alpha \in \text{atoms}(\varphi)$, the φ -prefix of α is the sequence of quantifiers in φ that bind the variables in $\text{vars}(\alpha)$. The order of quantifiers in the φ -prefix matches the order of their occurrences in φ from left to right. In contrast, the order and multiplicity of variables within α , as well as the constants occurring in α , are irrelevant for determining its φ -prefix. In particular, if α contains only constant symbols, then its φ -prefix is the empty sequence.

The sentence φ belongs to the class $\bar{K}$ if there exists a selection $\forall x_1, \dots, \forall x_n$ of universal quantifier occurrences in φ , for some $n \geq 0$, such that:

1. None of the selected universal quantifiers is in the scope of an existential quantifier.
2. For every atom $\alpha \in \text{atoms}(\varphi)$, at least one of the following holds:
 - (i) the φ -prefix of α has length at most 1;
 - (ii) the φ -prefix of α consists exactly of the selected universal quantifiers $\forall x_1, \dots, \forall x_n$;
 - (iii) the last quantifier in the φ -prefix of α is existential.

The selected universal quantifiers are listed in their left-to-right order in φ , but need not form a contiguous block in the sequence of all quantifiers of φ . The variables $x_1, \dots, x_n$ bound by the

selected quantifiers are called the *uniform variables* of φ , and n is called the *uniformity* of φ .

To avoid ambiguity, we adopt the convention that n is always chosen as small as possible. Since for $n = 1$ condition (ii) is subsumed by condition (i), we have $n = 0$ or $n \geq 2$. Under this convention, the required selection $\forall x_1, \dots, \forall x_n$, whenever it exists, is unique. Moreover, if $n \geq 2$, it must coincide with the φ -prefix of some atom of φ .

The above definition of $\bar{K}$ does not essentially rely on negation normal form; it only requires that every quantifier occur positively. We therefore adopt a more liberal syntax for $\bar{K}$ and allow arbitrary Boolean combinations within quantifier-free subformulas. In particular, we admit $\rightarrow$ and $\leftrightarrow$, provided that no quantifier occurs in the premise of an implication or on either side of a bi-implication.

Prenex form. It is often convenient to consider sentences of $\bar{K}$ in prenex form. Given a sentence $\varphi \in \bar{K}$, its standard conversion to prenex form yields, up to a reordering of the initial universal quantifiers, a sentence of the form

$$\forall x_1 \dots \forall x_n Q_{n+1}x_{n+1} \dots Q_mx_m \psi, \quad (1)$$

where $0 \leq n \leq m$, each $Q_i \in \{\forall, \exists\}$, and ψ is equality-free and quantifier-free.

The conditions on φ -prefixes translate as follows: for every atom $\alpha \in \text{atoms}(\psi)$, at least one of the following holds:

- (i) $|\text{vars}(\alpha)| \leq 1$;
- (ii) $\text{vars}(\alpha) = \{x_1, \dots, x_n\}$, where $x_1, \dots, x_n$ are the uniform variables of φ ;
- (iii) the variable $x_i \in \text{vars}(\alpha)$ with maximal index i is existentially quantified.

Example. To illustrate the definition of $\bar{K}$, consider the sentence φ_1 :

$$\forall a, j, s \left((\text{article}(a) \wedge \text{journal}(j) \wedge \text{scientist}(s) \wedge \text{submittedToBy}(a, j, s)) \rightarrow \right. \\ \left. \exists r (\text{reviewer}(r) \wedge \text{reviewsFor}(r, j) \wedge \neg \text{hasConflict}(r, s) \wedge \text{qualifiedFor}(r, a)) \right).$$

The sentence φ_1 belongs to $\bar{K}$ with uniform variables a, j, s . The φ_1 -prefixes of its atoms satisfy the conditions of the definition as follows:

- (i) the prefixes $\forall a, \forall j, \forall s$, and $\exists r$ have length at most 1;
- (ii) the prefix $\forall a \forall j \forall s$ consists of all selected universal quantifiers;
- (iii) the remaining prefixes $\forall j \exists r, \forall s \exists r$, and $\forall a \exists r$ end with an existential quantifier.

As nonexample, consider the transitivity axiom

$$\varphi_2 := \forall x, y, z ((T(x, y) \wedge T(y, z)) \rightarrow T(x, z)).$$

Its atoms have the φ_2 -prefixes $\forall x \forall y, \forall y \forall z$, and $\forall x \forall z$. They have length greater than 1, do not end with an existential quantifier, and cannot all satisfy condition (ii). Hence, $\varphi_2 \notin \bar{K}$.

The class $\overline{DK}$. Following [17], let $\overline{DK}$ denote the class of conjunctions of $\bar{K}$ sentences; that is,

$$\varphi_1 \wedge \dots \wedge \varphi_k \in \overline{DK} \quad \text{whenever} \quad \varphi_1, \dots, \varphi_k \in \bar{K}.$$

A sentence of $\overline{DK}$ need not be equivalent to a sentence of $\bar{K}$: conjuncts of different uniformities cannot, in general, share a common sequence of uniform variables. If they have the same uniformity, however, they can be combined into a single $\bar{K}$ sentence. To see this, consider $\varphi_1, \varphi_2 \in \bar{K}$ of the same uniformity. Without loss of generality, both are in prenex form (1), with the same sequence $\bar{x}$ of uniform variables. Writing $\varphi_i = \forall \bar{x} \varphi'_i[\bar{x}]$, for $i \in \{1, 2\}$, we have

$$\varphi_1 \wedge \varphi_2 \iff \forall \bar{x} (\varphi'_1[\bar{x}] \wedge \varphi'_2[\bar{x}]).$$

A routine check shows that the sentence on the right belongs to $\bar{K}$.

The solvable Skolem class. In their monograph [3], Dreben and Goldfarb systematically study solvable classes of formulas, including several classes of Skolem sentences ($\forall^* \exists^*$). One of

these, referred to as Class 2.4 and later also studied in [13], is the only such class relevant to this paper. We therefore refer to it simply as the *solvable Skolem class*.

Sentences in the solvable Skolem class have the form

$$\forall x_1 \dots \forall x_n \forall x_{n+1} \dots \forall x_\ell \exists x_{\ell+1} \dots \exists x_m \psi, \quad (2)$$

where $0 \leq n \leq \ell \leq m$, ψ is equality-free and quantifier-free, and every $\alpha \in \text{atoms}(\psi)$ satisfies at least one of the following:

$$|\text{vars}(\alpha)| \leq 1 \quad \text{or} \quad \text{vars}(\alpha) = \{x_1, \dots, x_n\} \quad \text{or} \quad x_i \in \text{vars}(\alpha) \text{ for some } i > \ell.$$

The solvable Skolem class is a subfragment of $\overline{\text{K}}$: its sentences are in prenex form (1), with $x_1, \dots, x_n$ as uniform variables, subject to the additional restriction that all existential quantifiers follow all universal quantifiers. In contrast, $\overline{\text{K}}$ allows arbitrary quantifier alternation after the initial block binding the uniform variables.

The solvable Skolem class will play an important role in our decidability proof, as we show that every sentence of $\overline{\text{K}}$ can be reduced to this class while preserving satisfiability.

4 A Skolemisation Procedure

In this section, we introduce the *Skolem normal form* for sentences in $\overline{\text{DK}}$, which allows for a substantially simpler treatment of $\overline{\text{DK}}$ in subsequent proofs. We then show a polynomial-time reduction from the general case to Skolem normal form that preserves (finite) satisfiability.

Definition 1. We say that a sentence $\varphi \in \overline{\text{DK}}$ is in Skolem normal form if

$$\varphi = \bigwedge_{t \in \mathcal{I}} \forall \bar{x} \exists x_{m+1} \psi_t[\bar{x}, x_{m+1}],$$

where $\mathcal{I}$ is a finite set of indices, $\bar{x} = \langle x_1, \dots, x_m \rangle$ is a variable sequence of length $m \geq 1$, common to all t , x_{m+1} is a single variable, and $\psi_t[\bar{x}, x_{m+1}]$ is quantifier-free.

For every $t \in \mathcal{I}$, let $\varphi_t := \forall \bar{x} \exists x_{m+1} \psi_t$ and let $n_t \in \{0\} \cup [2, m]$ denote the uniformity of φ_t , with uniform variables $x_1, \dots, x_{n_t}$.

We partition the set of atoms of ψ_t as $\text{atoms}(\psi_t) = \alpha_t \cup \beta_t \cup \gamma_t$, where

$$\begin{aligned} \alpha_t &:= \{\alpha \in \text{atoms}(\psi_t) \mid \text{vars}(\alpha) \subseteq \{x_i\} \text{ for some } i \in [m]\}, \\ \beta_t &:= \{\alpha \in \text{atoms}(\psi_t) \mid n_t \geq 2 \text{ and } \text{vars}(\alpha) = \{x_1, \dots, x_{n_t}\}\}, \\ \gamma_t &:= \{\alpha \in \text{atoms}(\psi_t) \mid x_{m+1} \in \text{vars}(\alpha)\}. \end{aligned}$$

Let $\overline{\text{DK}}_{\text{snf}}$ denote the class of all Skolem normal form sentences of $\overline{\text{DK}}$. Whenever we consider a sentence of $\overline{\text{DK}}_{\text{snf}}$, we use the notation of Definition 1.

We now establish the required reduction from $\overline{\text{DK}}$ to $\overline{\text{DK}}_{\text{snf}}$. Since the reduction can be viewed as a form of *relational Skolemisation*, we refer to it as the *Skolemisation procedure*.

The usual Skolem form can subsequently be obtained by replacing the existential variable x_{m+1} in each conjunct φ_t by a fresh term $f_t(x_1, \dots, x_m)$. We will not need this form.

Proposition 2. There exists a polynomial-time function $\mathcal{F}: \overline{\text{DK}} \rightarrow \overline{\text{DK}}_{\text{snf}}$ such that, for every non-empty domain D , $\varphi \in \overline{\text{DK}}$ has a model over D if and only if $\mathcal{F}(\varphi)$ has a model over D .

We show the rewriting for a single $\overline{\text{K}}$ sentence; the extension to $\overline{\text{DK}}$ is straightforward.

Let $\varphi \in \overline{\text{K}}$ and let $\Sigma = \langle \bar{R}, \bar{c} \rangle$ be the signature of φ . By converting φ to the prenex form (1) and introducing dummy quantifiers whenever necessary, we may assume that

$$\varphi = \forall \bar{x} \forall x_{n+1} \exists x_{n+2} \dots \forall x_{m-1} \exists x_m \psi,$$

where $\bar{x} = \langle x_1, \dots, x_n \rangle$ are the uniform variables, $m - n > 0$ is even, the quantifier binding x_{n+i} is universal for odd i and existential for even i , and ψ is quantifier-free.

We proceed by induction on $m - n$. If $m - n = 2$, then $\varphi = \forall x_1 \dots \forall x_{n+1} \exists x_{n+2} \psi$ is already in $\overline{\text{DK}}_{\text{snf}}$, and we set $\mathcal{F}(\varphi) := \varphi$. Suppose now that $m - n > 2$. Let $\bar{x}' = \langle x_1, \dots, x_{n+2} \rangle$ and write $\chi[\bar{x}'] := \forall x_{n+3} \exists x_{n+4} \dots \forall x_{m-1} \exists x_m \psi$. Thus, $\varphi = \forall \bar{x} \forall x_{n+1} \exists x_{n+2} \chi$.

Define $\delta := \{\alpha \in \text{atoms}(\psi) \mid \text{vars}(\alpha) \subseteq \{x_1, \dots, x_{n+2}\}\}$. We extend Σ to a signature Σ^+ by introducing a fresh $(n+2)$ -ary relation symbol R_χ and, for every $\alpha \in \delta$, a fresh $(n+2)$ -ary relation symbol R_α . Let χ^+ be obtained from χ by replacing every occurrence of each atom $\alpha \in \delta$ by $R_\alpha(\bar{x}')$. The new atoms use precisely the variables $\bar{x}' = \langle x_1, \dots, x_{n+2} \rangle$, in this order, while the subscript α records the original relation symbol and its full sequence of arguments, including repetitions and constants. Over Σ^+ , define

$$\mu := \forall \bar{x} \forall x_{n+1} \exists x_{n+2} \left(R_\chi(\bar{x}') \wedge \bigwedge_{\alpha \in \delta} (\alpha \leftrightarrow R_\alpha(\bar{x}')) \right) \quad \text{and} \quad \nu := \forall \bar{x}' (R_\chi(\bar{x}') \rightarrow \chi^+).$$

The above sentences provide an intermediate rewriting step in the construction of $\mathcal{F}(\varphi)$. We return to $\mathcal{F}(\varphi)$ shortly and first verify the correctness of this step.

We first show that $\mu \wedge \nu \models \varphi$. By the definition of ν , we may replace the conjunct $R_\chi(\bar{x}')$ in the matrix of μ by χ^+ when deriving a consequence of $\mu \wedge \nu$. It follows that

$$\mu \wedge \nu \models \forall \bar{x} \forall x_{n+1} \exists x_{n+2} \left(\chi^+ \wedge \bigwedge_{\alpha \in \delta} (\alpha \leftrightarrow R_\alpha(\bar{x}')) \right).$$

For every $\alpha \in \delta$, the equivalence $\alpha \leftrightarrow R_\alpha(\bar{x}')$ allows each occurrence of $R_\alpha(\bar{x}')$ in χ^+ to be replaced by the corresponding atom α . By the definition of χ^+ , performing these replacements reconstructs χ . Since $\varphi = \forall \bar{x} \forall x_{n+1} \exists x_{n+2} \chi$, it follows that $\mu \wedge \nu \models \varphi$.

Conversely, suppose that $\mathfrak{A} \models \varphi$. We expand $\mathfrak{A}$ to a Σ^+ -structure by interpreting

$$R_\chi^\mathfrak{A} := \{\bar{a} \in A^{n+2} \mid \mathfrak{A} \models \chi[\bar{a}]\}, \quad R_\alpha^\mathfrak{A} := \{\bar{a} \in A^{n+2} \mid \mathfrak{A} \models \alpha[\bar{a}/\bar{x}']\} \quad \text{for } \alpha \in \delta.$$

One verifies that the resulting expansion satisfies $\mathfrak{A} \models \mu \wedge \nu$, thereby establishing the correctness of the rewriting step. We now return to the construction of $\mathcal{F}(\varphi)$.

The sentence μ belongs to $\overline{\text{DK}}_{\text{snf}}$. The sentence ν need not belong to $\overline{\text{DK}}_{\text{snf}}$, but it is again a $\overline{\text{K}}$ sentence, now with uniform variables $\bar{x}'$. Moreover, ν has strictly smaller quantifier alternation depth than φ , so the induction hypothesis applies to it. We set $\mathcal{F}(\varphi) := \mu \wedge \mathcal{F}(\nu)$.

Finally, we introduce dummy universal quantifiers as necessary to ensure that all conjuncts have the same number of them, as required by Definition 1. This establishes the reduction.

A simple decidability argument. Observe that sentences in $\overline{\text{DK}}_{\text{snf}}$ are, in fact, conjunctions of sentences in the solvable Skolem class. Combining this observation with Proposition 2 yields a simple decidability argument for $\overline{\text{DK}}$, based on the earlier work of Goldfarb [13].

Using a probabilistic method, Goldfarb establishes the finite model property and decidability of the solvable Skolem class. Although his result is technically stated for a single sentence in the solvable Skolem class, whereas $\overline{\text{DK}}_{\text{snf}}$ consists of conjunctions of such sentences, his argument extends seamlessly to this broader setting. In this way, we obtain that $\overline{\text{DK}}$ has the finite model property and that its satisfiability problem is decidable.

5 The Complexity of $\overline{\text{DK}}$: A Simpler Proof

In this section, we show the optimal NEXPTIME upper bound on the complexity of $\overline{\text{K}}$ and $\overline{\text{DK}}$. As this result is not new, we begin with a brief overview of the existing literature.

- Dreben and Goldfarb [3] showed that the solvable Skolem class (referred to as Class 2.4 therein) has the finite model property, with a doubly exponential bound on the size of constructed models. Later, Goldfarb [13] provided an alternative proof by extending the

probabilistic method of Gurevich and Shelah [15] for the Gödel class.² However, he proved only that, as the domain size tends to infinity, the probability that a random structure satisfies a given sentence converges to 1. A closer inspection of the argument yields at best a doubly exponential upper bound on the required domain size. Moreover, none of [3, 13] addresses the complexity of satisfiability, and only a naïve 2-NEXPTIME upper bound can be inferred. The optimal NEXPTIME upper bound for the solvable Skolem class follows instead from the corresponding upper bound for $\bar{K}$ and is stated as a corollary in [7].

It also appears that standard proofs of the NEXPTIME upper bound for the Gödel class (see, e.g., [2, 29]) do not extend to the solvable Skolem class. In particular, these proofs exploit the fact that Gödel-class sentences admit only two universal quantifiers.

- A recent work [7] establishes that $\bar{K}$ has the exponential model property and is therefore NEXPTIME-complete. However, the approach is technically involved. In the absence of a suitable Skolemisation procedure, the authors work directly with $\bar{K}$ sentences featuring quantifier alternation, handled via a game-theoretic approach (the so-called *satisfiability game* in [7]). Satisfiability is characterised in terms of the existence of a winning strategy in this game. To obtain exponential bounds, equivalence relations on game positions are introduced to yield a compact representation of strategies. From such representations, finite models are constructed using technically intricate arguments: construction is based on so-called paradoxical tournament graphs with vertices labelled by game positions.

Our proof avoids the technical complications of [7] and yields the optimal NEXPTIME upper bound via a substantially simpler argument. Our plan is as follows.

In Subsection 5.1, we begin with two preliminary reductions that simplify the subsequent arguments. In Subsection 5.2, we provide a criterion for satisfiability of sentences in $\bar{DK}_{\text{snf}}$ in terms of the existence of a combinatorial certificate, called a *satisfiability witness*, and show that this condition is both necessary and sufficient. For the sufficiency direction, we construct finite random models by refining Goldfarb's approach with a more detailed analysis of probabilities.

A purely random construction is not sufficient, as the size of a satisfiability witness is, in general, doubly exponential. Therefore, in Subsection 5.3, we show that every satisfiable sentence admits a witness of exponential size. This part of the proof is inspired by [7]; however, thanks to our Skolemisation procedure, ideas of [7] are applied in a more direct setting.

5.1 Preliminary reductions

Standard names. In a structure $\mathfrak{A}$, distinct constants may be interpreted by the same element, $c_1^{\mathfrak{A}} = c_2^{\mathfrak{A}}$, which is inconvenient in the subsequent proofs. Hence, from this point on we adopt the *standard name assumption*: distinct constants denote distinct elements, and we identify each constant with its interpretation, $c^{\mathfrak{A}} = c$. This assumption is without loss of generality for both satisfiability and model-size bounds. Accordingly, we partition the domain as $A = A_0 \dot{\cup} \bar{c}$, where A_0 is the set of *unnamed* elements and $\bar{c}$ is the tuple of constants in the signature.

Simple assignments. To avoid case distinctions, decidability proofs for the Gödel class (see, e.g., [15]) typically begin with a simple reduction that allows one to consider only pairs of distinct elements. Namely, an input sentence $\forall x_1, x_2 \exists \bar{y} \psi[x_1, x_2, \bar{y}]$, with ψ quantifier-free, is transformed into $\forall x_1, x_2 \exists \bar{y} \exists \bar{z} (x_1 \neq x_2 \rightarrow (\psi[x_1, x_2, \bar{y}] \wedge \psi[x_1, x_1, \bar{z}]))$, where $\bar{z}$ is a tuple of fresh variables of the same length as $\bar{y}$. Over domains of size at least two, this transformation preserves satisfiability. Its advantage is that the subsequent argument need only consider the

²Goldfarb [12] also developed a probabilistic proof for a different fragment, called the *Maslov–Krom class*. This fragment is unrelated to the class K and should not be confused with it.

case $x_1 \neq x_2$, since the case $x_1 = x_2$ is accounted for by the second conjunct $\psi[x_1, x_1, \bar{z}]$. Equality is used only as part of the reduction and is not allowed to occur in ψ .

We generalise this idea to $\overline{\text{DK}}_{\text{snf}}$ by introducing *simple* assignments, which map distinct variables to pairwise distinct unnamed elements. The reduction follows the same general scheme as for the Gödel class: it considers all possible *equality patterns* among the universal variables and constants. Consequently, the subsequent arguments become significantly cleaner, as it suffices to consider only simple assignments.

However, unlike in the Gödel-class case, carrying out the reduction is not immediate. Indeed, with an unbounded number of universal variables, the number of equality patterns is exponential in the sentence length. We address this issue by carefully controlling the relevant syntactic parameters and expressing the resulting model-size bounds in terms of the length of the input sentence. We formalise this convention as follows.

Given a structure $\mathfrak{A}$ and $m \in \mathbb{N}$, we write $\langle A \rangle^m$ for the set of all m -tuples $\bar{a} \in A^m$ whose entries are pairwise distinct and unnamed. The following lemma establishes a reduction that allows satisfaction of sentences in $\overline{\text{DK}}_{\text{snf}}$ to be verified by considering only tuples from $\langle A \rangle^m$.

Lemma 3. *Every sentence $\varphi_{\text{in}} \in \overline{\text{DK}}_{\text{snf}}$ can be reduced to a sentence $\varphi \in \overline{\text{DK}}_{\text{snf}}$ over the same signature such that, with the notation of Definition 1 applied to φ , the following hold:*

1. *The sentences φ_{in} and φ are logically equivalent.*
2. *For every structure $\mathfrak{A}$ with at least m unnamed elements,*

$$\mathfrak{A} \models \varphi \iff \text{for every } t \in \mathcal{I} \text{ and } \bar{a} \in \langle A \rangle^m, \text{ there exists } b \in A \text{ such that } \mathfrak{A} \models \psi_t[\bar{a}, b].$$

Moreover, the reduction does not introduce new variables and the following bounds hold:

$$|\mathcal{I}| = 2^{\mathcal{O}(|\varphi_{\text{in}}| \log |\varphi_{\text{in}}|)}, \quad \max_{t \in \mathcal{I}} |\psi_t| = \mathcal{O}(|\varphi_{\text{in}}|), \quad \left| \bigcup_{t \in \mathcal{I}: n_t \geq 2} \alpha_t \right| = \mathcal{O}(|\varphi_{\text{in}}|^2).$$

5.2 A Satisfiability Criterion

For the rest of this section, we fix an *input sentence* $\varphi_{\text{in}} \in \overline{\text{DK}}_{\text{snf}}$ and let φ be the sentence obtained from φ_{in} by Lemma 3. Let $\Sigma = \langle \bar{R}, \bar{c} \rangle$ denote their common signature. Unless stated otherwise, we consider only Σ -structures with at least m unnamed elements. From this point on, the notation of Definition 1 always refers to the components of φ , while φ_{in} is retained only to record the size of the input. This distinction is important because the reduction incurs an exponential blow-up, whereas the model-size bounds established below must be stated in terms of $|\varphi_{\text{in}}|$ rather than $|\varphi|$. For this purpose, Lemma 3 also records precise bounds on the relevant syntactic parameters of φ . The significance of these bounds will become clear in Subsection 5.3.

Satisfiability witnesses. Following [13], we characterise satisfiability of φ in terms of the existence of a collection of outer-types satisfying suitable realisability and extension properties.

We first characterise, in combinatorial terms, when a collection of outer-types corresponds to the set of all k -outer-types for $k \leq m+1$ realised in a structure in which every realised 1-type has at least $m+1$ distinct realisations. This is captured by the notion of φ -realisability.

Definition 4. *A set of outer-types $\mathcal{S} \subseteq \tau_o(\Sigma)$ is φ -realisable if the following conditions hold:*

- *$\mathcal{S}$ contains a unique 0-type, and every other outer-type $\tau_o \in \mathcal{S}$ is consistent with it.*
- *$\mathcal{S}$ contains at least one 1-type.*
- *$\mathcal{S}$ is closed under automorphisms of outer-types induced by permutations of variables.*
- *For every $k \in [m+1]$ and every k -outer-type $\tau_o \in \mathcal{S}$, $\tau_o \langle x_i \rangle \in \mathcal{S}$ for every $i \in [k]$.*
- *For every $k \in [m+1]$ and all 1-types $\pi_1, \dots, \pi_k \in \mathcal{S}$, there exists a k -outer-type $\tau_o \in \mathcal{S}$ such that $\tau_o \langle x_i \rangle = \pi_i$ for every $i \in [k]$.*

Fix for a moment a φ -realisable set $\mathcal{S}$ of outer-types. To construct a model of φ using $\mathcal{S}$, it suffices, for each $t \in \mathcal{I}$, to ensure that every m -type compatible with $\mathcal{S}$ can be extended to an $(m+1)$ -type satisfying ψ_t and remaining compatible with $\mathcal{S}$. Intuitively, the m -type describes the facts under an assignment to the universal variables $x_1, \dots, x_m$, while the $(m+1)$ -type describes the facts under an assignment extended to the existential variable x_{m+1} .

We formalise this in the following definition.

Definition 5. Let $\mathcal{S} \subseteq \boldsymbol{\tau}_o(\Sigma)$. An m -type τ is an $\mathcal{S}$ -requirement if $\boldsymbol{\tau}_o(\tau) \subseteq \mathcal{S}$. Given $t \in \mathcal{I}$, an $(\mathcal{S}, t)$ -extension of τ is an $(m+1)$ -type σ such that $\boldsymbol{\tau}_o(\sigma) \subseteq \mathcal{S}$ and $\sigma \models \tau \wedge \psi_t$.

The set $\mathcal{S}$ has the φ -extension property if, for every $\mathcal{S}$ -requirement τ and every $t \in \mathcal{I}$, there exists an $(\mathcal{S}, t)$ -extension of τ .

Thus, an $\mathcal{S}$ -requirement τ is any m -type that induces only outer-types from $\mathcal{S}$, and its $(\mathcal{S}, t)$ -extension σ is any $(m+1)$ -type that extends τ in such a way that ψ_t is satisfied and all induced outer-types still belong to $\mathcal{S}$. Finally, the φ -extension property guarantees that every $\mathcal{S}$ -requirement can be extended in this way for each conjunct φ_t .

We combine the notions of φ -realisability and the φ -extension property in the following definition of a satisfiability witness. As we show subsequently, φ is satisfiable if and only if there exists a collection of outer-types with these two properties.

Definition 6. A set of outer-types $\mathcal{W} \subseteq \boldsymbol{\tau}_o(\Sigma)$ is called a satisfiability witness for φ if it is φ -realisable and has the φ -extension property.

Completeness. The notion of a satisfiability witness is *complete*: whenever φ is satisfiable, it admits a satisfiability witness. This is established by the following argument.

We begin with an auxiliary claim describing a type-surgery operation that allows us to transfer witnessing extensions between m -types while controlling the induced outer-types.

Claim 7. Fix $t \in \mathcal{I}$, and let τ_1, τ_2 be m -types such that

1. $\tau_1 \models \alpha \iff \tau_2 \models \alpha$ for every $\alpha \in \boldsymbol{\alpha}_t \cup \boldsymbol{\beta}_t$;
2. $\tau_1 \langle x_i \rangle = \tau_2 \langle x_i \rangle$ for every $i \in [m]$.

Let σ_1 be an $(m+1)$ -type extending τ_1 and satisfying ψ_t , and define $\sigma_2 := (\sigma_1 \setminus \tau_1) \cup \tau_2$. Then σ_2 is an $(m+1)$ -type extending τ_2 and satisfying ψ_t . Moreover, $\boldsymbol{\tau}_o(\sigma_2) \subseteq \boldsymbol{\tau}_o(\tau_2) \cup \boldsymbol{\tau}_o(\sigma_1)$.

We now return to the completeness argument. Let $\mathfrak{A} \models \varphi$. We first transform $\mathfrak{A}$ into a new structure $\mathfrak{B}$ by taking its *infinite unfolding*, constructed as follows. The domain B consists of countably many copies of each element of A and is identified with $A \times \mathbb{N}$, where $(c, 0)$ is identified with the constant $c \in \bar{c}$. For every relation symbol $R \in \bar{R}$ of arity k , we define

$$R^{\mathfrak{B}} := \{ \langle (a_1, i_1), \dots, (a_k, i_k) \rangle \mid a_1, \dots, a_k \in A, i_1, \dots, i_k \in \mathbb{N}, R^{\mathfrak{A}}(a_1, \dots, a_k) \}.$$

By a standard argument, $\mathfrak{B} \models \varphi$. Moreover, for every $t \in \mathcal{I}$ and every $\bar{a} \in \langle B \rangle^m$, there exists an unnamed element $b_{\bar{a}, t} \in B$ not occurring in $\bar{a}$ such that $\mathfrak{B} \models \psi_t[\bar{a}, b_{\bar{a}, t}]$. Indeed, every element $(a, i) \in B$, including one identified with a constant, is indistinguishable from its copy $(a, j) \in B$. Hence, any witness for ψ_t can be replaced by an unnamed copy not occurring in $\bar{a}$.

Let $\mathcal{W}$ consist of all k -outer-types, for $k \leq m+1$, realised in $\mathfrak{B}$. We claim that $\mathcal{W}$ is a satisfiability witness for φ .

By construction of $\mathfrak{B}$, every realised 1-type has infinitely many distinct realisations. It directly follows that $\mathcal{W}$ is φ -realisable. It remains to verify the φ -extension property.

Fix $t \in \mathcal{I}$ and let τ be a $\mathcal{W}$ -requirement. Consider the atoms in $\boldsymbol{\alpha}_t \cup \boldsymbol{\beta}_t$, that is, those that contain only universal variables. The atoms in $\boldsymbol{\alpha}_t$ depend only on the 1-types induced by τ , while those in $\boldsymbol{\beta}_t$ depend only on the outer-type induced on the uniform variables $x_1, \dots, x_{n_t}$.

Hence, although τ itself need not be realised in $\mathfrak{B}$, we can choose a tuple $\bar{a} \in \langle B \rangle^m$ that agrees with τ on $\alpha_t \cup \beta_t$. Indeed, we first choose distinct unnamed elements $a_1, \dots, a_{n_t}$ realising $\tau \langle x_1, \dots, x_{n_t} \rangle$. We then choose distinct unnamed elements $a_{n_t+1}, \dots, a_m$, also distinct from the preceding ones, such that $\text{tp}^{\mathfrak{B}}[a_i] = \tau \langle x_i \rangle$ for every $i \in [n_t + 1, m]$.

Let $\tau' := \text{tp}^{\mathfrak{B}}[\bar{a}]$ and $\sigma' := \text{tp}^{\mathfrak{B}}[\bar{a}, \bar{b}_{\bar{a}, t}]$. Then σ' extends τ' and satisfies ψ_t . By the choice of $\bar{a}$, the types τ and τ' agree on all atoms in $\alpha_t \cup \beta_t$ and induce the same sequences of 1-types. Hence, Claim 7, applied to τ and τ' , shows that $\sigma := (\sigma' \setminus \tau') \cup \tau$ extends τ , satisfies ψ_t , and $\tau_o(\sigma) \subseteq \tau_o(\tau) \cup \tau_o(\sigma')$. Since τ is a $\mathcal{W}$ -requirement, we have $\tau_o(\tau) \subseteq \mathcal{W}$. Moreover, σ' is realised in $\mathfrak{B}$, and hence $\tau_o(\sigma') \subseteq \mathcal{W}$. Therefore, $\tau_o(\sigma) \subseteq \mathcal{W}$, so σ is a $(\mathcal{W}, t)$ -extension of τ .

Consequently, $\mathcal{W}$ has the φ -extension property and is therefore a satisfiability witness for φ .

Random models. We now establish the more interesting direction, namely that the notion of a satisfiability witness is *sound* for finite satisfiability: whenever φ admits a satisfiability witness, it has a finite model. The proof follows Goldfarb's probabilistic approach [13], but refines the probability analysis to obtain an explicit bound on the size of the constructed models.

Fix a satisfiability witness $\mathcal{W}$ for φ . Let $p \in \mathbb{N}$ be a parameter, to be determined later, specifying the number of unnamed elements in the domain. We construct a finite Σ -structure $\mathfrak{A}$ with domain $A := [p] \cup \bar{c}$ in three steps:

1. For every $a \in [p]$, choose uniformly at random a 1-type $\pi \in \mathcal{W} \cap \tau^1(\Sigma)$, and set $\text{tp}^{\mathfrak{A}}[a] := \pi$.
2. For every $k \in [2, m+1]$ and every $\{a_1 < \dots < a_k\} \subseteq [p]$, choose uniformly at random a k -outer-type $\tau_o \in \mathcal{W} \cap \tau_o^k(\Sigma)$ such that $\tau_o \langle x_i \rangle = \text{tp}^{\mathfrak{A}}[a_i]$ for every $i \in [k]$. Since $\mathcal{W}$ is φ -realisable, at least one such k -outer-type exists. Since τ_o is consistent with the 1-types, it can be assigned to $\langle a_1, \dots, a_k \rangle$ without affecting the outer-type assigned to any other tuple with a different set of underlying elements. Thus, set $\text{tp}_o^{\mathfrak{A}}[a_1, \dots, a_k] := \tau_o$.
3. Set to false the truth values of all facts with more than $m+1$ distinct unnamed elements.

The construction of $\mathfrak{A}$ is now complete. In particular, recall that types and outer-types also include literals involving constants. Thus, all facts involving them are implicitly defined.

We now show that $\mathbb{P}[\mathfrak{A} \models \varphi] > 0$ once the parameter p is sufficiently large. For the proof, let $s_\gamma := \max_{t \in \mathcal{I}} |\gamma_t|$ and $\delta := |\mathcal{W}|^{-(s_\gamma+1)}$.

Claim 8. *Let $t \in \mathcal{I}$ and let $\bar{a} \in \langle A \rangle^m$. Then $\mathbb{P}[\mathfrak{A} \not\models \exists x_{m+1} \psi_t[\bar{a}, x_{m+1}]] \leq (1 - \delta)^{p-m}$.*

Proof. Write $\bar{a} = \langle a_1, \dots, a_m \rangle$ and condition on the type $\tau := \text{tp}^{\mathfrak{A}}[\bar{a}]$. By construction of $\mathfrak{A}$, every outer-type induced by τ belongs to $\mathcal{W}$, and hence τ is a $\mathcal{W}$ -requirement. Since $\mathcal{W}$ is a satisfiability witness, τ admits a $(\mathcal{W}, t)$ -extension σ .

Fix an element $a_{m+1} \in [p] \setminus \bar{a}$. We lower-bound the probability that a_{m+1} is a witness for $\bar{a}$, that is, that $\mathfrak{A} \models \psi_t[\bar{a}, a_{m+1}]$. Since $\sigma \models \psi_t$, it suffices that every atom $\alpha \in \text{atoms}(\psi_t)$ has the same truth value on $\langle a_1, \dots, a_{m+1} \rangle$ as in σ , that is, $\text{tp}^{\mathfrak{A}}[a_1, \dots, a_{m+1}] \models \alpha \iff \sigma \models \alpha$.

Since σ extends τ , the atoms in $\alpha_t \cup \beta_t$ already have the required truth values. It therefore remains to consider the atoms in γ_t . For an atom $\alpha \in \gamma_t$ with $\text{vars}(\alpha) = \{x_{i_1}, \dots, x_{i_k}, x_{m+1}\}$, the required truth value is ensured whenever

$$\text{tp}_o^{\mathfrak{A}}[a_{i_1}, \dots, a_{i_k}, a_{m+1}] = \sigma \langle x_{i_1}, \dots, x_{i_k}, x_{m+1} \rangle. \quad (3)$$

Let $\mathcal{E}(a_{m+1})$ be the event that (3) holds for every $\alpha \in \gamma_t$. For this event to occur, the 1-type of a_{m+1} must agree with the 1-type induced by σ on x_{m+1} , and the outer-type of every relevant tuple $\langle a_{i_1}, \dots, a_{i_k}, a_{m+1} \rangle$ must agree with the corresponding outer-type induced by σ .

Each of the above requirements is an admissible choice in the random construction, and each of them is made uniformly from a set of at most $|\mathcal{W}|$ possibilities. As there are at most $|\gamma_t| + 1 \leq s_\gamma + 1$ such requirements, the product rule for probabilities gives

$$\mathbb{P}[\mathcal{E}(a_{m+1}) \mid \text{tp}^{\mathfrak{A}}[\bar{a}] = \tau] \geq |\mathcal{W}|^{-(s_\gamma+1)} = \delta. \quad (4)$$

Now let $B := [p] \setminus \bar{a}$. The events $\{\mathcal{E}(b) \mid b \in B\}$ are mutually independent, conditional on $\text{tp}^{\mathfrak{A}}[\bar{a}] = \tau$. Indeed, for distinct $b, b' \in B$, the events $\mathcal{E}(b)$ and $\mathcal{E}(b')$ depend on disjoint random choices, except for the 1-types of the elements of $\bar{a}$, which are fixed by the conditioning. Consequently, the probability that no element $b \in B$ provides a witness for $\bar{a}$ is at most

$$\mathbb{P}\left[\bigcap_{b \in B} \neg \mathcal{E}(b) \mid \text{tp}^{\mathfrak{A}}[\bar{a}] = \tau\right] = \prod_{b \in B} \mathbb{P}[\neg \mathcal{E}(b) \mid \text{tp}^{\mathfrak{A}}[\bar{a}] = \tau] \leq (1 - \delta)^{|B|}. \quad (5)$$

Since $|B| = p - m$, we obtain $\mathbb{P}[\mathfrak{A} \not\models \exists x_{m+1} \psi_t[\bar{a}, x_{m+1}] \mid \text{tp}^{\mathfrak{A}}[\bar{a}] = \tau] \leq (1 - \delta)^{p-m}$. As the final bound is independent of τ , the condition $\text{tp}^{\mathfrak{A}}[\bar{a}] = \tau$ may be dropped. $\square$

We now bound the probability that the sampled structure fails to satisfy φ . By the form of φ , we have $\mathfrak{A} \models \varphi$ precisely when $\mathfrak{A} \models \exists x_{m+1} \psi_t[\bar{a}, x_{m+1}]$ for every $t \in \mathcal{I}$ and every tuple $\bar{a} \in \langle A \rangle^m$. For every such pair t and $\bar{a}$, Claim 8 gives $\mathbb{P}[\mathfrak{A} \not\models \exists x_{m+1} \psi_t[\bar{a}, x_{m+1}]] \leq (1 - \delta)^{p-m}$. Taking the union bound over all $t \in \mathcal{I}$ and $\bar{a} \in \langle A \rangle^m$, and using $|\langle A \rangle^m| \leq p^m$, we obtain

$$\mathbb{P}[\mathfrak{A} \not\models \varphi] \leq \sum_{t \in \mathcal{I}} \sum_{\bar{a} \in \langle A \rangle^m} \mathbb{P}[\mathfrak{A} \not\models \exists x_{m+1} \psi_t[\bar{a}, x_{m+1}]] \leq |\mathcal{I}| \cdot p^m \cdot (1 - \delta)^{p-m}. \quad (6)$$

For fixed φ and $\mathcal{W}$, the right-hand side of (6) tends to 0 as $p \rightarrow \infty$, since its polynomial factor p^m is dominated by the exponentially decreasing factor $(1 - \delta)^{p-m}$. Thus, for sufficiently large p , $\mathbb{P}[\mathfrak{A} \models \varphi] > 0$. Consequently, there exists a finite model of φ .

A more careful analysis of (6) gives the following quantitative version of the result, which will be needed for our complexity bounds.

Lemma 9. *Let $\varphi_{\text{in}} \in \overline{\text{DK}}_{\text{snf}}$, and let φ be obtained from φ_{in} by the reduction of Lemma 3. Given a satisfiability witness $\mathcal{W}$ for φ , the sentence φ admits a finite model $\mathfrak{A}$ with domain size*

$$|A| = 2^{\mathcal{O}(|\varphi_{\text{in}}| \log |\varphi_{\text{in}}|)} \cdot |\mathcal{W}|^{\mathcal{O}(|\varphi_{\text{in}}|)}.$$

5.3 Exponential-Size Certificates

We retain the notation from Subsection 5.2: the input sentence $\varphi_{\text{in}} \in \overline{\text{DK}}_{\text{snf}}$, its conversion φ given by Lemma 3, and the signature $\Sigma = \langle \bar{R}, \bar{c} \rangle$. In this subsection, using techniques from [7], we show that whenever φ is satisfiable, it admits a satisfiability witness of size exponential in $|\varphi_{\text{in}}|$. Stating the bound in terms of φ_{in} is important, because φ is exponentially larger than φ_{in} .

Obtaining the desired bound is not immediate. A straightforward estimate on the number of possible k -outer-types gives a doubly exponential bound in the size of the signature Σ . In the presence of constants and relation symbols of unbounded arity, this bound is already attained at $k = 1$. Thus, simply counting all outer-types that may occur in a satisfiability witness yields only a doubly exponential bound. To obtain a singly exponential bound, we act as follows.

We start with an arbitrary satisfiability witness $\mathcal{W}$ for φ and reduce its size in two stages. First, we reduce the number of 1-types occurring in $\mathcal{W}$, obtaining a new witness $\mathcal{W}'$. This step may, in fact, increase the number of k -outer-types for $k > 1$. Second, we take the witness $\mathcal{W}'$ and reduce the number of k -outer-types for $k > 1$ while preserving the number of 1-types. This way we obtain the final witness $\mathcal{W}''$ which is exponentially bounded in $|\varphi_{\text{in}}|$.

Reduction of 1-types. Recall that the notation of Definition 1 refers to φ , and, in particular, that $\text{atoms}(\psi_t) = \alpha_t \cup \beta_t \cup \gamma_t$. For the construction, we define the auxiliary sets

$$\alpha^* := \bigcup_{t \in \mathcal{I}: n_t \geq 2} \alpha_t, \quad \gamma'_t := \{\alpha \in \gamma_t \mid \text{vars}(\alpha) = \{x_{m+1}\}\}. \quad (7)$$

For brevity, we introduce the following notation for the relevant set sizes:

$$s_\alpha^* = |\alpha^*|, \quad s_\alpha = \max_{t \in \mathcal{I}} |\alpha_t|, \quad s_\beta = \max_{t \in \mathcal{I}} |\beta_t|, \quad s'_\gamma = \max_{t \in \mathcal{I}} |\gamma'_t|.$$

We reduce the number of 1-types occurring in $\mathcal{W}$ by introducing a family of equivalence relations. For $t \in \mathcal{I}$ and 1-types $\pi_1, \pi_2 \in \mathcal{W} \cap \boldsymbol{\tau}^1(\Sigma)$, define $\pi_1 \sim_t \pi_2$ if

$$\pi_1 \models \alpha[x_1/\cdot] \iff \pi_2 \models \alpha[x_1/\cdot] \quad \text{for every } \alpha \in \boldsymbol{\alpha}^* \cup \boldsymbol{\gamma}'_t.$$

Above, the notation $[x_1/\cdot]$ stands for the substitution replacing every variable with x_1 .

Let $N_{\sim_t}$ denote the number of equivalence classes of $\sim_t$. By definition,

$$\sum_{t \in \mathcal{I}} N_{\sim_t} \leq |\mathcal{I}| \cdot 2^{s_{\alpha}^* + s_{\gamma}'}. \quad (8)$$

For every $t \in \mathcal{I}$, we fix a choice function $\text{chc}_t: (\mathcal{W} \cap \boldsymbol{\tau}^1(\Sigma))/\sim_t \rightarrow \mathcal{W} \cap \boldsymbol{\tau}^1(\Sigma)$ that selects a representative from each equivalence class of $\sim_t$. For convenience, for a 1-type $\pi \in \mathcal{W}$, we write $\text{chc}(\pi; t) := \text{chc}_t([\pi]_{\sim_t})$ for the chosen $\sim_t$ -representative of π .

More generally, let $\tau_o \in \mathcal{W}$ be a k -outer-type and let $t_1, \dots, t_k \in \mathcal{I}$. For every $i \in [k]$, let $\pi_i := \tau_o \langle x_i \rangle$ and $\pi'_i := \text{chc}(\pi_i; t_i)$. We write $\text{chc}(\tau_o; t_1, \dots, t_k)$ for the k -outer-type obtained from τ_o by replacing each induced 1-type π_i by π'_i , and leaving all other literals unchanged.

Observe that replacing the induced 1-types leaves unchanged every literal involving at least two distinct variables. Moreover, since all 1-types under consideration belong to $\mathcal{W}$, they induce the same 0-type, and hence the variable-free literals remain unchanged as well. Therefore, the outer-types τ_o and $\text{chc}(\tau_o; t_1, \dots, t_k)$ may differ only in literals involving precisely one variable.

We now define $\mathcal{W}' := \{\text{chc}(\tau_o; t_1, \dots, t_k) \mid k \in [0, m+1], \tau_o \in \mathcal{W} \cap \boldsymbol{\tau}_o^k(\Sigma), t_1, \dots, t_k \in \mathcal{I}\}$.

Claim 10. *The set $\mathcal{W}'$ is a satisfiability witness for φ .*

A routine verification of the definition shows that $\mathcal{W}'$ is φ -realisable. In contrast, establishing the φ -extension property requires a rather technical argument. We therefore consider only a special case that illustrates the underlying intuition; the full proof is deferred to the full version.

We consider $\mathcal{W}'$ -requirements that are direct images of $\mathcal{W}$ -requirements and restrict attention to conjuncts of uniformity at least 2. Hence, fix a $\mathcal{W}$ -requirement τ and indices $t_1, \dots, t_m \in \mathcal{I}$, and consider the $\mathcal{W}'$ -requirement $\tau' = \text{chc}(\tau; t_1, \dots, t_m)$, where chc is extended from outer-types to types in the analogous way. Fix $t \in \mathcal{I}$, and suppose that $n_t \geq 2$. Under these simplifying assumptions, we construct a $(\mathcal{W}', t)$ -extension of τ' .

Since $\mathcal{W}$ is a satisfiability witness, τ admits a $(\mathcal{W}, t)$ -extension σ . We claim that $\sigma' := \text{chc}(\sigma; t_1, \dots, t_m, t)$ is a $(\mathcal{W}', t)$ -extension of τ' . By construction, σ' extends τ' and induces only outer-types from $\mathcal{W}'$. Thus, it remains only to verify that $\sigma' \models \psi_t$.

Since $\sigma \models \psi_t$, it suffices to show that every atom of ψ_t has the same truth value in σ' as in σ . The transformation from σ to σ' changes only induced 1-types, so atoms involving at least two distinct variables are unaffected. Hence, it remains only to consider the atoms in $\boldsymbol{\alpha}_t \cup \boldsymbol{\gamma}'_t$.

Let $\alpha \in \boldsymbol{\alpha}_t$, with $\text{vars}(\alpha) \subseteq \{x_i\}$ for some $i \in [m]$. Its truth value can be affected only by the 1-type induced on x_i , which is replaced according to $\sim_{t_i}$. Since $n_t \geq 2$, we have $\boldsymbol{\alpha}_t \subseteq \boldsymbol{\alpha}^*$, and hence the definition of $\sim_{t_i}$ yields $\sigma \models \alpha \iff \sigma' \models \alpha$.

Now let $\alpha \in \boldsymbol{\gamma}'_t$. Its truth value depends only on the 1-type induced on x_{m+1} , which is replaced according to $\sim_t$. Hence, by the definition of $\sim_t$, $\sigma \models \alpha \iff \sigma' \models \alpha$.

Reduction of witness size. We now reduce the size of $\mathcal{W}'$ by introducing another family of equivalence relations, this time on the $\mathcal{W}'$ -requirements.

For $t \in \mathcal{I}$ and $\mathcal{W}'$ -requirements τ_1, τ_2 , we define $\tau_1 \approx_t \tau_2$ if

- $\tau_1 \models \alpha \iff \tau_2 \models \alpha$ for every $\alpha \in \boldsymbol{\alpha}_t \cup \boldsymbol{\beta}_t$;
- $\tau_1 \langle x_i \rangle = \tau_2 \langle x_i \rangle$ for every $i \in [m]$.

Let $N_{\approx_t}$ denote the number of equivalence classes of $\approx_t$. There are $|\mathcal{W}' \cap \boldsymbol{\tau}^1(\Sigma)|^m$ sequences of 1-types of length m and at most $2^{s_{\alpha} + s_{\beta}}$ truth assignments to the atoms in $\boldsymbol{\alpha}_t \cup \boldsymbol{\beta}_t$. Hence,

$$\sum_{t \in \mathcal{I}} N_{\approx_t} \leq |\mathcal{I}| \cdot |\mathcal{W}' \cap \boldsymbol{\tau}^1(\Sigma)|^m \cdot 2^{s_{\alpha} + s_{\beta}}. \quad (9)$$

We construct the final set $\mathcal{W}''$ of outer-types in two stages:

1. For each $t \in \mathcal{I}$ and each equivalence class $\mathcal{C}$ of $\approx_t$, choose a representative $\tau \in \mathcal{C}$ and a $(\mathcal{W}', t)$ -extension σ of τ , and include all outer-types from $\tau_o(\sigma)$.
2. For each $k \in [m+1]$ and each sequence of 1-types $\pi_1, \dots, \pi_k \in \mathcal{W}' \cap \boldsymbol{\tau}^1(\Sigma)$, choose a single k -outer-type $\tau_o \in \mathcal{W}'$ such that $\tau_o \langle x_i \rangle = \pi_i$ for every $i \in [k]$, and include all k -outer-types obtainable from τ_o by permuting variables.

Claim 11. *The set $\mathcal{W}''$ is a satisfiability witness for φ .*

Proof. We first verify the φ -extension property. Let τ be any $\mathcal{W}''$ -requirement and let $t \in \mathcal{I}$. Since $\mathcal{W}'' \subseteq \mathcal{W}'$, the type τ is also a $\mathcal{W}'$ -requirement. Let τ' be the representative of the $\approx_t$ -class of τ selected in stage 1, and let σ' be the selected $(\mathcal{W}', t)$ -extension of τ' . By Claim 7, $\sigma := (\sigma' \setminus \tau') \cup \tau$ extends τ and satisfies ψ_t . Moreover, every outer-type induced by σ is induced either by σ' or by τ . The former belong to $\mathcal{W}''$ by stage 1, while the latter belong to $\mathcal{W}''$ because τ is a $\mathcal{W}''$ -requirement. Hence, σ is in fact a $(\mathcal{W}'', t)$ -extension of τ .

It remains to verify that $\mathcal{W}''$ is φ -realisable. The outer-types included in stage 1 already satisfy all conditions of φ -realisability except, potentially, the condition requiring every sequence of 1-types of length at most $m+1$ to be induced by some outer-type. The outer-types included in stage 2 ensure that this remaining condition is satisfied. $\square$

Size bound. An $(m+1)$ -type can induce at most $\sum_{i=0}^{m+1} \binom{m+1}{i} \cdot i! \leq (m+1)^{m+2}$ distinct outer-types. For each equivalence class $\mathcal{C}$ of $\approx_t$, stage 1 includes all outer-types induced by one selected $(m+1)$ -type extending the representative of $\mathcal{C}$. Moreover, stage 2 includes at most $(m+1)!$ additional outer-types for every sequence of 1-types of length at most $m+1$. Therefore,

$$|\mathcal{W}''| \leq \sum_{t \in \mathcal{I}} N_{\approx_t} \cdot (m+1)^{m+2} + \sum_{k \in [m+1]} |\mathcal{W}' \cap \boldsymbol{\tau}^1(\Sigma)|^k \cdot (m+1)!. \quad (10)$$

Combining inequalities (8), (9), and (10) yields

$$|\mathcal{W}''| \leq |\mathcal{I}|^{m+1} \cdot 2^{m \cdot (s_\alpha^* + s_\gamma') + s_\alpha + s_\beta} \cdot (m+1)^{m+2} + (m+1) \cdot |\mathcal{I}|^{m+1} \cdot 2^{(m+1) \cdot (s_\alpha^* + s_\gamma')} \cdot (m+1)!. \quad (11)$$

By Lemma 3, $s_\alpha^* = \mathcal{O}(|\varphi_{\text{in}}|^2)$, $|\mathcal{I}| = 2^{\mathcal{O}(|\varphi_{\text{in}}| \log |\varphi_{\text{in}}|)}$, and m, s_α, s_β , and s_γ' are all $\mathcal{O}(|\varphi_{\text{in}}|)$. Substituting these bounds into (11) gives $|\mathcal{W}''| = 2^{\mathcal{O}(|\varphi_{\text{in}}|^3)}$.

We have therefore established an exponential bound on the size of minimal satisfiability witnesses. By applying Lemma 9, we obtain that every satisfiable sentence in $\overline{\text{DK}}$ admits a model of exponential size. Hence, $\overline{\text{DK}}$ has the exponential model property and its satisfiability problem is in NEXPTIME. For the lower bound, recall that $\overline{\text{DK}}$ subsumes the Scott normal form of equality-free FO^2 , whose satisfiability problem is NEXPTIME-hard; see, e.g., [29].

Theorem 12. *The class $\overline{\text{DK}}$ has the exponential model property. The satisfiability problem for $\overline{\text{DK}}$ is NEXPTIME-complete.*

6 The Decidability of $\overline{\text{DK}}$ with Equivalence

In this section, we study an extension of $\overline{\text{DK}}$ with a single equivalence relation. Formally, let $\overline{\text{DK}} + E$ denote the class of $\overline{\text{DK}}$ sentences over signatures containing a distinguished binary relation symbol E , called the *equivalence symbol*. In addition to this distinguished symbol E , signatures may contain arbitrary non-distinguished relation symbols and constants. Models are restricted to structures $\mathfrak{A}$ in which $E^{\mathfrak{A}}$ is an equivalence relation. Syntactically, E may be used *freely*: E -atoms need not satisfy the restrictions of $\overline{\text{K}}$ on φ -prefixes.

We establish a 2-NEXPTIME upper bound for satisfiability of $\overline{DK} + E$. Together with the matching lower bound proved in the full version, this yields 2-NEXPTIME-completeness.

The Skolemisation procedure of Proposition 2 remains sound for $\overline{DK} + E$. Hence, we work with $\overline{DK}_{\text{snf}} + E$ rather than the full class $\overline{DK} + E$. The two preliminary reductions from Subsection 5.1 extend to $\overline{DK}_{\text{snf}} + E$ as well. Accordingly, from this point on we adopt the standard name assumption, fix an input sentence $\varphi_{\text{in}} \in \overline{DK}_{\text{snf}} + E$, and let φ be the sentence obtained from φ_{in} by Lemma 3. As before, φ_{in} is retained only for measuring the input size, while all subsequent arguments concern φ . We use for φ the notation introduced in Definition 1.

Let $\Sigma_0 = \langle \bar{R}, \bar{c} \rangle$ be the *base signature* of φ , consisting of the free relation symbols and constants, and let $\Sigma_E = \langle \bar{R} \cup \{E\}, \bar{c} \rangle$ be the full signature of φ , obtained by extending Σ_0 with the equivalence symbol E .

Class labels. Our first step is to encode the equivalence classes of E by unary predicates. For each $j \in \mathbb{N}$, introduce a fresh unary predicate L_j , called a *class label*, and let $L := \{L_j \mid j \in \mathbb{N}\}$. We use the auxiliary signature $\Sigma_L := \langle \bar{R} \cup L, \bar{c} \rangle$, obtained from the base signature Σ_0 by adding the class-label predicates. Note that, by construction, the signature Σ_L is countably infinite.

Given a Σ_E -structure $\mathfrak{A}$ in which $E^{\mathfrak{A}}$ is an equivalence relation, we associate with it a Σ_L -structure by assigning a distinct class label to each equivalence class. More precisely, enumerate the classes in $A/E^{\mathfrak{A}}$ and assign to each $C \in A/E^{\mathfrak{A}}$ one previously unused class label L_j by interpreting $L_j^{\mathfrak{A}} = C$. If $A/E^{\mathfrak{A}}$ is finite, all unused class labels L_j are interpreted as empty.

Thus, now every element carries exactly one class label, and two elements are E -equivalent precisely when they carry the same label; that is, $aE^{\mathfrak{A}}b \iff \mathfrak{A} \models L_j(a) \wedge L_j(b)$ for some $j \in \mathbb{N}$.

We then forget the interpretation $E^{\mathfrak{A}}$ and view $\mathfrak{A}$ as a Σ_L -structure. No information is lost in doing so, since $E^{\mathfrak{A}}$ can be recovered entirely from the class labels using the above equivalence. Conversely, any Σ_L -structure in which every element carries exactly one class label induces in this way an equivalence relation, which we take as the interpretation of E . Accordingly, from this point on we work exclusively with the signature Σ_L and regard E as implicitly encoded. Under this convention, φ can still be evaluated in Σ_L -structures. Importantly, φ has no direct access to the predicates L_j and can only use the derived relation E .

Types and outer-types. We henceforth work with types and outer-types over the auxiliary signature Σ_L . They are defined as before, with the additional requirement that every variable or constant occurring in a type or outer-type satisfies exactly one class-label predicate. Thus, despite the infinite signature Σ_L , individual types and outer-types admit finite descriptions: each is represented by its Σ_0 -reduct together with the class labels of its variables and constants.

As with Σ_L -structures, we adopt the convention that in a type or outer-type the equivalence relation E is induced by the class-label predicates L_j .

The following terminology will be convenient. A class label L_j *occurs* in a type or outer-type τ if τ contains a positive L_j -literal involving a variable or a constant; the corresponding variable or constant then *carries* the class label L_j . A 1-type π *carries* a class label L_j if $\pi \models L_j(x_1)$. More generally, L_j *occurs* in a set of outer-types if it occurs in one of its members.

Satisfiability witnesses. We adapt Definitions 4 and 5, and consequently Definition 6, to the present setting by taking all types and outer-types occurring in those definitions over Σ_L . Apart from this change of signature, the definitions remain unchanged. In particular, satisfiability witnesses may now be infinite.

The notion of a satisfiability witness remains complete, that is, whenever φ is satisfiable, it admits a satisfiability witness. This is established by the argument from Subsection 5.2, which applies unchanged to the class-label encoding.

It remains to establish the converse direction with suitable size bounds. We show that whenever φ admits a satisfiability witness, it admits a finite satisfiability witness of size doubly

exponential in $|\varphi_{\text{in}}|$. We then apply the random model construction of Subsection 5.2 to obtain a finite model of doubly exponential size. This yields the desired 2-NEXPTIME upper bound.

Local reduction of 1-types. Fix a satisfiability witness $\mathcal{W}$ for φ . In $\mathcal{W}$, the number of 1-types carrying any fixed class label L_j may be doubly exponential. Our first step is to reduce this number to singly exponential. This reduction is crucial for obtaining the optimal 2-NEXPTIME upper bound; without it, the resulting complexity bound would be only 3-NEXPTIME.

We adapt the strategy from Subsection 5.3. Recall that there, for every $t \in \mathcal{I}$, we introduced an equivalence relation $\sim_t$ on 1-types. In the present setting, we refine $\sim_t$ by additionally requiring equivalent 1-types to carry the same class label.

Let the sets α^* and γ'_t be as in (7), and let $[x_1/\cdot]$ denote the substitution identifying every variable with x_1 . For $t \in \mathcal{I}$ and 1-types $\pi_1, \pi_2 \in \mathcal{W} \cap \tau^1(\Sigma_L)$, we set $\pi_1 \sim_t \pi_2$ if

1. $\pi_1 \models \alpha[x_1/\cdot] \iff \pi_2 \models \alpha[x_1/\cdot]$ for every $\alpha \in \alpha^* \cup \gamma'_t$;
2. $\pi_1 \models L_j(x_1)$ and $\pi_2 \models L_j(x_1)$ for some $j \in \mathbb{N}$.

We now define $\mathcal{W}'$ exactly as in Subsection 5.3, using the updated definition of $\sim_t$. The argument of Claim 10 (see the full version) extends naturally to the additional predicates L_j and shows that $\mathcal{W}'$ is a satisfiability witness for φ .

It remains to bound the number of 1-types carrying a fixed label L_j . Let $s_\alpha^* := |\alpha^*|$ and $s'_{\gamma} := \max_{t \in \mathcal{I}} |\gamma'_t|$. Fix $j \in \mathbb{N}$. For a given $t \in \mathcal{I}$, there are at most $2^{s_\alpha^* + s'_{\gamma}}$ distinct $\sim_t$ -classes carrying the label L_j . As $\mathcal{W}'$ retains one representative for each $\sim_t$ -class for each t , we obtain

$$|\{\pi \in \mathcal{W}' \cap \tau^1(\Sigma_L) \mid \pi \models L_j(x_1)\}| \leq |\mathcal{I}| \cdot 2^{s_\alpha^* + s'_{\gamma}}. \quad (12)$$

By Lemma 3, the right-hand side of (12) is $2^{\mathcal{O}(|\varphi_{\text{in}}|^2)}$. Therefore, the number of 1-types in $\mathcal{W}'$ carrying any fixed class label is singly exponential in $|\varphi_{\text{in}}|$, as desired. At this stage, however, $\mathcal{W}'$ may still be infinite, since infinitely many distinct class labels may occur.

Profiles of class labels. Our next goal is therefore to transform the satisfiability witness $\mathcal{W}'$ into one with only finitely many class labels. The key observation is that class labels carrying the same collection of 1-types, up to the choice of the class label itself, should be interchangeable. To make this idea precise, we associate with each class label a *profile* recording this collection of 1-types and, additionally, the constants carrying that label. We then close $\mathcal{W}'$ under interchanges of labels having the same profile. This will allow us to discard all but a bounded number of labels of each profile: whenever a discarded class label is needed to provide an existential witness, it can be replaced by one of the retained labels of the same profile.

For every $j \in \mathbb{N}$, we call the pair $P_j := (R_j, C_j)$ the *profile* of the class label L_j , where

$$\begin{aligned} R_j &:= \{\pi \in \tau^1(\Sigma_0) \mid \pi \text{ is the } \Sigma_0\text{-reduct of a 1-type in } \mathcal{W}' \text{ that contains } L_j(x_1)\}, \\ C_j &:= \{c \in \bar{c} \mid \text{the unique 0-type of } \mathcal{W}' \text{ contains } L_j(c)\}. \end{aligned}$$

The set R_j records the underlying 1-types over the base signature Σ_0 that carry the class label L_j , while C_j records the constants carrying that label. Hence, the number of distinct profiles P_j is clearly finite. More precisely, it is at most doubly exponential in $|\varphi_{\text{in}}|$.

Indeed, by (12), the number of distinct sets R_j is bounded by

$$|\{R_j \mid j \in \mathbb{N}\}| \leq |\tau^1(\Sigma_0)|^{2^{\mathcal{O}(|\varphi_{\text{in}}|^2)}} \leq \left(2^{2^{\mathcal{O}(|\varphi_{\text{in}}| \log |\varphi_{\text{in}}|)}}\right)^{2^{\mathcal{O}(|\varphi_{\text{in}}|^2)}} \leq 2^{2^{\mathcal{O}(|\varphi_{\text{in}}|^2)}}. \quad (13)$$

The second component C_j contributes only linearly many additional possibilities: if $C_j \neq \emptyset$, then no other class label has the same profile as L_j .

A highly symmetric witness. We now close $\mathcal{W}'$ under permutations of class labels having the same profile. We say that a bijection $\rho: \mathbb{N} \rightarrow \mathbb{N}$ is *profile-preserving* if $P_j = P_{\rho(j)}$ for every $j \in \mathbb{N}$. Such a bijection acts on types and outer-types by renaming class labels: every literal

$\pm L_j(a)$ is replaced by the corresponding literal $\pm L_{\rho(j)}(a)$. We then set

$$\mathcal{W}'' := \{\rho(\tau) \mid \tau \in \mathcal{W}' \text{ and } \rho \text{ is a profile-preserving bijection}\}.$$

We claim that the set $\mathcal{W}''$ is a satisfiability witness. It is clearly φ -realisable. To verify the φ -extension property, fix $t \in \mathcal{I}$ and a $\mathcal{W}''$ -requirement τ . Let $\tau_o := \tau\langle x_1, \dots, x_{n_t} \rangle$, and choose a profile-preserving bijection ρ such that $\tau_o \in \rho(\mathcal{W}')$. Moreover, let $\pi_i := \tau\langle x_i \rangle$ for every $i \in [n_t + 1, m]$. Since ρ is profile-preserving, each 1-type π_i satisfies $\pi_i \in \rho(\mathcal{W}')$.

We now construct a $\mathcal{W}'$ -requirement τ' in two steps:

1. Set $\tau'\langle x_1, \dots, x_{n_t} \rangle := \rho^{-1}(\tau_o)$, and, for every $i \in [n_t + 1, m]$, set $\tau'\langle x_i \rangle := \rho^{-1}(\pi_i)$.
2. Using φ -realisability of $\mathcal{W}'$, assign outer-types from $\mathcal{W}'$ to the remaining sets of variables, consistently with the already fixed literals.

Let σ' be a $(\mathcal{W}', t)$ -extension of τ' , and set $\sigma := (\rho(\sigma') \setminus \rho(\tau')) \cup \tau$. By step 1, the types τ and $\rho(\tau')$ induce the same 1-types and agree on all atoms in $\alpha_t \cup \beta_t$. Hence, Claim 7, which extends unchanged to the class-label encoding, yields that σ extends τ , satisfies ψ_t , and $\tau_o(\sigma) \subseteq \tau_o(\tau) \cup \tau_o(\rho(\sigma'))$. Since τ is a $\mathcal{W}''$ -requirement, $\tau_o(\tau) \subseteq \mathcal{W}''$. Moreover, $\tau_o(\sigma') \subseteq \mathcal{W}'$, and therefore $\tau_o(\rho(\sigma')) \subseteq \rho(\mathcal{W}') \subseteq \mathcal{W}''$. Thus $\tau_o(\sigma) \subseteq \mathcal{W}''$, so σ is a $(\mathcal{W}'', t)$ -extension of τ .

A finite witness. We now retain only finitely many class labels of each profile. For every profile P , let $\mathcal{J}_P := \{j \in \mathbb{N} \mid P_j = P\}$, and choose $\mathcal{K}_P \subseteq \mathcal{J}_P$ such that $|\mathcal{K}_P| = \min\{|\mathcal{J}_P|, m+1\}$.

Write $\mathcal{K} := \bigcup_P \mathcal{K}_P$. Let $\mathcal{W}'''$ consist of the outer-types in $\mathcal{W}''$ in which every variable carries a class label L_j with $j \in \mathcal{K}$. Formally, for every $k \in [0, m+1]$ and every k -outer-type τ_o ,

$$\tau_o \in \mathcal{W}''' \iff \tau_o \in \mathcal{W}'' \text{ and } \tau_o \models L_{j_1}(x_1) \wedge \dots \wedge L_{j_k}(x_k) \text{ for some } j_1, \dots, j_k \in \mathcal{K}.$$

Note that class labels carried by constants are retained: if $C_j \neq \emptyset$, then $\mathcal{J}_{P_j} = \{j\}$.

We claim that $\mathcal{W}'''$ is again a satisfiability witness for φ . It is straightforward to verify that $\mathcal{W}'''$ is φ -realisable. To verify the φ -extension property, fix $t \in \mathcal{I}$ and a $\mathcal{W}'''$ -requirement τ . Since $\mathcal{W}''' \subseteq \mathcal{W}''$, the m -type τ is also a $\mathcal{W}''$ -requirement, and hence admits a $(\mathcal{W}'', t)$ -extension σ . Let L_j be the class label carried by the variable x_{m+1} in σ , and let $P = P_j$. If $j \in \mathcal{K}_P$, then σ is already a $(\mathcal{W}''', t)$ -extension of τ .

Suppose otherwise. Then $|\mathcal{J}_P| > m+1$ and $|\mathcal{K}_P| = m+1$. Since the variables occurring in τ carry at most m distinct class labels, there exists $k \in \mathcal{K}_P$ that is not carried by any of them. As $P_j = P_k$, the transposition exchanging j and k is profile-preserving. Moreover, it fixes τ , while replacing the label of x_{m+1} by the retained label L_k . Since $\mathcal{W}''$ is closed under profile-preserving bijections, the image of σ under this transposition is a $(\mathcal{W}'', t)$ -extension of τ . Since all labels occurring in τ are retained in $\mathcal{W}'''$, it is in fact a $(\mathcal{W}''', t)$ -extension of τ .

Size bound. We now bound the size of $\mathcal{W}'''$. Each outer-type in $\mathcal{W}'''$ is described by its Σ_0 -reduct and the class labels carried by its variables and constants. For every $k \leq m+1$, the number of all k -outer-types over Σ_0 is $2^{2^{\mathcal{O}(|\varphi_{\text{in}}| \log |\varphi_{\text{in}}|)}}$. Since at most $m+1$ class labels are retained for each profile P_j , the number of class labels occurring in $\mathcal{W}'''$ is at most $m+1$ times the bound in (13), and hence at most $2^{2^{\mathcal{O}(|\varphi_{\text{in}}|^2)}}$. Therefore, $|\mathcal{W}'''|$ is doubly exponential in $|\varphi_{\text{in}}|$.

We now take the reduct of $\mathcal{W}'''$ obtained by forgetting the class-label predicates, which are no longer needed. The resulting collection can be viewed as a satisfiability witness over the ordinary finite signature. The random model construction of Subsection 5.2 therefore applies without change. Hence, Lemma 9 yields a doubly exponential bound on model size.

Theorem 13. *The class $\overline{\text{DK}} + E$ has the doubly exponential model property. The satisfiability problem for $\overline{\text{DK}} + E$ is 2-NEXPTIME-complete.*

Acknowledgment. The author thanks Emanuel Kieroński for discussions, and the reviewers of LPAR 2026 for their comments. Funding: Polish National Science Center 2021/41/B/ST6/00996.

References

- [1] Wilhelm Ackermann. Über die Erfüllbarkeit gewisser Zähl ausdrücke. *Mathematische Annalen*, 100:638–649, 1928.
- [2] Egon Börger, Erich Grädel, and Yuri Gurevich. *The classical decision problem*. Perspectives in Mathematical Logic. Springer, 1997.
- [3] Burton Dreben and Warren D. Goldfarb. *The Decision Problem: Solvable Classes of Quantificational Formulas*. Addison-Wesley, 1979.
- [4] Christian G. Fermüller, Alexander Leitsch, Ullrich Hustadt, and Tanel Tammet. Resolution decision procedures. In John Alan Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning (in 2 volumes)*, pages 1791–1849. Elsevier and MIT Press, 2001.
- [5] Christian G. Fermüller, Alexander Leitsch, Tanel Tammet, and N. K. Zamov. *Resolution Methods for the Decision Problem*, volume 679 of *Lecture Notes in Computer Science*. Springer, 1993.
- [6] Oskar Fiuk and Emanuel Kieronski. Alternating quantifiers in uniform one-dimensional fragments with an excursion into three-variable logic. *Log. Methods Comput. Sci.*, 21(1), 2025.
- [7] Oskar Fiuk, Emanuel Kieronski, and Vincent Michielini. On the complexity of Maslov's class K. In Paweł Sobociński, Ugo Dal Lago, and Javier Esparza, editors, *Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2024, Tallinn, Estonia, July 8-11, 2024*, pages 35:1–35:14. ACM, 2024.
- [8] Сергей Ю. Маслов. Обратный метод установления выводимости для логических исчислений. *Тр. МИАН СССР*, 98:25–95, 1968.
- [9] Harald Ganzinger, Christoph Meyer, and Margus Veanes. The two-variable guarded fragment with transitive relations. In *14th Annual IEEE Symposium on Logic in Computer Science, Trento, Italy, July 2-5, 1999*, pages 24–34. IEEE Computer Society, 1999.
- [10] Kurt Gödel. Zum Entscheidungsproblem des logischen Funktionenkalküls. *Monatshefte für Mathematik und Physik*, 40:433–443, 1933.
- [11] Warren D. Goldfarb. The unsolvability of the Gödel class with identity. *J. Symb. Log.*, 49(4):1237–1252, 1984.
- [12] Warren D. Goldfarb. Random models and the Maslov class. *J. Symb. Log.*, 54(2):460–466, 1989.
- [13] Warren D. Goldfarb. Random models and solvable Skolem classes. *J. Symb. Log.*, 58(3):908–914, 1993.
- [14] Erich Grädel. On the restraining power of guards. *J. Symb. Log.*, 64(4):1719–1742, 1999.
- [15] Yuri Gurevich and Saharon Shelah. Random models and the Gödel case of the decision problem. *J. Symb. Log.*, 48(4):1120–1124, 1983.
- [16] Lauri Hella and Antti Kuusisto. One-dimensional fragment of first-order logic. In Rajeev Goré, Barteld P. Kooi, and Agi Kurucz, editors, *Advances in Modal Logic 10, invited and contributed papers from the tenth conference on "Advances in Modal Logic," held in Groningen, The Netherlands, August 5-8, 2014*, pages 274–293. College Publications, 2014.
- [17] Ullrich Hustadt and Renate A. Schmidt. Maslov's class K revisited. In Harald Ganzinger, editor, *Automated Deduction - CADE-16, 16th International Conference on Automated Deduction, Trento, Italy, July 7-10, 1999, Proceedings*, volume 1632 of *Lecture Notes in Computer Science*, pages 172–186. Springer, 1999.
- [18] Ullrich Hustadt, Renate A. Schmidt, and Lilia Georgieva. A survey of decidable first-order fragments and description logics. *Journal of Relational Methods in Computer Science*, 1:2004, 2004.
- [19] Emanuel Kieronski and Antti Kuusisto. Complexity and expressivity of uniform one-dimensional fragment with equality. In Erzsébet Csuhaj-Varjú, Martin Dietzfelbinger, and Zoltán Ésik, editors, *Mathematical Foundations of Computer Science 2014 - 39th International Symposium, MFCS 2014, Budapest, Hungary, August 25-29, 2014. Proceedings, Part I*, volume 8634 of *Lecture Notes in Computer Science*, pages 365–376. Springer, 2014.

- [20] Emanuel Kieronski, Jakub Michaliszyn, Ian Pratt-Hartmann, and Lidia Tendera. Two-variable first-order logic with equivalence closure. *SIAM J. Comput.*, 43(3):1012–1063, 2014.
- [21] Emanuel Kieronski and Martin Otto. Small substructures and decidability issues for first-order logic with two variables. *J. Symb. Log.*, 77(3):729–765, 2012.
- [22] Emanuel Kieronski and Lidia Tendera. On finite satisfiability of two-variable first-order logic with equivalence relations. In *Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS 2009, 11-14 August 2009, Los Angeles, CA, USA*, pages 123–132. IEEE Computer Society, 2009.
- [23] Antti Kuusisto. On the uniform one-dimensional fragment. In Maurizio Lenzerini and Rafael Peñaloza, editors, *Proceedings of the 29th International Workshop on Description Logics, Cape Town, South Africa, April 22-25, 2016*, volume 1577 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2016.
- [24] Vladimir Lifschitz. What is the inverse method? *J. Autom. Reason.*, 5(1):1–23, mar 1989.
- [25] Leopold Löwenheim. Über Möglichkeiten im Relativkalkül. *Mathematische Annalen*, 76:447–470, 1915.
- [26] Sergei J. Maslov. The inverse method for establishing deducibility for logical calculi. *The Calculi of Symbolic Logic I: Proceedings of the Steklov Institute of Mathematics*, 98, 1971.
- [27] Martin Otto. Two variable first-order logic over ordered domains. *J. Symb. Log.*, 66(2):685–702, 2001.
- [28] Ian Pratt-Hartmann. Finite satisfiability for two-variable, first-order logic with one transitive relation is decidable. *Math. Log. Q.*, 64(3):218–248, 2018.
- [29] Ian Pratt-Hartmann. *Fragments of First-Order Logic*. Oxford Logic Guides. Oxford University Press, United Kingdom, 2023.
- [30] Ian Pratt-Hartmann, Wiesław Szust, and Lidia Tendera. The fluted fragment revisited. *J. Symb. Log.*, 84(3):1020–1048, 2019.
- [31] Thomas Schwentick and Thomas Zeume. Two-variable logic with two order relations. *Log. Methods Comput. Sci.*, 8(1), 2012.
- [32] Dana Scott. A decision method for validity of sentences in two variables. *J. Symb. Log.*, 27:477, 1962.
- [33] Marco Voigt. *Decidable fragments of first-order logic and of first-order linear arithmetic with uninterpreted predicates*. PhD thesis, Saarland University, Saarbrücken, Germany, Saarbrücken, Germany, 2019.