

MULTI-STRUCTURAL GAMES AND BEYOND

MARCO CARMOSINO ^a, RONALD FAGIN ^a, NEIL IMMERMANN ^b, PHOKION G. KOLAITIS ^{a,c},
JONATHAN LENCHNER ^a, AND RIK SENGUPTA ^{a,b}

^a IBM Research

e-mail address: mlc@ibm.com, fagin@us.ibm.com, lenchner@us.ibm.com, rik@ibm.com

^b University of Massachusetts Amherst

e-mail address: immerman@umass.edu

^c University of California Santa Cruz

e-mail address: kolaitis@ucsc.edu

ABSTRACT. Multi-structural (MS) games are combinatorial games that capture the number of quantifiers of first-order sentences. On the face of their definition, MS games differ from Ehrenfeucht-Fraïssé (EF) games in two ways: first, MS games are played on two sets of structures, while EF games are played on a pair of structures; second, in MS games, Duplicator can make any number of copies of structures. In the first part of this paper, we perform a finer analysis of MS games and develop a closer comparison of MS games with EF games. In particular, we point out that the use of sets of structures is of the essence and that when MS games are played on pairs of structures, they capture Boolean combinations of first-order sentences with a fixed number of quantifiers. After this, we focus on another important difference between MS games and EF games, namely, the necessity for Spoiler to play on top of a previous move in order to win some MS games. Via an analysis of the types realized during MS games, we delineate the expressive power of the variant of MS games in which Spoiler never plays on top of a previous move. In the second part we focus on simultaneously capturing number of quantifiers and number of variables in first-order logic. We show that natural variants of the MS game do *not* achieve this. We then introduce a new game, the quantifier-variable tree game, and show that it simultaneously captures the number of quantifiers and number of variables. We conclude by generalizing this game to a family of games, the *syntactic games*, that simultaneously capture reasonable syntactic measures and the number of variables.

1. INTRODUCTION

Combinatorial games are an effective tool to analyze the expressive power of logics on sets of structures. The prototypical example of such combinatorial games are the Ehrenfeucht-Fraïssé (EF) games [Ehr61, Fra54], played by two players, called Spoiler and Duplicator; EF games capture the quantifier rank of first-order sentences and have been used to analyze the expressive power of first-order logic (FO). Specifically, for every $r \in \mathbb{N}$, two structures \mathbf{A} and \mathbf{B} satisfy the same FO-sentences of quantifier rank r if and only if Duplicator wins the r -round EF game on (\mathbf{A}, \mathbf{B}) . Since this holds true for finite and infinite structures

Key words and phrases: logic, complexity, games.



alike, EF games can be used in finite model theory, unlike other tools from mathematical logic that fail in the finite realm, such as the compactness theorem. Variants of EF games in which the players start by playing unary relations were used to analyze the expressive power of monadic NP, which is the collection of problems in NP that are definable by sentences of existential monadic second-order logic [Fag75, dR84, Sch94, FSV95]. A different variant of EF games are the pebble games; they capture the number of variables and were used to analyze the expressive power of finite-variable first-order logic and infinitary logics [Bar77, Imm82, KV92].

Multi-structural (MS) games were introduced in [Imm81], re-discovered recently in [FLRV21], and investigated further in [FLVW22] and [CFI⁺24]. MS games capture the number of quantifiers of FO-sentences, which is another important parameter in understanding the expressive power of FO. If one is interested in showing that a property P of finite structures, such as connectivity or acyclicity, is not expressible by *any* fixed FO-sentence ψ , then both EF games and MS games work equally well. The original motivation of introducing MS games, however, came from the potential use of MS games to obtain lower bounds when a sequence $\{\psi_n\}_{n \geq 1}$ of FO-sentences is considered so that a property P is expressed by ψ_n on structures with at most n elements. Let $h(n)$ be a function from the natural numbers to the natural numbers. In [Imm81], $\text{QN}[h(n)]$ is defined to be the class of all properties P for which there is a *uniform* sequence $\{\psi_n\}_{n \geq 1}$, of FO-sentences such that ψ_n expresses P on structures with at most n elements, ψ_n has $O(h(n))$ quantifiers, and $h(n)$ is constructible in $\text{DSPACE}[h(n)]$, where DSPACE stands for deterministic space. From results in [Imm81], it follows that on the class of all ordered finite structures,

$$\text{NL} \subseteq \text{QN}[\log(n)].$$

Thus, proving that some NP-complete or some P-complete problem requires $\omega(\log(n))$ quantifiers would separate NL from NP or from P, and this could be achieved by playing MS games on ordered finite structures. In contrast, as pointed out in [Imm81], FO-sentences of quantifier rank $O(\log(n))$ capture isomorphism on ordered finite structures, and hence EF games cannot be used to prove lower bounds for sequences of FO-sentences with $\omega(\log(n))$ quantifiers. Even though the early optimism of using combinatorial games to separate complexity classes has yet to materialize, it is still worth exploring MS games further with the aim of developing a deeper understanding of their features and potential uses.

On the face of their definitions, EF games and MS games differ in the following ways. First, EF games are played on a pair (\mathbf{A}, \mathbf{B}) of structures, while MS games are played on a pair $(\mathcal{A}, \mathcal{B})$ of *sets* of structures; second, in each round of the MS game, Duplicator can make any number of copies of structures on the side that she has to play; and, third, in the EF game, Duplicator must maintain a partial isomorphism between \mathbf{A} and \mathbf{B} , while in the MS game, she must maintain a partial isomorphism between some structure from \mathcal{A} and some structure from \mathcal{B} . Spoiler wins the r -round EF game on (\mathbf{A}, \mathbf{B}) iff there is a FO-sentence of quantifier rank r that is true on \mathbf{A} and false on \mathbf{B} , while Spoiler wins the r -round MS game on $(\mathcal{A}, \mathcal{B})$ iff there is a FO-sentence with r quantifiers that is *separating* for $(\mathcal{A}, \mathcal{B})$, i.e., it is true on every structure in \mathcal{A} and false on every structure in \mathcal{B} .

In this paper, we carry out a finer analysis of MS games and shed additional light on how EF games compare with MS games. In particular, when applying MS games to obtain inexpressibility results, one can use infinite sets \mathcal{A} and \mathcal{B} of structures. We show that the use of infinite sets does not yield stronger inexpressibility results; however, the sets used cannot always be assumed to be singletons (i.e., we cannot always start with

a pair of structures). We also show that r -round MS games restricted to singleton sets capture Boolean combinations of FO-sentences with r quantifiers. Finally, we add insight into Duplicator’s ability to make copies of structures during the game: if enough copies of each structure are made before the first round, then Duplicator does not need to make any additional copies during gameplay.

In [FLRV21] and [FLVW22], it was pointed out that, unlike the EF game, Spoiler may get an advantage in the MS game by playing “on top,” i.e., by placing a pebble on top of an existing pebble or a constant. Here, we first give a self-contained proof of the fact that, in general, Spoiler may not win an MS game without sometimes playing on top. We then explore the expressive power of the variant of the MS game in which Spoiler never plays on top. By analyzing the *types* (i.e., the combinations of atomic and negated atomic formulas satisfied by the pebbles played and the constants), we characterize when Spoiler wins the r -round MS game on $(\mathcal{A}, \mathcal{B})$ without playing on top in terms of properties of separating sentences for $(\mathcal{A}, \mathcal{B})$.

We next study combinatorial games that simultaneously capture number of quantifiers and number of variables. It is well known that EF games can be adapted to simultaneously capture quantifier rank and number of variables by limiting the number of pebbles used. In contrast, we show that natural variants of MS games obtained by limiting the number of pebbles fail to simultaneously capture number of quantifiers and number of variables. For this reason, we introduce a new game, the *quantifier-variable tree* game (QVT game). The QVT game is inspired by the Adler-Immerman game [AI01], which was introduced to study formula size. Variants of the Adler-Immerman game were subsequently investigated in [GS05] to study the succinctness of logics, and in [HV15] to obtain lower bounds for formula size in propositional logic and in FO. Our main result about the QVT game asserts that Spoiler wins the r -round, k -pebble QVT game on a pair $(\mathcal{A}, \mathcal{B})$ of sets of structures if and only if there is a separating sentence for $(\mathcal{A}, \mathcal{B})$ with r quantifiers and k variables.

Finally, we generalize the QVT game to a class of two-player games called *syntactic games*, that simultaneously capture some other measure of first-order formulas and the number of variables. We focus on *compositional* syntactic measures, i.e., measures on first-order formulas defined inductively for formulas in terms of its value on their subformulas. Quantifier count, quantifier rank, and formula size are all examples of such measures. Our syntactic games provide a unified game-theoretic setting in which to study these measures. A similar game for formulas of the infinitary logic $\mathcal{L}_{\omega_1, \omega}$ was studied in [VW13].

In conclusion, our results yield new insights into MS games, their variants, and generalizations. The next step in this investigation would be to use these games to determine the optimal value of a complexity measure and number of variables to express various combinatorial properties as a function of the size of the relevant structures.

2. PRELIMINARIES

Throughout this paper, we consider a relational schema τ consisting of finitely many relation and constant symbols. We denote the set of relation symbols in τ by $\text{Pred}(\tau)$. We write τ -structures in boldface, e.g., \mathbf{A} . We denote the universe (i.e., set of elements) of \mathbf{A} by A , and sets of τ -structures using calligraphic typeface, e.g., \mathcal{A} .

2.1. Pebbled structures. We have a set \mathcal{C} of *pebble colors*, with arbitrarily many *pebbles* of each color available. A τ -structure \mathbf{A} is *pebbled* if zero or more elements from A have one or more pebbles on them, so that *at most* one pebble of each color is present in \mathbf{A} . If pebbles of color $1, \dots, t$ are placed on elements $a_1, \dots, a_t \in A$, we refer to this pebbled structure as $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$. For readability, when the context is obvious, we refer to $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ as the t -pebbled (or simply pebbled) structure \mathbf{A} . Note that the unpebbled structure \mathbf{A} corresponds to the pebbled structure $\langle \mathbf{A} \mid \rangle$, with an empty set of pebbles. When $t = 0$, the notation $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ refers to the unpebbled structure $\langle \mathbf{A} \mid \rangle$.

Definition 2.1. Consider two pebbled structures $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ and $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle$, and let $f : A \rightarrow B$ be the map such that:

- for $1 \leq i \leq t$, we have that f maps $a_i \mapsto b_i$.
- for each constant symbol c in τ , we have that f maps $c^{\mathbf{A}} \mapsto c^{\mathbf{B}}$.

We say the pair $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ and $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle$ of pebbled structures *matches* (and call it a *matching pair*) iff the map f defined above is an isomorphism between the substructures of \mathbf{A} and \mathbf{B} induced by its domain and range.

Of course, note that the pebbled structures $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ and $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle$ may form a matching pair even if \mathbf{A} and \mathbf{B} are not isomorphic. In fact, they may even form a matching pair without any pebbles placed on them.

2.2. The multi-structural game. Fix $r \in \mathbb{N}$, and two nonempty sets \mathcal{A} and \mathcal{B} of τ -structures. We define the r -round multi-structural game on $(\mathcal{A}, \mathcal{B})$ as follows. We refer to \mathcal{A} and \mathcal{B} as the *left* and *right* sides, respectively. Informally, in each round of the game, one player (Spoiler) places a pebble on every structure on either the left side or the right side. Then the other player (Duplicator) places a pebble of the same color on every structure on the other side. Duplicator (but not Spoiler) may copy structures to make different placements on distinct copies of the same structure. Duplicator tries to ensure that, at the end of every round, there is some matching pair (Definition 2.1); Spoiler tries to eliminate all matching pairs in r or fewer rounds.

We now give a formal description of the game. We view all structures henceforth as pebbled structures. The game proceeds over up to r rounds, by building a sequence of *configurations* $(\mathcal{A}_0, \mathcal{B}_0), (\mathcal{A}_1, \mathcal{B}_1), \dots$, where $\mathcal{A}_0 = \mathcal{A}$ and $\mathcal{B}_0 = \mathcal{B}$. Initially, if there are no matching pairs, then Spoiler wins the 0-round game. Otherwise, inductively, in each round $t = 1, \dots, r$, with configuration $(\mathcal{A}_{t-1}, \mathcal{B}_{t-1})$, Spoiler selects one of the following moves:

- **Play-Left:**
 - (a) For each pebbled structure $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle \in \mathcal{A}_{t-1}$, Spoiler places a pebble colored t on an element $a_t \in A$, creating the pebbled structure $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$. Call the resulting set \mathcal{A}_t .
 - (b) For each pebbled structure $\langle \mathbf{B} \mid b_1, \dots, b_{t-1} \rangle \in \mathcal{B}_{t-1}$, Duplicator *may* make any number of copies of this pebbled structure, and then for each such copy, *must* place a pebble colored t on an element $b_t \in B$, creating the pebbled structure $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle$. Call the resulting set \mathcal{B}_t .
 - (c) At the end of this round, if no t -pebbled structure in \mathcal{A}_t forms a matching pair with a t -pebbled structure in \mathcal{B}_t , then Spoiler wins the game. Otherwise, if $t < r$, play continues.

- **Play-Right:** This move is dual to **Play-Left**; Spoiler places pebbles colored t on the pebbled structures on the right side, and Duplicator responds on the left side.

At the end of round r , Duplicator wins the game if there is still some r -pebbled structure on the left and some r -pebbled structure on the right forming a matching pair.

After t rounds of the r -round MS game on $(\mathcal{A}, \mathcal{B})$ have been played, we have the collection \mathcal{A}_t of pebbled structures $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ on the left, and the collection \mathcal{B}_t of pebbled structures $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle$ on the right, where \mathbf{A} and \mathbf{B} range over \mathcal{A} and \mathcal{B} respectively, and a_i and b_i are the elements pebbled by the players on \mathbf{A} and \mathbf{B} in round i , for $1 \leq i \leq t$. A *strategy* for Spoiler is a function that takes as input such a configuration $(\mathcal{A}_t, \mathcal{B}_t)$, and provides him with the next move in the game. Specifically, given such a configuration, a strategy tells Spoiler whether to play **Play-Left** or **Play-Right**; if it tells Spoiler to play **Play-Left**, then for each pebbled structure $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle \in \mathcal{A}_t$, the strategy provides an element $a_{t+1} \in A$ for Spoiler to play the pebble colored $t+1$ on. Similarly, if it tells Spoiler to play **Play-Right**, then for each pebbled structure $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle \in \mathcal{B}_t$, the strategy provides an element $b_{t+1} \in B$ for Spoiler to play the pebble colored $t+1$ on. A *winning* strategy for Spoiler in the r -round MS game on $(\mathcal{A}, \mathcal{B})$ is a strategy such that Spoiler always wins the game if he follows it.

We say that Duplicator follows the *oblivious* strategy if in each round, for every pebbled structure on the side she plays on, she makes as many copies as there are elements in the universe of the structure, and then plays a pebble on a different element in each copy. It is obvious that if Duplicator wins the r -round MS game on $(\mathcal{A}, \mathcal{B})$, then Duplicator can win by following the oblivious strategy. From now on, unless we say otherwise, we will assume that Duplicator follows the oblivious strategy in the MS game.

We next state an important definition.

Definition 2.2. Given two sets \mathcal{A} and \mathcal{B} of τ -structures, a *separating sentence* for $(\mathcal{A}, \mathcal{B})$ is a FO-sentence ψ such that every $\mathbf{A} \in \mathcal{A}$ has $\mathbf{A} \models \psi$, and every $\mathbf{B} \in \mathcal{B}$ has $\mathbf{B} \models \neg\psi$.

Clearly, ψ is a separating sentence for $(\mathcal{A}, \mathcal{B})$ if and only if $\neg\psi$ is a separating sentence for $(\mathcal{B}, \mathcal{A})$.

Next, we state without proof the fundamental theorem of MS games, which is in [Imm81, Theorem 10] and [FLRV21, Theorem 1.2].

Theorem 2.3 [Imm81, FLRV21]. *Spoiler has a winning strategy for the r -round MS game on $(\mathcal{A}, \mathcal{B})$ if and only if there is an r -quantifier separating sentence for $(\mathcal{A}, \mathcal{B})$.*

As discussed in [Imm81], Duplicator's ability to make copies of the structures in the MS game is crucial. Consider the next example (Figure 1).

The game is played on $\{\text{LO}(3)\}$, a singleton set consisting of a linear order of size 3 on the left, and $\{\text{LO}(2)\}$, a singleton set consisting of a linear order of size 2 on the right. Spoiler's only interesting move is to play the middle element B2 on the left, as indicated in red in the figure (it can be easily seen that all other moves by Spoiler lead to defeats). If this were now an EF game, Duplicator would lose — a response of L1 would be met with a Spoiler play on B1 in round 2, while a response of L2 would be met with a Spoiler play on B3 in round 2. However, in the MS game, Duplicator can make a second copy of $\text{LO}(2)$, and play *both* possible moves. It should now be obvious that it is impossible for Spoiler to win in the one remaining move, illustrating a clear difference between EF and MS games. Indeed, there is a separating sentence for $(\{\text{LO}(3)\}, \{\text{LO}(2)\})$ of quantifier rank 2, namely,




Figure 1: The beginning of a 2-round MS game on $(\mathcal{A}, \mathcal{B})$, where $\mathcal{A} = \{\text{LO}(3)\}$, the singleton linear order of size 3, and $\mathcal{B} = \{\text{LO}(2)\}$, the singleton linear order of size 2. Spoiler plays his first move (indicated in red) on B2, the middle element of $\text{LO}(3)$. Duplicator then makes a copy of $\text{LO}(2)$ and plays each of the two possible moves in response. This example is a slight variation on the examples in [FLRV21, FLVW22].

$\exists x(\exists y(y < x) \wedge \exists y(x < y))$, which uses three quantifiers, but no separating sentence with two quantifiers!

2.3. Strategies. A perusal of the proof of Theorem 2.3 in [Imm81] or in [FLRV21] reveals that if ψ is a separating sentence for $(\mathcal{A}, \mathcal{B})$ with r quantifiers, then Spoiler has a winning strategy for the r -round MS game on $(\mathcal{A}, \mathcal{B})$ where he “follows” ψ . To explain in more precise terms what “following” ψ means, assume that ψ is of the form $Q_1x_1 \dots Q_r x_r \theta$, where θ is quantifier-free. If Q_t is \exists , then Spoiler plays Play-Left, while if Q_t is \forall , then Spoiler plays Play-Right. Assume by induction that at the start of round $t \geq 1$, the configuration has the property that for every $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle$ on the left and $\langle \mathbf{B} \mid b_1, \dots, b_{t-1} \rangle$ on the right in this configuration, we have:

$$\begin{aligned} \mathbf{A} &\models Q_t x_t \dots Q_r x_r \theta(a_1/x_1, \dots, a_{t-1}/x_{t-1}) \\ \mathbf{B} &\models \neg Q_t x_t \dots Q_r x_r \theta(b_1/x_1, \dots, b_{t-1}/x_{t-1}). \end{aligned}$$

If Q_t is \exists , then there is some $a_t \in A$ such that $\mathbf{A} \models Q_{t+1}x_{t+1} \dots Q_r x_r \theta(a_1/x_1, \dots, a_t/x_t)$; in this case, Spoiler picks such an element a_t and plays the pebble on it. If Q_t is \forall , then $\mathbf{B} \models \exists x_t \neg Q_{t+1}x_{t+1} \dots Q_r x_r \theta(b_1/x_1, \dots, b_{t-1}/x_{t-1})$, and so there is an element $b_t \in B$ such that $\mathbf{B} \models \neg Q_{t+1}x_{t+1} \dots Q_r x_r \theta(b_1/x_1, \dots, b_t/x_t)$; in this case, Spoiler picks such an element b_t and plays the pebble on it. Since in each of these cases, more than one witness to the existential quantifier may exist, there may exist different strategies for Spoiler obtained by “following” ψ ; however, each of these strategies is a winning strategy for Spoiler. We say that a strategy for Spoiler is *obtained* from ψ if it is one of the strategies for Spoiler obtained by following ψ in this way.

The preceding discussion is summarized in the first part of the next result; the second part of the result can be extracted from the proof of Theorem 2.3.

Theorem 2.4. *Let $r \in \mathbb{N}$, and let \mathcal{A} and \mathcal{B} be two sets of τ -structures. Then, the following are true.*

- (a) *If $\psi \equiv Q_1x_1 \dots Q_r x_r \theta$ is a separating sentence for $(\mathcal{A}, \mathcal{B})$ with r quantifiers, then Spoiler can win the r -round MS game on $(\mathcal{A}, \mathcal{B})$ using a strategy obtained from ψ . Moreover, for every such strategy and for all pebbled structures $\langle \mathbf{A} \mid a_1, \dots, a_r \rangle$ on the left and $\langle \mathbf{B} \mid b_1, \dots, b_r \rangle$ on the right in the final configuration arising by using this strategy, we have $\mathbf{A} \models \theta(a_1/x_1, \dots, a_r/x_r)$ and $\mathbf{B} \models \neg \theta(b_1/x_1, \dots, b_r/x_r)$.*

- (b) *If Spoiler has a winning strategy \mathcal{S} for the r -round MS game on $(\mathcal{A}, \mathcal{B})$, then there is a separating sentence ψ for $(\mathcal{A}, \mathcal{B})$ with r quantifiers such that the strategy \mathcal{S} is one of the strategies for Spoiler obtained from ψ .*

2.4. Types. The following concepts will recur throughout this paper. An *atomic* formula is an expression of one of the following forms: an equality $x_i = x_j$ between two distinct variables; an equality $c_i = x_j$ between a constant symbol and a variable; an equality $c_i = c_j$ between two distinct constant symbols; an expression $R(y_1, \dots, y_m)$, where R is an m -ary relation symbol in the given schema, for some $m \geq 1$, and y_1, \dots, y_m are not-necessarily-distinct variables or constant symbols. A *type* over an r -tuple of distinct variables, x_1, \dots, x_r , is a conjunction of atomic and negated atomic formulas such that for each atomic formula with variables from x_1, \dots, x_r , exactly one of the atomic formula and its negation appears as a conjunct. We will often denote such a type by $t(x_1, \dots, x_r)$. If we wish to emphasize the numbers of variables, we can also refer to such a type as an *r -type*.

3. EHRENFUCHT-FRAÏSSÉ GAMES VS. MULTI-STRUCTURAL GAMES

Ehrenfeucht-Fraïssé games capture indistinguishability with respect to quantifier rank, while multi-structural games capture indistinguishability with respect to number of quantifiers. Each of these two families of games yields a method for establishing lower bounds in the corresponding fragment of first-order logic. Here, we compare the two methods and uncover differences between them.

A *property* of τ -structures is a Boolean query on τ , i.e., a function P such that for every τ -structure \mathbf{A} , either $P(\mathbf{A}) = 1$, or $P(\mathbf{B}) = 0$, and P is invariant under isomorphisms. If $P(\mathbf{A}) = 1$, we say that \mathbf{A} *satisfies* P . In the rest of the presentation, we shall say Spoiler (resp. Duplicator) *wins* a certain EF game to mean that he (resp. she) has a winning strategy in that game.

Method 3.1 (The EF Game Method). *Let $r \in \mathbb{N}$, and let P be a property of τ -structures. Then the following statements are equivalent.*

- (a) *No FO(τ)-sentence of quantifier rank r defines P .*
- (b) *There are τ -structures \mathbf{A} and \mathbf{B} such that \mathbf{A} satisfies P , \mathbf{B} does not satisfy P , and Duplicator wins the r -round EF game on (\mathbf{A}, \mathbf{B}) .*

The direction (b) \Rightarrow (a) captures the “soundness” of Method 3.1. It follows immediately from the basic property of EF games that if Duplicator wins the r -round EF game on (\mathbf{A}, \mathbf{B}) , then every FO-sentence of quantifier rank at most r that is true on \mathbf{A} is also true on \mathbf{B} . The direction (a) \Rightarrow (b) captures the “completeness” of Method 3.1. Its proof uses the fact that FO-sentences of quantifier rank r are closed under disjunctions, and also the following properties of the equivalence relation \equiv_r^{qr} , where $\mathbf{A} \equiv_r^{\text{qr}} \mathbf{B}$ means that \mathbf{A} and \mathbf{B} satisfy the same FO-sentences of quantifier rank at most r :

- (a) \equiv_r^{qr} has finitely many equivalence classes.
- (b) Each equivalence class of \equiv_r^{qr} is definable by a FO-sentence of quantifier rank r .

We now turn our attention to the MS game. The following result is an immediate consequence of Theorem 2.3 and the determinacy of the MS game.

Method 3.2 (The MS Game Method). *Let $r \in \mathbb{N}$, and let P be a property of τ -structures. Then the following statements are equivalent.*

- (a) No FO-sentence with r quantifiers defines P .
- (b) Duplicator wins the r -round MS game on $(\mathcal{A}(P), \mathcal{B}(P))$, where $\mathcal{A}(P)$ is the set of all τ -structures that satisfy P and $\mathcal{B}(P)$ is the set of all τ -structures that do not.

Clearly, for every property P , at least one of $\mathcal{A}(P)$ and $\mathcal{B}(P)$ is infinite. Thus, it is natural to ask: can the MS Game Method be used to prove inexpressibility results by considering finite sets of structures only?

Let $r \in \mathbb{N}$ and consider the number-of-quantifiers equivalence relation \equiv_r^{qn} , where $\mathbf{A} \equiv_r^{\text{qn}} \mathbf{B}$ means that \mathbf{A} and \mathbf{B} satisfy the same FO-sentences with r quantifiers. The following proposition contains some basic facts about \equiv_r^{qn} .

Proposition 3.3. *For every $r \in \mathbb{N}$, the following hold.*

- (a) \equiv_r^{qn} has finitely many equivalence classes.
- (b) For every τ -structure \mathbf{A} , the \equiv_r^{qn} -equivalence class of \mathbf{A} is definable by the conjunction of all FO-sentences with r quantifiers satisfied by \mathbf{A} .
- (c) Each \equiv_r^{qn} -equivalence class is definable by a FO-sentence, but it need not be definable by a FO-sentence with r quantifiers.

Proof. The equivalence relation \equiv_r^{qn} has finitely many equivalence classes¹ because, up to logical equivalence, there are only finitely many FO-sentences with r quantifiers. Recall that we made the assumption throughout that our fixed schema τ has finitely many relation symbols.

Take a structure \mathbf{A} and consider all FO-sentences with (at most) r quantifiers that are true on \mathbf{A} . Since, up to logical equivalence, there are finitely many such sentences, the conjunction $\theta_{\mathbf{A}}$ of these sentences is a FO-sentence. Moreover, $\theta_{\mathbf{A}}$ defines the \equiv_r^{qn} -equivalence class of \mathbf{A} , because FO-sentences with r quantifiers are closed under negation. Indeed, if \mathbf{B} satisfies $\theta_{\mathbf{A}}$, then clearly every FO-sentence with r quantifiers that is true on \mathbf{A} is also true on \mathbf{B} . The converse is also true, since, otherwise, we would have a FO-sentence ψ with r quantifiers that is true on \mathbf{B} but not on \mathbf{A} . But then $\neg\psi$ would be true on \mathbf{A} and false on \mathbf{B} , which is a contradiction.

In general, the sentence $\theta_{\mathbf{A}}$ need not be logically equivalent to any FO-sentences with r quantifiers. To see this, consider the case $r = 1$ and a relational schema consisting of a binary relation symbol R . Let \mathbf{A} be the structure consisting of a self-loop $R(a, a)$ and an isolated node b . Then \mathbf{A} satisfies $\exists x R(x, x) \wedge \exists y \neg R(y, y)$, but no FO-sentence with a single quantifier logically implies $\exists x R(x, x) \wedge \exists y \neg R(y, y)$, and hence the \equiv_r^{qn} -equivalence class of \mathbf{A} cannot be expressed by a FO-sentence with a single quantifier. This generalizes to $r > 1$. \square

Note also that, unlike FO-sentences of quantifier rank r , FO-sentences with r quantifiers are not closed under finite disjunctions or finite conjunctions.

The next proposition shows that it suffices to consider finite sets in MS games.

Proposition 3.4. *Let $r \in \mathbb{N}$, let \mathcal{A} and \mathcal{B} be two sets of τ -structures, and let \mathcal{A}' and \mathcal{B}' be the sets of τ -structures obtained from \mathcal{A} and \mathcal{B} by keeping exactly one structure from each equivalence class \equiv_r^{qn} with members in \mathcal{A} and \mathcal{B} , respectively. Then the following hold.*

- (a) The sets \mathcal{A}' and \mathcal{B}' are finite.
- (b) Duplicator wins the r -round MS game on $(\mathcal{A}, \mathcal{B})$ if and only if Duplicator wins the r -round MS game on $(\mathcal{A}', \mathcal{B}')$.

¹Technically speaking, this is true only once we have fixed the schema τ . For the sake of readability, we make this assumption throughout.

Proof. The first part follows immediately from Proposition 3.3, by the fact that the equivalence relation \equiv_r^{qn} has finitely many equivalence classes.

For the second part, since $\mathcal{A}' \subseteq \mathcal{A}$ and $\mathcal{B}' \subseteq \mathcal{B}$, if Duplicator wins the r -round MS game on \mathcal{A}' and \mathcal{B}' , then Duplicator also wins the r -round MS game on \mathcal{A} and \mathcal{B} . For the other direction, assume that Duplicator wins the r -round MS game on \mathcal{A} and \mathcal{B} . We have to show that Duplicator wins the r -round MS game on \mathcal{A}' and \mathcal{B}' . If this were not true, then Spoiler wins the r -round MS game on \mathcal{A}' and \mathcal{B}' . Hence, by Theorem 2.3, there is a FO-sentence ψ with r quantifiers such that ψ is true on every structure in \mathcal{A}' and false on every structure in \mathcal{B}' . Since every structure in \mathcal{A} (resp. \mathcal{B}) is \equiv_r^{qn} to a structure in \mathcal{A}' (resp. \mathcal{B}'), it follows that ψ is true on every structure in \mathcal{A} and false on every structure in \mathcal{B} . Hence, by Theorem 2.3, Spoiler wins the r -round game on \mathcal{A} and \mathcal{B} , which is a contradiction. \square

By combining Method 3.2 and Proposition 3.4, we obtain a method that brings Method 3.2 closer to Method 3.1.

Method 3.5 (The MS Game Method revisited). *Let $r \in \mathbb{N}$, and let P be a property of τ -structures. Then the following statements are equivalent.*

- (a) *No FO-sentence with r quantifiers defines P .*
- (b) *There are finite sets \mathcal{A} and \mathcal{B} of τ -structures such that every structure in \mathcal{A} satisfies P , no structure in \mathcal{B} satisfies P , and Duplicator wins the r -round MS game on $(\mathcal{A}, \mathcal{B})$.*

Can Method 3.5 be further refined so that inexpressibility results about FO-sentences with r quantifiers are proved by playing MS games on pairs of *singleton* sets? It is clear that if Spoiler wins the r -round MS game on $(\mathcal{A}, \mathcal{B})$, then for every $\mathbf{A} \in \mathcal{A}$ and every $\mathbf{B} \in \mathcal{B}$, Spoiler wins the r -round MS game on $(\{\mathbf{A}\}, \{\mathbf{B}\})$. The next example shows that the converse need not be true, even if $r = 1$.

Example 3.6. Consider a schema with two unary predicates R and G , and consider the structures \mathbf{A} , \mathbf{B} , and \mathbf{C} , where

- \mathbf{A} has $\{a_1, a_2\}$ as its universe, $R^{\mathbf{A}} = \{a_1\}$, $G^{\mathbf{A}} = \{a_2\}$.
- \mathbf{B} has $\{b\}$ as its universe, $R^{\mathbf{B}} = \{b\}$, $G^{\mathbf{B}} = \emptyset$.
- \mathbf{C} has $\{c\}$ as its universe, $R^{\mathbf{C}} = \emptyset$, $G^{\mathbf{C}} = \{c\}$.

It is easy to check that Spoiler wins the 1-round MS game on $(\{\mathbf{A}\}, \{\mathbf{B}\})$ and also on $(\{\mathbf{A}\}, \{\mathbf{C}\})$, but Duplicator wins the 1-round MS game on $(\{\mathbf{A}\}, \{\mathbf{B}, \mathbf{C}\})$.

Thus, Duplicator may win the MS game on two sets of structures, but Duplicator may *not* win the MS game on any pair of structures from these two sets.

What do MS games restricted to singleton sets of structures capture? The next result gives the answer.

Theorem 3.7. *Let $r \in \mathbb{N}$, and let P be a property of τ -structures. Then the following statements are equivalent.*

- (a) *The property P is definable by a Boolean combination of FO-sentences each with r quantifiers.*
- (b) *For all τ -structures \mathbf{A} and \mathbf{B} , if \mathbf{A} satisfies property P and Duplicator wins the r -round MS game on $(\{\mathbf{A}\}, \{\mathbf{B}\})$, then \mathbf{B} satisfies property P .*

Proof. Assume that P is a property definable by a Boolean combination ψ of FO-sentences each of which has r quantifiers. WLOG assume that ψ is in disjunctive normal form, i.e., ψ is a disjunction $\bigvee_{i=1}^m \theta_i$, where each θ_i is a conjunction of FO-sentences each with r quantifiers

(observe that the negation of a FO-sentence with r quantifiers is a FO-sentence with r quantifiers). Assume now that \mathbf{A} is a structure that satisfies P and that Duplicator wins the r -round MS game on $(\{\mathbf{A}\}, \{\mathbf{B}\})$. Since \mathbf{A} satisfies P , there is some $i \leq m$ such that $\mathbf{A} \models \theta_i$. Since Duplicator wins the r -round MS game on $(\{\mathbf{A}\}, \{\mathbf{B}\})$, Theorem 2.3 implies that every FO-sentence with r quantifiers that is true on \mathbf{A} is also true on \mathbf{B} , so that every conjunct of θ_i is true on \mathbf{B} , and hence $\mathbf{B} \models \theta_i$. This implies that $\mathbf{B} \models \psi$ and so \mathbf{B} satisfies P .

Conversely, assume that for all structures \mathbf{A} and \mathbf{B} , if \mathbf{A} satisfies P and Duplicator wins the r -round MS game on $(\{\mathbf{A}\}, \{\mathbf{B}\})$, then \mathbf{B} satisfies P . For every structure \mathbf{A} satisfying P , let $\theta_{\mathbf{A}}$ be the conjunction of all FO-sentences with r quantifiers satisfied by \mathbf{A} . By Proposition 3.3, $\theta_{\mathbf{A}}$ is a FO-sentence that defines the \equiv_r^{qn} -equivalence class of \mathbf{A} . Let ψ be the (finite) disjunction $\bigvee_{\mathbf{A} \models P} \theta_{\mathbf{A}}$. Clearly, ψ is a Boolean combination of FO-sentences each with r quantifiers. We claim that ψ defines P . Assume first that \mathbf{B} is a structure that satisfies P . Then $\mathbf{B} \models \theta_{\mathbf{B}}$, and $\theta_{\mathbf{B}}$ is a disjunct of ψ , hence $\mathbf{B} \models \psi$. Conversely, assume that \mathbf{B} is a structure such that $\mathbf{B} \models \psi$. Hence, there is a structure \mathbf{A} such that \mathbf{A} satisfies P and $\mathbf{B} \models \theta_{\mathbf{A}}$. Since $\theta_{\mathbf{A}}$ defines the \equiv_r^{qn} -equivalence class of \mathbf{A} , it follows that \mathbf{A} and \mathbf{B} satisfy the same FO-sentences with r quantifiers. Therefore, by Theorem 2.3, we have that Duplicator wins the r -round MS game on $(\{\mathbf{A}\}, \{\mathbf{B}\})$, hence, by our hypothesis, we have that \mathbf{B} satisfies P . \square

As an immediate consequence of Theorem 3.7, we obtain the following method about MS games on singleton sets.

Method 3.8. *Let $r \in \mathbb{N}$, and let P be a property of τ -structures. Then the following statements are equivalent.*

- (a) *There is no Boolean combination of FO-sentences each with r quantifiers that defines P .*
- (b) *There are τ -structures \mathbf{A} and \mathbf{B} such that \mathbf{A} satisfies P , \mathbf{B} does not satisfy P , and Duplicator wins the r -round MS game on $(\{\mathbf{A}\}, \{\mathbf{B}\})$.*

Because of the disjunctive normal form and since FO-sentences with r quantifiers are closed under negation, we could replace “Boolean combination” by “disjunction of conjunctions” in the statements of Theorem 3.7 and Method 3.8.

In summary, we now have the following rather complete picture about the similarities and the differences between EF games and MS games:

- (a) r -round EF games played on a pair of structures capture expressibility by a FO-sentence of quantifier rank r .
- (b) r -round MS games played on sets of structures capture expressibility by a FO-sentence with r quantifiers.

The sets of structures can be taken to be finite, but not singletons.

- (c) r -round MS games played on a pair of structures (i.e., on singletons) capture expressibility by a Boolean combination of FO-sentences, each with r quantifiers.

The difference between (a) and (c) is because FO-sentences of quantifier rank r are closed under disjunctions and conjunctions, whereas FO-sentences with r quantifiers are not.

We also have the following trade-off between the EF Game Method and the MS Game Method in establishing inexpressibility results about a property P : for EF games, we can work with pairs of structures, but we need to find them and to show that they have the desired properties in Method 3.1; in contrast, for MS games we can collect all the structures that satisfy property P on one side and all the structures that do not satisfy P on the other, but then it may be harder to show that Duplicator wins the MS game on these sets.

Duplicator’s “superpower” in MS games arises from the ability to make copies of structures while the game is played. The next theorem says that, if enough copies of each structure are made at the beginning of the game, then Duplicator does not need to make any further copies during gameplay.

Theorem 3.9. *Fix $r \in \mathbb{N}$, and let \mathcal{A} and \mathcal{B} be two sets of finite τ -structures. Then, the following are equivalent.*

- (a) *Duplicator wins the r -round MS game on $(\mathcal{A}, \mathcal{B})$.*
- (b) *Duplicator wins the r -round MS game **without** duplicating on $(\mathcal{A}^+, \mathcal{B}^+)$, where \mathcal{A}^+ is the multiset consisting of $|\mathcal{A}|^r$ copies of each $\mathbf{A} \in \mathcal{A}$, and \mathcal{B}^+ is the multiset consisting of $|\mathcal{B}|^r$ copies of each $\mathbf{B} \in \mathcal{B}$.*

Proof. Fix a number r of rounds. Let us refer to the ordinary MS game on $(\mathcal{A}, \mathcal{B})$ as the *original* game, and to the MS game on $(\mathcal{A}^+, \mathcal{B}^+)$ where Duplicator is not allowed to make copies as the *new* game. Recall that the game proceeds by building a sequence of configurations $(\mathcal{A}_0, \mathcal{B}_0), (\mathcal{A}_1, \mathcal{B}_1), \dots$, where $\mathcal{A}_0 = \mathcal{A}$ and $\mathcal{B}_0 = \mathcal{B}$. In particular, the configuration at the beginning of round t for $1 \leq t \leq r$ is $(\mathcal{A}_{t-1}, \mathcal{B}_{t-1})$, consisting of $(t-1)$ -pebbled structures (including when $t = 1$). We define the configurations $(\mathcal{A}_t^+, \mathcal{B}_t^+)$ of the new game analogously.

Claim 1. If Spoiler has a winning strategy in the original game on $(\mathcal{A}, \mathcal{B})$, then Spoiler has a winning strategy in the new game on $(\mathcal{A}^+, \mathcal{B}^+)$.

This direction is quite easy: if Spoiler has a winning strategy in the original game on $(\mathcal{A}, \mathcal{B})$, then by the Fundamental Theorem of MS Games, this means there is some separating FO sentence $\varphi \in \mathcal{L}(\tau)$ such that φ is true of all $\mathbf{A} \in \mathcal{A}$, and φ is false of all $\mathbf{B} \in \mathcal{B}$. Since the *set* of structures in \mathcal{A}^+ (resp. \mathcal{B}^+) is precisely the set of structures in \mathcal{A} (resp. \mathcal{B}), this must also mean that φ is a separating sentence for $(\mathcal{A}^+, \mathcal{B}^+)$ as well. Therefore, Spoiler has a winning strategy in $(\mathcal{A}^+, \mathcal{B}^+)$ for the original game. Then, since the only difference between the old and new games is that Duplicator is more restricted in the new game, Spoiler certainly has a winning strategy in the new game for $(\mathcal{A}^+, \mathcal{B}^+)$ as well.

For an alternative argument involving an explicit description of Spoiler’s winning strategy, we proceed as follows.

Fix an arbitrary instance $(\mathcal{A}, \mathcal{B})$ of the original game, and assume that Spoiler has a winning strategy on this instance in the original game (and WLOG, assume that Duplicator plays obliviously, so that Spoiler’s strategy is *non-adaptive* in the following sense: his sequence of moves in the game is determined at the start of the game, before any move has been played). Consider the corresponding instance $(\mathcal{A}^+, \mathcal{B}^+)$ of the new game. We shall define a strategy for Spoiler to win this instance $(\mathcal{A}^+, \mathcal{B}^+)$ in the new game.

Suppose WLOG in round t , Spoiler plays in \mathcal{A}_{t-1} in the original game; this is well-defined since Spoiler’s strategy is non-adaptive. We shall ensure that in round t of the new game, Spoiler plays in \mathcal{A}_{t-1}^+ . A symmetric argument will hold if Spoiler plays in \mathcal{B}_{t-1} .

To lift Spoiler’s move in round t to a move in the new game, we will ensure that the following invariant holds for all t satisfying $1 \leq t \leq r + 1$:

Invariant 1. Every $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle \in \mathcal{A}_{t-1}^+$ in the new game is present in \mathcal{A}_{t-1} in the original game. Similarly, every $\langle \mathbf{B} \mid b_1, \dots, b_{t-1} \rangle \in \mathcal{B}_{t-1}^+$ in the new game is present in \mathcal{B}_{t-1} in the original game.

As long as Invariant 1 holds, Spoiler can take *all* copies of any particular pebbled structure in the new game, and play the move he plays on that structure in the original game; this would be a well-specified strategy.

To see that Invariant 1 holds under this strategy, we shall use induction on t . Clearly, it is true for $t = 1$, by the definition of \mathcal{A}^+ and \mathcal{B}^+ . Suppose Invariant 1 is true of $t - 1$, and assume Spoiler follows the strategy above in round t .

Consider a pebbled structure $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle \in \mathcal{A}_t^+$ in the new game. This structure arose because in round t , Spoiler played pebble t on element $a_t \in A$ in the pebbled structure $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle$. Therefore, the structure $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle$ was in \mathcal{A}_{t-1}^+ . By induction, this structure $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle$ was also in \mathcal{A}_{t-1} in the original game. By our specified strategy, Spoiler played pebble t on $a_t \in A$ in the new game precisely because he also made that same move in the original game. Therefore, the structure $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ exists at the end of round t in the original game (and hence is in \mathcal{A}_t).

Now consider a pebbled structure $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle \in \mathcal{B}_t^+$ in the new game. This structure arose in round t when Duplicator played pebble t on element $b_t \in B$ on the pebbled structure $\langle \mathbf{B} \mid b_1, \dots, b_{t-1} \rangle$. Hence, $\langle \mathbf{B} \mid b_1, \dots, b_{t-1} \rangle$ is in \mathcal{B}_{t-1}^+ . By induction, this pebbled structure $\langle \mathbf{B} \mid b_1, \dots, b_{t-1} \rangle$ was also in \mathcal{B}_{t-1} in the original game. Since Spoiler plays his round t move in \mathcal{A}_{t-1} , we know Duplicator moves obliviously in \mathcal{B}_{t-1} in the original game, and makes all possible moves, including a move that plays pebble t on $b_t \in B$. (Note that we make no assumptions about Duplicator's strategy in the *new* game.) Hence, the pebbled structure $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle$ arises during round t in the original game (and hence is in \mathcal{B}_t).

Therefore, Invariant 1 is true of t , and so by induction, Spoiler's strategy is well-specified.

By assumption, we used a winning strategy for Spoiler in the original game on $(\mathcal{A}, \mathcal{B})$; therefore, at the end of r rounds in the original game, no pair of pebbled structures from \mathcal{A}_r and \mathcal{B}_r forms a matching pair. Suppose in the new game there are pebbled structures $\langle \mathbf{A} \mid a_1, \dots, a_r \rangle \in \mathcal{A}_r^+$ and $\langle \mathbf{B} \mid b_1, \dots, b_r \rangle \in \mathcal{B}_r^+$ that form a matching pair. By Invariant 1 above, these also appear in \mathcal{A}_r and \mathcal{B}_r in the original game, which is a contradiction. Therefore, after r rounds, no pair of pebbled structures from \mathcal{A}_r^+ and \mathcal{B}_r^+ form a matching pair, and hence Spoiler wins the new game on $(\mathcal{A}^+, \mathcal{B}^+)$.

Claim 2. If Duplicator has a winning strategy in the original game on $(\mathcal{A}, \mathcal{B})$, then she has a winning strategy in the new game on $(\mathcal{A}^+, \mathcal{B}^+)$.

Fix an arbitrary instance $(\mathcal{A}, \mathcal{B})$ of the original game, and consider the corresponding instance $(\mathcal{A}^+, \mathcal{B}^+)$ of the new game. We shall define a strategy for Duplicator in this instance $(\mathcal{A}^+, \mathcal{B}^+)$ of the new game, and prove that it is well-specified. Once we do this, we shall show that this strategy in the new game essentially “simulates” the oblivious strategy in the original instance $(\mathcal{A}, \mathcal{B})$ in the original game; therefore, if Duplicator has a winning strategy in the original instance, then her oblivious strategy would win, and we would then be able to conclude that she wins in the new game as well.

To describe Duplicator's strategy, we need to describe her move in round t of the new game, for each t satisfying $1 \leq t \leq r$. For this, we will ensure that the following invariant holds for all t satisfying $0 \leq t \leq r + 1$:

Invariant 2. For each $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle \in \mathcal{A}_t^+$, there are $|A|^{r-t}$ copies of it in \mathcal{A}_t^+ . Similarly, for each $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle \in \mathcal{B}_t^+$, there are $|B|^{r-t}$ copies of it in \mathcal{B}_t^+ .

Note that Invariant 2 is certainly true for $t = 0$, by the definition of \mathcal{A}^+ and \mathcal{B}^+ . We proceed inductively, so suppose Invariant 2 is true of $t - 1$. Assume WLOG that in round t

of the new game, Spoiler plays in \mathcal{A}_{t-1}^+ . For each $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle \in \mathcal{A}_{t-1}^+$, Spoiler chooses some $a_t \in A$, and plays pebble t on a_t , creating $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle \in \mathcal{A}_t^+$. By induction, there are $|A|^{r-t+1}$ copies of $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle$, but at most $|A|$ distinct new pebbled structures that can be created as a result of playing pebble t . So, by the pigeonhole principle, there must be some t -pebbled structure $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ arising from $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle$ that has at least $|A|^{r-t}$ copies after this move. We call this move by Spoiler (i.e., playing pebble t on $a_t \in A$) his *favorite move* on the pebbled structure $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle \in \mathcal{A}_{t-1}^+$. Duplicator keeps these $|A|^{r-t}$ copies of $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ in \mathcal{A}_t^+ (breaking ties arbitrarily in case there are multiple such structures with at least that many copies), and deletes all other t -pebbled structures arising from $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle$. Note that deleting structures only makes it harder for Duplicator to win. This ensures that each set of $|A|^{r-t+1}$ copies of $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle \in \mathcal{A}_{t-1}^+$ gives rise to $|A|^{r-t}$ copies of *some* $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle \in \mathcal{A}_t^+$ (as a result of Spoiler playing his favorite move on those pebbled structures), and these are the only copies arising from that set.

Duplicator also has to respond in \mathcal{B}_{t-1}^+ . By induction, each structure $\langle \mathbf{B} \mid b_1, \dots, b_{t-1} \rangle \in \mathcal{B}_{t-1}^+$ has $|B|^{r-t+1}$ copies. For each $b_t \in B$, Duplicator plays pebble t on element b_t in exactly $|B|^{r-t}$ of these copies. This ensures that each set of $|B|^{r-t+1}$ copies of $\langle \mathbf{B} \mid b_1, \dots, b_{t-1} \rangle \in \mathcal{B}_{t-1}^+$ gives rise to $|B|$ new sets of t -pebbled structures in \mathcal{B}_t^+ , each consisting of $|B|^{r-t}$ copies of some $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle \in \mathcal{B}_t^+$ (and each corresponding to a distinct choice of b_t), and these are the only copies arising from that set.

Therefore, Invariant 2 is true of t ; by induction, Duplicator's strategy is well-specified.

We shall now prove that, if Duplicator has a winning strategy for $(\mathcal{A}, \mathcal{B})$ in the original game, then her strategy described above is winning in the new game for $(\mathcal{A}^+, \mathcal{B}^+)$. Assume henceforth that Duplicator has a winning strategy for $(\mathcal{A}, \mathcal{B})$ in the original game. In particular, her oblivious strategy wins in the original game.

Consider *any* sequence of moves by Spoiler in the new game on $(\mathcal{A}^+, \mathcal{B}^+)$, followed by Duplicator's responses as specified by her strategy above. This describes a complete run of the new game over r rounds. We shall now lift this to a complete run of the original game over r rounds. To do so, we shall maintain the following invariant for $0 \leq t \leq r$:

Invariant 3. The multiset \mathcal{A}_t^+ contains $|A|^{r-t}$ copies of a pebbled structure $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ if and only if \mathcal{A}_t contains one copy of the same structure (and similarly for \mathcal{B}_t^+ and \mathcal{B}_t).

Note that Invariant 3 is certainly true for $i = 0$, by the definition of \mathcal{A}^+ and \mathcal{B}^+ . We proceed inductively, so suppose Invariant 3 is true of $t - 1$. Assume WLOG that in round t of the new game, Spoiler plays in \mathcal{A}_{t-1}^+ . We know from Invariant 2 that each $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle \in \mathcal{A}_{t-1}^+$ comes in a set of $|A|^{r-t+1}$ copies. By induction, there is one copy of $\langle \mathbf{A} \mid a_1, \dots, a_{t-1} \rangle$ in \mathcal{A}_{t-1} in the original game. We shall specify Spoiler's move on this copy in the original game, in order to construct our valid run. By Duplicator's specified strategy in the new game, this set of $|A|^{r-t+1}$ copies gives rise (through Spoiler's favorite move) to a set of $|A|^{r-t}$ copies of some $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle \in \mathcal{A}_t^+$, with all other copies deleted. In the original game, we shall have Spoiler play his favorite move on the copy as well; namely, we shall have him play pebble t on element $a_t \in A$, creating $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle \in \mathcal{A}_t$. We repeat this process for all other structures in \mathcal{A}_{t-1} , and hence specify Spoiler's move in round t completely. We have Duplicator respond *obliviously* in \mathcal{B}_{t-1} in round t in the original game. By Invariant 2, we know each $\langle \mathbf{B} \mid b_1, \dots, b_{t-1} \rangle \in \mathcal{B}_{t-1}^+$ comes in a set of $|B|^{r-t+1}$ copies, which gives rise to $|B|$ sets each comprised of $|B|^{r-t}$ copies of $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle \in \mathcal{B}_t^+$

(each such set corresponding to a distinct choice of b_t). By induction, there was one copy of $\langle \mathbf{B} \mid b_1, \dots, b_{t-1} \rangle$ in \mathcal{B}_{t-1} . When Duplicator responds obliviously, she creates one copy of each $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle \in \mathcal{B}_t$ in the original game (one copy for each distinct choice for b_t). Therefore, each such copy rising in the original game corresponds to one of the sets of $|B|^{r-t}$ copies in the new game. This completely specifies Duplicator’s move in round t completely, and also proves that Invariant 3 is true of t as well, finishing the induction.

This specifies one particular run of the original game on $(\mathcal{A}, \mathcal{B})$, where in each round, Spoiler plays his favorite move for that round on all structures in his side, and Duplicator responds obliviously on the other side. Since by assumption, this is a winning instance for Duplicator, this means that her oblivious strategy wins against any Spoiler’s strategy, and so in particular on this run of the game composed of Spoiler’s sequence of favorite moves as well. In particular, there is some $\langle \mathbf{A} \mid a_1, \dots, a_r \rangle \in \mathcal{A}_r$ and some $\langle \mathbf{B} \mid b_1, \dots, b_r \rangle \in \mathcal{B}_r$ forming a matching pair. By Invariant 3, this matching pair is present in $(\mathcal{A}_r^+, \mathcal{B}_r^+)$ as well, and so Duplicator wins this run of the new game.

Since Spoiler’s moves in the new game were arbitrary, it follows that Duplicator has a winning strategy in the new game, and we are done. \square

4. PLAY-ON-TOP MOVES IN MS GAMES

In [FLRV21], the authors draw attention to a surprising difference between the EF game and the MS game. Specifically, in the EF game, Spoiler gains no advantage by ever placing a pebble on top of an existing pebble or a constant. As it turns out, there are instances of the MS game where Spoiler wins, but only if he may sometimes play “on top”, i.e., he places a pebble on a constant or on an existing pebble. In this section, we delineate the expressive power of the variant of the MS game in which Spoiler never plays on top.

We begin with an example which will be used later, in which Spoiler wins the 3-round MS game without playing on top. For $s = 3, 4$, let $\text{LO}(s)$ be the linear order of size s , and consider the sentence

$$\Phi_{3,\vee} : \forall x_1 \exists x_2 \exists x_3 (x_1 < x_2 < x_3 \vee x_2 < x_3 < x_1). \quad (4.1)$$

This sentence says that every element has two smaller elements or two larger elements, and so $\text{LO}(4) \models \Phi_{3,\vee}$, but $\text{LO}(3) \models \neg \Phi_{3,\vee}$. Since $\Phi_{3,\vee}$ is a separating sentence for $(\{\text{LO}(4)\}, \{\text{LO}(3)\})$, Spoiler wins the 3-round MS game on $(\{\text{LO}(4)\}, \{\text{LO}(3)\})$. Moreover, it is easy to verify that, by following the sentence $\Phi_{3,\vee}$, Spoiler can win the MS game on $(\{\text{LO}(4)\}, \{\text{LO}(3)\})$ without ever playing on top.

Ideas and results in [FLRV21] and [FLVW22] yield examples where Spoiler wins, but only by playing on top. We give a self-contained proof that Spoiler sometimes needs to play on top, using an example that involves smaller structures and only three rounds (which is the fewest rounds for which playing on top may make a difference if there are no constants in τ).

See Figure 2, whose left side contains $\text{RT}(4)$, a rooted tree whose longest branch has 4 nodes and, $\text{LO}(4)$. The right side contains $\text{RT}(3)$, whose longest branch has 3 nodes, and $\text{LO}(3)$. Let τ have the single binary relation symbol $<$, where $x < y$ means that x is a descendent of y in the trees as drawn, so that larger nodes consistently appear higher on the page.

The following proposition shows that Spoiler needs to play on top to win this instance in the 3-round MS game.




Figure 2: A rooted tree $\text{RT}(4)$ whose longest branch has 4 nodes, and a linear order $\text{LO}(4)$ of 4 nodes, drawn as a rooted tree (left); a rooted tree $\text{RT}(3)$ whose longest branch has 3 nodes and a linear order $\text{LO}(3)$ of 3 nodes (right).

Proposition 4.1. *Consider the sets $\mathcal{A} = \{\text{RT}(4), \text{LO}(4)\}$ and $\mathcal{B} = \{\text{RT}(3), \text{LO}(3)\}$ as depicted in Figure 5. Spoiler wins the 3-round MS game on $(\mathcal{A}, \mathcal{B})$, but cannot win this game without playing on top.*

Proof. Consider the FO-sentence

$$\Phi_{3,\exists} : \exists x_1 \forall x_2 \exists x_3 \left(\begin{array}{l} x_2 < x_1 \rightarrow x_3 > x_1 \wedge \\ x_2 > x_1 \rightarrow (x_3 \neq x_2 \wedge x_3 > x_1) \wedge \\ x_2 = x_1 \rightarrow x_3 < x_1 \end{array} \right), \quad (4.2)$$

which asserts that there is an element with one smaller element and two larger elements. Clearly, $\Phi_{3,\exists}$ has 3 quantifiers and is a separating sentence for $(\mathcal{A}, \mathcal{B})$; hence Theorem 2.3 implies that Spoiler wins the 3-round MS game on $(\mathcal{A}, \mathcal{B})$. We will show that every winning strategy for Spoiler in the 3-round MS game requires that Spoiler plays on top. As a stepping stone, we will first establish the following claim.

Claim 1: For every winning strategy for Spoiler in the 3-round MS game on $(\mathcal{A}, \mathcal{B})$, the following hold:

- (a) Spoiler's round 1 move must be Play-Left.
- (b) Spoiler's round 1 play on $\text{RT}(4)$ must be on either B11 or B12; furthermore, Spoiler's round 1 play on $\text{LO}(4)$ must be on either B2 or B3.
- (c) Spoiler's round 2 move must be Play-Right.
- (d) Spoiler's round 3 move must be Play-Left.

Note that (a), (c), and (d) also mean that every 3-quantifier separating sentence for $(\mathcal{A}, \mathcal{B})$ must have $\exists \forall \exists$ as its quantifier prefix.

Proof of Claim 1. (a) We begin by showing that Spoiler's first-round move must be on the left side. For this, it suffices to show that if Spoiler's first-round move is on the right side, then Duplicator can maintain a matching pair with one copy of $\text{RT}(4)$ and one copy of $\text{RT}(3)$ for three rounds. Indeed, in response to a first-round move by Spoiler on L1 in $\text{RT}(3)$, Duplicator plays on B1 in one copy of $\text{RT}(4)$, and then easily survives two more rounds just on this pair of pebbled structures (e.g., in response to a second-round Spoiler play on B12, Duplicator plays on L11 in one copy and L12 in another copy of the $\text{RT}(3)$). First-round Spoiler plays on L11 or L21 are met by a Duplicator play on B21 in one copy of $\text{RT}(4)$, and first-round Spoiler plays on L12 or L22 are met by a

Duplicator play on B22 in a copy of RT(4); in each case, Duplicator readily survives two more rounds on just this pair of structures. Thus, Spoiler’s round 1 move must be Play-Left.

- (b) If Spoiler’s first-round move on RT(4) is not on B11 or on B12, then once again, Duplicator survives three rounds just on the pair $(\{RT(4)\}, \{RT(3)\})$ by keeping a matching pair between two copies via the following responses in a copy of RT(3): $B1 \leftrightarrow L1$, $B13 \leftrightarrow L12$, $B21 \leftrightarrow L21$, $B22 \leftrightarrow L22$.

Furthermore, if Spoiler’s first-round move on LO(4) is on B1 or B4, Duplicator survives three rounds just on the pair $(\{LO(4)\}, \{LO(3)\})$ by responding on L1 or L3 respectively. Therefore, Spoiler’s first-round move on LO(4) must be on B2 or B3.

- (c) Figure 3 now depicts the state of the game after the first round (with Spoiler having played on B12 on the RT(4) and on B3 on the LO(4)), showing Duplicator’s oblivious response on the right side as well.




Figure 3: Configuration after round 1 of the MS game, with a particular choice of Spoiler’s round 1 moves in red on the left, and Duplicator’s oblivious responses in red on the right.

We now show that Spoiler’s second-round moves must be on the right side. We can focus entirely on the linear orders on the two sides. It suffices to consider Spoiler playing B3 in round 1 (as B2 would be symmetric). Figure 4 depicts the linear orders on the two sides after round 1.




Figure 4: Linear orders with moves in red after the first round of the MS game.

If Spoiler's second-round move on $\text{LO}(4)$ is on $B3$, he already plays on top, and we are done. If Spoiler's second-round move is on $B1$ in $\text{LO}(4)$, then Duplicator would play $L1$ in the second and third copies, and survive another round on one of these pairs. If Spoiler's second-round move is on $B2$, Duplicator would play $L1$ in the second copy and $L2$ in the third copy, and survive another round on one of these pairs. Finally, if Spoiler's second-round move is on $B4$, then Duplicator would play $L3$ on the second copy and survive one more round on this pair. It follows that Spoiler's round 2 move must be **Play-Right**.

- (d) If Spoiler's third-round move is also on the right side, then it is easy to check that Duplicator can maintain a matching pair with the $\text{LO}(4)$ on the left and the third copy of $\text{LO}(3)$ shown on the right. This proves Claim 1. \square

Now that Claim 1 has been established, we are ready to show that Spoiler cannot win the 3-round MS game on $(\mathcal{A}, \mathcal{B})$ without playing on top. We can make this argument by considering just the linear orders. By symmetry, we may assume that Spoiler's round 1 move on the $\text{LO}(4)$ is on $B3$. Let us focus on the second copy of $\text{LO}(3)$ on the right. If Spoiler's second-round move on that copy is on $L3$, then Duplicator responds on $B4$ in $\text{LO}(4)$ and wins on just this pair easily. Therefore, if Spoiler is to not move on top, he must play on $L1$ on this second copy of $\text{LO}(3)$, and on $L1$ or $L2$ on the third copy of $\text{LO}(3)$. If Spoiler plays on $L2$, then Duplicator responds on $B2$ in $\text{LO}(4)$ and is guaranteed a matching pair between $\text{LO}(4)$ and one of the copies of $\text{LO}(3)$; if Spoiler plays on $L1$, then Duplicator responds on $B1$, and the same conclusion holds.

This completes the proof that Spoiler must play on top in order to win the 3-round MS game on $(\mathcal{A}, \mathcal{B})$. \square

We now analyze the variant of the MS game in which Spoiler never plays on top. Recall that a *type* over an r -tuple of distinct variables, x_1, \dots, x_r , is a conjunction of atomic and

negated atomic formulas such that for each atomic formula with variables from x_1, \dots, x_r , exactly one of the atomic formula and its negation appears as a conjunct.

We assume that every FO-sentence with r quantifiers is in prenex normal form, that its quantifier-free part is a disjunction of types, and that the quantifier prefix is $Q_1x_1 \dots Q_rx_r$, where each Q_j is the quantifier \exists or the quantifier \forall . In other words, we assume that we work with FO-sentences of the form $Q_1x_1 \dots Q_rx_r\theta$, where θ is a disjunction of types.

Definition 4.2. Let ψ be a FO-sentence of the form $Q_1x_1 \dots Q_rx_r\theta$, where θ is a disjunction of types.

- A type $t(x_1, \dots, x_r)$ occurring as a disjunct of θ is *non-replicating* in ψ if conditions (a) and (b) below hold:
 - (a) For every two distinct variables x_i and x_j , if the equality $x_i = x_j$ appears in $t(x_1, \dots, x_r)$, then exactly one of the following two conditions holds:
 - (i) Both variables x_i and x_j are universally quantified;
 - (ii) One of the two variables is universally quantified, the other variable is existentially quantified, and the existential quantifier appears before the universal quantifier in the quantifier prefix (that is, if $Q_i = \exists$ and $Q_j = \forall$, then we must have $i < j$).
 - (b) If an equality $c_i = x_j$ appears in $t(x_1, \dots, x_r)$, then the variable x_j is universally quantified.
- A type $t(x_1, \dots, x_r)$ occurring as a disjunct of θ is *replicating* in ψ if $t(x_1, \dots, x_r)$ is not non-replicating in ψ .
- ψ is a *non-replicating* sentence if every type occurring as a disjunct of $\theta(x_1, \dots, x_r)$ of ψ is non-replicating in ψ .
- ψ is a *replicating* sentence if it is not a non-replicating sentence (i.e., at least one type occurring as a disjunct of $\theta(x_1, \dots, x_r)$ of ψ is replicating in ψ).

We illustrate the notion of a non-replicating sentence with several examples in which the $<$ relation is assumed to range over the elements of rooted trees (with not all elements necessarily related); thus, instead of spelling out full types explicitly, we will often give a part of the type that determines the full type.

Example 4.3. If $\psi \equiv Q_1x_1 \dots Q_rx_r\theta$ is such that every type in θ is equality-free, then ψ is non-replicating.

Example 4.4. Consider the FO-sentence

$$\Phi_{3,\forall} : \forall x_1 \exists x_2 \exists x_3 (x_1 < x_2 < x_3 \vee x_2 < x_3 < x_1), \quad (4.3)$$

encountered at the beginning of this section. We claim that both $\Phi_{3,\forall}$ and $\neg\Phi_{3,\forall}$ are non-replicating sentences.

To see this, first observe that the quantifier-free part of $\Phi_{3,\forall}$ consists of two equality-free types, hence $\Phi_{3,\forall}$ is a non-replicating sentence by Example 4.3.

Secondly, the negation $\neg\Phi_{3,\forall}$ of $\Phi_{3,\forall}$ is the FO-sentence

$$\exists x_1 \forall x_2 \forall x_3 \neg(x_1 < x_2 < x_3 \vee x_2 < x_3 < x_1). \quad (4.4)$$

This is a non-replicating sentence because the possible equalities in the types of its quantifier-free part are $x_1 = x_2$, $x_1 = x_3$, and $x_2 = x_3$. The first two involve the existentially quantified variable x_1 and the universally quantified variables x_2 or x_3 (hence, they satisfy condition (ii) in Definition 4.2(a)), while the third involves the two universally quantified variables x_2 and x_3 (hence, it satisfies condition (i) in Definition 4.2(a)).

Example 4.5. Consider the FO-sentence

$$\Phi_{3,\exists} : \exists x_1 \forall x_2 \exists x_3 \left(\begin{array}{l} x_2 < x_1 \rightarrow x_3 > x_1 \wedge \\ x_2 > x_1 \rightarrow (x_3 \neq x_2 \wedge x_3 > x_1) \wedge \\ x_2 = x_1 \rightarrow x_3 < x_1 \end{array} \right). \quad (4.5)$$

This sentence asserts that there is an element with one smaller element and two larger elements. We claim that $\Phi_{3,\exists}$ is a non-replicating sentence, but its negation $\neg\Phi_{3,\exists}$ is replicating.

First, it is easy to see that the quantifier-free part of $\Phi_{3,\exists}$ is logically equivalent to a disjunction of types in which the only equality is $x_1 = x_2$. Since x_1 is existentially quantified and x_2 is universally quantified in $\Phi_{3,\exists}$, the sentence $\Phi_{3,\exists}$ is non-replicating. The negation $\neg\Phi_{3,\exists}$ of $\Phi_{3,\exists}$ is the FO-sentence

$$\forall x_1 \exists x_2 \forall x_3 \neg \left(\begin{array}{l} x_2 < x_1 \rightarrow x_3 > x_1 \wedge \\ x_2 > x_1 \rightarrow (x_3 \neq x_2 \wedge x_3 > x_1) \wedge \\ x_2 = x_1 \rightarrow x_3 < x_1 \end{array} \right), \quad (4.6)$$

By pushing the negation inside, it is easy to see that, as a disjunction of types, the quantifier-free part of $\neg\Phi_{3,\exists}$ includes the type $x_1 = x_2 < x_3$ as a disjunct. Since x_1 is universally quantified and x_2 is existentially quantified in $\neg\Phi_{3,\exists}$, the sentence $\neg\Phi_{3,\exists}$ is replicating.

We are now ready to state and prove the main result about the expressive power of the variant of the MS game in which Spoiler never plays on top.

Theorem 4.6. *Let $r \in \mathbb{N}$, and let \mathcal{A} and \mathcal{B} be two sets of τ -structures. Then the following statements are equivalent:*

- (a) *Spoiler wins the r -round MS game on $(\mathcal{A}, \mathcal{B})$ without ever playing on top.*
- (b) *There is a separating sentence ψ for $(\mathcal{A}, \mathcal{B})$ of the form $Q_1 x_1 \dots Q_r x_r \theta(x_1, \dots, x_r)$, where θ is quantifier-free; moreover, if \mathcal{S} is a winning strategy for Spoiler obtained from ψ , then for every pebbled structure $\langle \mathbf{A} \mid a_1, \dots, a_r \rangle$ with $\mathbf{A} \in \mathcal{A}$, and for every pebbled structure $\langle \mathbf{B} \mid b_1, \dots, b_r \rangle$ with $\mathbf{B} \in \mathcal{B}$ arising by using this strategy, the following hold:*
 - (i) *The disjunct of θ satisfied by (a_1, \dots, a_r) is a non-replicating type in ψ .*
 - (ii) *The disjunct of $\neg\theta$ satisfied by (b_1, \dots, b_r) is a non-replicating type in $\neg\psi$.*

Proof. Assume that Spoiler wins the r -round MS game on \mathcal{A} and \mathcal{B} without ever playing on top. By Theorem 2.4, there is a FO-sentence $\psi \equiv Q_1 x_1 \dots Q_r x_r \theta(x_1, \dots, x_r)$ with r quantifiers such that ψ is separating for $(\mathcal{A}, \mathcal{B})$; moreover, if \mathcal{S} is a winning strategy for Spoiler obtained from ψ , then for every pebbled structure $\langle \mathbf{A} \mid a_1, \dots, a_r \rangle$ with $\mathbf{A} \in \mathcal{A}$, and every pebbled structure $\langle \mathbf{B} \mid b_1, \dots, b_r \rangle$ with $\mathbf{B} \in \mathcal{B}$ arising by using this strategy, we have that $\mathbf{A} \models \theta(a_1, \dots, a_r)$ and $\mathbf{B} \models \neg\theta(b_1, \dots, b_r)$. Fix two such pebbled structures $\langle \mathbf{A} \mid a_1, \dots, a_r \rangle$ with $\mathbf{A} \in \mathcal{A}$ and $\langle \mathbf{B} \mid b_1, \dots, b_r \rangle$ with $\mathbf{B} \in \mathcal{B}$. Let $t(x_1, \dots, x_r)$ be the type that occurs as a disjunct of $\theta(x_1, \dots, x_r)$ and is satisfied by (a_1, \dots, a_r) . We claim that $t(x_1, \dots, x_r)$ is a non-replicating type in ψ . Otherwise, one of the following three things would happen: (1) $t(x_1, \dots, x_r)$ would contain an equality $x_i = x_j$, where both the i -th and the j -th quantifier of ψ are \exists ; this implies that $a_i = a_j$, but since both a_i and a_j were played by Spoiler, this means that Spoiler played on top, which is a contradiction. (2) $t(x_1, \dots, x_r)$ would contain an equality $x_i = x_j$, where the i -th quantifier of ψ is \forall , the j -th quantifier of

ψ is \exists , and $i < j$; this implies that $a_i = a_j$, but since a_j was played by Spoiler, this means that Spoiler played on top, which is a contradiction. (3) $t(x_1, \dots, x_r)$ would contain an equality $c_i = x_j$, where the j -th quantifier of ψ is \exists ; this implies that $c_i = a_j$, but a_j was played by Spoiler, hence Spoiler played on top, which is a contradiction.

A similar argument shows that the disjunct of $\neg\theta$ satisfied by (b_1, \dots, b_r) is a non-replicating type in $\neg\psi$.

Conversely, assume that there is a separating sentence ψ for $(\mathcal{A}, \mathcal{B})$ of the form $Q_1x_1 \dots Q_r x_r \theta(x_1, \dots, x_r)$, where θ is quantifier-free; further, assume that if \mathcal{S} is a winning strategy for Spoiler obtained from ψ , then the properties asserted in (b) hold. We claim that Spoiler never plays on top using strategy \mathcal{S} . Consider a pebbled structure $\langle \mathbf{A} \mid a_1, \dots, a_r \rangle$ with $\mathbf{A} \in \mathcal{A}$ and a pebbled structure $\langle \mathbf{B} \mid b_1, \dots, b_r \rangle$ with $\mathbf{B} \in \mathcal{B}$ arising by using this strategy. Let us examine the j -th move of Spoiler, where $1 \leq j \leq r$. There are two cases to consider, namely, the case in which $Q_j = \exists$ and the case in which $Q_j = \forall$. If $Q_j = \exists$, then Spoiler has to place the j -th pebble on an element of \mathbf{A} . Since Spoiler plays using the strategy \mathcal{S} , Part (a) of Theorem 2.4 tells that $\mathbf{A} \models \theta(a_1, \dots, a_r)$. Consequently, the tuple (a_1, \dots, a_r) must satisfy a type $t(x_1, \dots, x_r)$ occurring as a disjunct of θ . Thus, $t(x_1, \dots, x_r)$ must be a non-replicating type in ψ . Since the j -th quantifier of ψ is \exists , the type $t(x_1, \dots, x_r)$ cannot contain an equality of the form $x_i = x_j$ with $i < j$ or an equality of the form $c_i = x_j$; instead, it must contain $\neg(x_i = x_j)$ and $\neg(c_i = x_j)$. Therefore, $a_i \neq a_j$, for every $i < j$, and also $c_i \neq a_j$, where c_i is a constant; hence, Spoiler does not play on top on \mathbf{A} . If $Q_j = \forall$, the argument is similar using the fact that the type in $\neg\theta$ satisfied by (b_1, \dots, b_r) must be non-replicating in $\neg\psi$. \square

Corollary 4.7. *Let $r \in \mathbb{N}$, and let \mathcal{A} and \mathcal{B} be two sets of τ -structures. If there is a separating sentence ψ for $(\mathcal{A}, \mathcal{B})$ with r quantifiers such that both ψ and $\neg\psi$ are non-replicating sentences, then Spoiler wins the r -round MS game on $(\mathcal{A}, \mathcal{B})$ without ever playing on top.*

Consider the variant of the MS game in which Spoiler never plays on top on the left and the variant of the MS game in which Spoiler never plays on top on the right. The proof of Theorem 4.6 can be adapted to yield analogous results for these two variants. For example, we have the following results, whose proofs we omit.

Theorem 4.8. *Let $r \in \mathbb{N}$, and let \mathcal{A} and \mathcal{B} be two sets of τ -structures. Then the following statements are equivalent:*

- (a) *Spoiler wins the r -round MS game on $(\mathcal{A}, \mathcal{B})$ without ever playing on top on the left side.*
- (b) *There is a separating sentence ψ for $(\mathcal{A}, \mathcal{B})$ of the form $Q_1x_1 \dots Q_r x_r \theta(x_1, \dots, x_r)$, where θ is quantifier-free; moreover, if \mathcal{S} is a winning strategy for Spoiler obtained from ψ , then for every $\langle \mathbf{A} \mid a_1, \dots, a_r \rangle$ with $\mathbf{A} \in \mathcal{A}$ arising by using this strategy, the disjunct of θ satisfied by (a_1, \dots, a_r) is a non-replicating type in ψ .*

Corollary 4.9. *Let $r \in \mathbb{N}$, and let \mathcal{A} and \mathcal{B} be two sets of τ -structures. If there is a separating sentence ψ for $(\mathcal{A}, \mathcal{B})$ with r quantifiers such that ψ is a non-replicating sentence, then Spoiler wins the r -round MS game on $(\mathcal{A}, \mathcal{B})$ without ever playing on top on the left side.*

We revisit Example 4.4. Since the sentence $\Phi_{3,\forall}$ asserts that every element has two smaller elements or two larger elements, $\Phi_{3,\forall}$ is a separating sentence for $(\{\text{LO}(4)\}, \{\text{LO}(3)\})$. According to Example 4.4, both $\Phi_{3,\forall}$ and $\neg\Phi_{3,\forall}$ are non-replicating sentences, hence Corollary

4.7 implies that Spoiler can win the 3-round MS game on $(\{\text{LO}(4)\}, \{\text{LO}(3)\})$ without playing on top. As mentioned earlier, we can also verify this directly. Note that $\Phi_{3,\forall}$ is a separating sentence for $(\{\text{LO}(4)\}, \{\text{RT}(3), \text{LO}(3)\})$ as well, hence Corollary 4.7 implies that Spoiler can win the 3-round MS game on $(\{\text{LO}(4)\}, \{\text{RT}(3), \text{LO}(3)\})$ without playing on top.

For a different application, we claim that Spoiler can win the 3-round MS game on $(\{\text{RT}(4), \text{LO}(4)\}, \{\text{LO}(3)\})$ without playing on top. For this, we first let $\text{inc}(x, y)$ be the formula $\neg(x < y) \wedge \neg(y < x) \wedge (x \neq y)$, asserting that x and y are distinct incomparable elements, and then let $\Psi_{3,\forall}$ be the sentence

$$\Psi_{3,\forall} : \forall x_1 \exists x_2 \exists x_3 (x_1 < x_2 < x_3 \vee x_2 < x_3 < x_1 \vee \text{inc}(x_1, x_2) \wedge \text{inc}(x_1, x_3) \wedge x_2 \neq x_3),$$

which asserts that every element has two bigger elements, or two smaller elements, or two different elements that are incomparable with it. Clearly, $\Psi_{3,\forall}$ is a separating sentence for $(\{\text{RT}(4), \text{LO}(4)\}, \{\text{LO}(3)\})$; moreover, it is easy to see that both $\Psi_{3,\forall}$ and its negation are non-replicating sentences, hence Corollary 4.7 implies that Spoiler can win the 3-round MS game on $(\{\text{RT}(4), \text{LO}(4)\}, \{\text{LO}(3)\})$ without playing on top.

Finally, Proposition 4.1 and Corollary 4.7 imply that there is no FO-sentence ψ with 3 quantifiers such that ψ is a separating sentence for $(\{\text{RT}(4), \text{LO}(4)\}, \{\text{RT}(3), \text{LO}(3)\})$ and both ψ and $\neg\psi$ are non-replicating sentences.

Corollary 4.7 gives a sufficient condition for Spoiler to win the r -round MS game without ever playing on top, namely, it suffices to have a separating sentence such that both the sentence and its negation are non-replicating. It is natural to ask whether this condition is also necessary, in which case it would replace the more complicated necessary and sufficient condition in Theorem 4.8. The next result answers this question in the negative.




Figure 5: A rooted tree $\text{RT}(4)$ whose longest branch has 4 nodes (left); a rooted tree $\text{RT}(3)$ whose longest branch has 3 nodes (right).

Theorem 4.10. *Let $\text{RT}(4)$ and $\text{RT}(3)$ be the rooted trees in Figure 5. Then the following statements are true.*

- (a) *Spoiler can win the 3-round MS game on $(\{\text{RT}(4)\}, \{\text{RT}(3)\})$ without playing on top.*
- (b) *There is no FO-sentence ψ with 3 quantifiers such that*
 - (i) *ψ is a separating sentence for $(\{\text{RT}(4)\}, \{\text{RT}(3)\})$;*
 - (ii) *both ψ and $\neg\psi$ are non-replicating sentences.*

Proof. We claim that Spoiler can win the 3-round MS game on $(\{\text{RT}(4)\}, \{\text{RT}(3)\})$ without playing on top by playing his first-round moves on the left, his second-round moves on the right, and his third-round moves on the left as described below.

- Spoiler plays first by placing a pebble on B12 in $\text{RT}(4)$ on the left.

- After Duplicator makes five copies of $RT(3)$ on the right and places a pebble on a different element of each copy, Spoiler places pebbles on each copy of $RT(3)$ as follows:
 - (a) In response to L1 on the first copy of $RT(3)$, Spoiler pebbles L11.
 - (b) In response to L11 on the second copy of $RT(3)$, Spoiler pebbles L1,
 - (c) In response to L12 on the third copy of $RT(3)$, Spoiler pebbles L21.
 - (d) In response to L21 on the fourth copy of $RT(3)$, Spoiler pebbles L1.
 - (e) In response to L22 on the fifth copy of $RT(3)$, Spoiler pebbles L11.
- After Duplicator makes six copies of $RT(4)$ on the left and places a pebble on a different element of each copy, Spoiler places pebbles on each copy of $RT(4)$ as indicated in Figure 6, where the first-round moves of the two players are colored red, the second-round-moves of the two players are colored blue, and Spoiler's third-round moves are colored green.

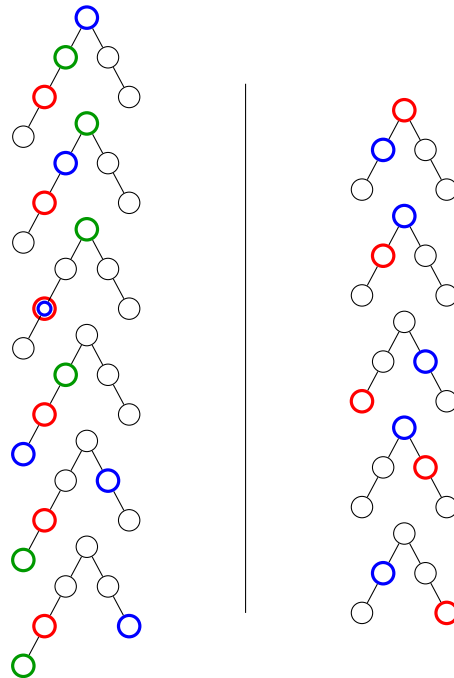


Figure 6: The configuration of the 3-round MS game after Spoiler's moves in the 3 rounds. Note that Spoiler has played Pebble-Left in round 1 with red, Pebble-Right in round 2 with blue, and Pebble-Left in round 3 with green.

From Figure 6, it is clear that Duplicator can not play on the right in the third round so that a partial isomorphism is maintained between one of the six copies of $RT(4)$ and one of the five copies of $RT(3)$. For instance, in four copies of $RT(4)$, the green pebble played by Spoiler are above the red one; in order to maintain a partial isomorphism with one of these copies, there are only two copies of $RT(3)$ for Duplicator as candidates – and in both of them, any choice of a green pebble above a red pebble would break all isomorphisms with the copies of $RT(4)$ on the left. Similarly, there are two copies of $RT(4)$ with a green pebble below a red pebble, so Duplicator has only two candidates among the copies of $RT(3)$ to place a green pebble below a red pebble – and in both of them, the blue pebble is comparable with at least one other pebble, which is not true in the corresponding copies of $RT(4)$ on the left. Thus, Spoiler wins the 3-round MS game on $(\{RT(4)\}, \{RT(3)\})$. Moreover, only B12 in the

third copy of $\text{RT}(4)$ has two pebbles placed on it with the red pebble placed by Spoiler in the first round and the blue pebble placed by Duplicator in the second round. Thus, Spoiler wins the 3-round MS game on $(\{\text{RT}(4)\}, \{\text{RT}(3)\})$ without playing on top. This completes the proof of the first part of the theorem.

To prove the second part of the theorem, we claim that Spoiler wins the 3-round MS game on $(\{\text{RT}(4)\}, \{\text{RT}(3)\})$ without playing on top only if Spoiler plays on the left in the first round, on the right in the second round, and on the left in the third round; moreover, Spoiler's round first move on $\text{RT}(4)$ must be on either B11 or B12. The argument about Spoiler's first-round move is similar to the one in the proof of Proposition 4.1.

We claim that if Spoiler's first-round move is on the right, then Duplicator can maintain a partial isomorphism between one copy of $\text{RT}(4)$ and one copy of $\text{RT}(3)$ for three rounds. Indeed, in response to a first-round move by Spoiler on L1 in $\text{RT}(3)$, Duplicator plays on B1 in one copy of $\text{RT}(4)$, and then easily survives two more rounds (e.g., in response to a second-round Spoiler play on B12, Duplicator plays on L11 in one copy and L12 in another copy of the $\text{RT}(3)$). First-round Spoiler plays on L11 or L21 are met by a Duplicator play on B21 in one copy of $\text{RT}(4)$, and first-round Spoiler plays on L12 or L22 are met by a Duplicator play on B22 in a copy of $\text{RT}(4)$; in each case, Duplicator readily survives two more rounds on just this pair of structures. If Spoiler's first-round move on $\text{RT}(4)$ is not on B11 or on B12, then once again, Duplicator survives three rounds just on $(\{\text{RT}(4)\}, \{\text{RT}(3)\})$ by maintaining a partial isomorphism between two copies via the following responses in a copy of $\text{RT}(3)$: $B1 \leftrightarrow L1$, $B13 \leftrightarrow L12$, $B21 \leftrightarrow L21$, $B22 \leftrightarrow L22$.

We have established that Spoiler wins only if his first-round move is on B12 or on B12 on the left. Using this, it is easy to verify that Spoiler must play his second-round moves on the right and his third-round moves on the left. Otherwise, Duplicator easily survives two more rounds. For example, if Spoiler plays B12 in the first round and B1 in the second round, then Duplicator plays L12 and L1 in the third copy of $\text{RT}(3)$, and then survives the third round.

So far, we have shown that Spoiler wins the 3-round MS game on $(\{\text{RT}(4)\}, \{\text{RT}(3)\})$ without playing on top only if Spoiler plays on the left in the first round, on the right in the second round, and on the left in the third round. Suppose now that there is a FO-sentence ψ with 3 quantifiers that is separating for $(\{\text{RT}(4)\}, \{\text{RT}(3)\})$ and such that both ψ and $\neg\psi$ are non-replicating. Corollary 4.7 implies that if Spoiler follows ψ , then Spoiler wins the 3-move MS game on $(\{\text{RT}(4)\}, \{\text{RT}(3)\})$ without playing on top. Therefore, ψ must have a quantifier prefix of the form $\exists\forall\exists$. Proposition 4.11, however, implies that at least one of the sentences ψ and $\neg\psi$ must be a non-replicating sentence, which is a contradiction. The proof of the theorem is now complete. \square

The proof of Theorem 4.10 shows that there are situations in which for Spoiler to win without ever playing on top, he must alternate between the left side and the right side more than once. In such cases, the separating sentence extracted from Spoiler's winning strategy has at least three alternations of quantifiers. The next proposition asserts that no sentence with three alternations of quantifiers has the property that both the sentence and its negation are non-replicating sentences.

Proposition 4.11. *Let $Q_1x_1 \dots Q_r x_r \theta$ be a FO-sentence with r variables such that $r \geq 3$ and θ is quantifier-free. If both ψ and $\neg\psi$ are non-replicating sentences, then neither ψ nor $\neg\psi$ contains the quantifiers $\exists x_i, \forall x_j, \exists x_k$ with $i < j < k$ in its prefix.*

Proof. Towards a contradiction and without loss of generality assume that ψ and $\neg\psi$ are non-replicating sentences such that the quantifier prefix of ψ contains the quantifiers $\exists x_i, \forall x_j, \exists x_k$ with $i < j < k$. The quantifier-free parts θ of ψ and $\neg\theta$ of $\neg\psi$ can be written as disjunctions of consistent types, so that each consistent type is a disjunct either of θ or of $\neg\theta$ (but not of both). Since ψ is a non-replicating sentence, no type occurring as a disjunct of θ can contain the equalities $x_i = x_k$ or $x_j = x_k$. Hence, at least one of the types $t(x_1, \dots, x_r)$ occurring as a disjunct of $\neg\theta$ must contain the equalities $x_i = x_k$ and $x_j = x_k$. Since $t(x_1, \dots, x_r)$ is a consistent type, it follows that $t(x_1, \dots, x_r)$ must also contain the equality $x_i = x_j$. Since x_i is universally quantified in $\neg\psi$ and x_j is existentially quantified in $\neg\psi$, it follows that $\neg\psi$ is a replicating sentence, contradicting the assumption made. \square

5. RESTRICTING THE NUMBER OF VARIABLES

Suppose in addition to the number of quantifiers, we also wish to simultaneously capture the number of *variables* needed to express a certain property, or distinguish two sets of τ -structures. How do we achieve this?

Let $\text{FO}^k(\tau)$ be the set of well-formed FO formulas over the schema τ , that only use the variables $\{x_1, \dots, x_k\}$.

5.1. Limitations of MS Games. It might at first seem reasonable to try to adapt the MS game (which already captures the number of quantifiers) to also capture the number of variables, simply by limiting the number of pebble colors that can be used. This approach works in a straightforward manner with EF games [Imm99, Definition 6.2, Theorem 6.10]. There are two natural ways to define such adaptations of the MS game.

Game 5.1 (MS Game with Repebbling). *Define the r -round, k -color MS game with repebbling on $(\mathcal{A}, \mathcal{B})$ as identical to the r -round MS game on $(\mathcal{A}, \mathcal{B})$, except that the set \mathcal{C} of pebble colors satisfies $|\mathcal{C}| = k$, and so, Spoiler needs to play with at most k pebble colors (forcing him to possibly re-use the same pebble color in two different rounds). When a pebble of a given color is re-used, the previously played pebble of the same color is “picked up” from all structures. In every round, Duplicator has to respond with the same pebble color as the one used by Spoiler in that round. As before, the winning conditions are the same, i.e., at the end of each of r rounds, Duplicator needs to exhibit a pebbled structure on the left and a pebbled structure on the right forming a matching pair, while Spoiler needs to exhibit a configuration within r rounds where no pair of pebbled structures from the left and right form a matching pair.*

While Game 5.1 seems like a reasonable candidate for capturing FO-distinguishability with r quantifiers and k variables, this turns out to not work. Recall that the discussion following Figure 1 observed that there is a 3-quantifier 2-variable sentence that is separating for $(\{\text{LO}(3)\}, \{\text{LO}(2)\})$, namely: $\exists x(\exists y(x < y) \wedge \exists y(y < x))$. So, if Game 5.1 were to capture distinguishability with r quantifiers and k variables, then Spoiler would need to have a winning strategy in the 3-round, 2-color MS game with repebbling on this instance. The following lemma, whose proof is straightforward and omitted, shows that this is not the case.

Lemma 5.2. *Duplicator has a winning strategy in the 3-round, 2-color MS game with repebbling on $(\{\text{LO}(3)\}, \{\text{LO}(2)\})$.*

In order to define the second variant of the MS game, we first need a definition.

Definition 5.3. During a complete play of the r -round, k -color MS game with repegbling on $(\mathcal{A}, \mathcal{B})$, we say that a pebbled structure $\langle \mathbf{S} \mid s_1, \dots, s_t \rangle$ is a *parent* of a pebbled structure $\langle \mathbf{S} \mid s'_1, \dots, s'_{t'} \rangle$, if both of the following statements hold:

- the configuration containing $\langle \mathbf{S} \mid s_1, \dots, s_t \rangle$ is from the round immediately preceding the configuration containing $\langle \mathbf{S} \mid s'_1, \dots, s'_{t'} \rangle$,
- $\langle \mathbf{S} \mid s'_1, \dots, s'_{t'} \rangle$ is the result of playing a new pebble color or reusing a pebble color on $\langle \mathbf{S} \mid s_1, \dots, s_t \rangle$.

Note that this means $t' = t$ (if a color was reused), or $t' = t + 1$ (if a new color was used).

Observe that in the MS game with repegbling, there may be $\mathbf{A}, \mathbf{A}' \in \mathcal{A}$ and $\mathbf{B}, \mathbf{B}' \in \mathcal{B}$ such that the following are both true for some $0 \leq t < r$:

- $\langle \mathbf{A}' \mid a'_1, \dots, a'_{t'} \rangle$ and $\langle \mathbf{B}' \mid b'_1, \dots, b'_{t'} \rangle$ are a matching pair in round $t + 1$.
- there are no matching pairs $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ and $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle$ in round t such that $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ is a parent of $\langle \mathbf{A}' \mid a'_1, \dots, a'_{t'} \rangle$, and $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle$ is a parent of $\langle \mathbf{B}' \mid b'_1, \dots, b'_{t'} \rangle$.

Our second variant of the MS game will constrain this heavily. First we need a definition.

Definition 5.4. During a complete play of the r -round, k -color MS game with repegbling on $(\mathcal{A}, \mathcal{B})$, two t -pebbled structures $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ on the left and $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle$ on the right form a *hereditary match* if the following hold:

- (a) $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ and $\langle \mathbf{B} \mid b_1, \dots, b_t \rangle$ are within the same configuration, say at the end of round t' .
- (b) There is a sequence $\mathbf{A}_0, \dots, \mathbf{A}_{t'}$ of pebbled structures with $\mathbf{A}_0 = \mathbf{A}$ and $\mathbf{A}_{t'} = \langle \mathbf{A} \mid a_1, \dots, a_t \rangle$, such that \mathbf{A}_i is the parent of \mathbf{A}_{i+1} for each $0 \leq i < t'$.
- (c) There is a sequence $\mathbf{B}_0, \dots, \mathbf{B}_{t'}$ of pebbled structures with $\mathbf{B}_0 = \mathbf{B}$ and $\mathbf{B}_{t'} = \langle \mathbf{B} \mid b_1, \dots, b_t \rangle$, such that \mathbf{B}_i is the parent of \mathbf{B}_{i+1} for each $0 \leq i < t'$.
- (d) \mathbf{A}_i and \mathbf{B}_i form a matching pair, for $0 \leq i \leq t'$.

Game 5.5 (Hereditary MS Game with Repegbling). *Define the r -round, k -color hereditary MS game with repegbling on $(\mathcal{A}, \mathcal{B})$ as identical to the r -round, k -color MS game with repegbling on $(\mathcal{A}, \mathcal{B})$, except that, at the end of each of r rounds, Duplicator needs to exhibit a pebbled structure on the left and a pebbled structure on the right forming a hereditary match (Definition 5.4), while Spoiler needs to exhibit a configuration within r rounds where no pebbled structure on the left forms a hereditary match with any pebbled structure on the right.*

Note that a Duplicator win in round r of Game 5.5 immediately certifies a sequence of matching pairs that are valid for each of the previous rounds, and so if Duplicator wins Game 5.5 on an instance $(\mathcal{A}, \mathcal{B})$, she certainly wins Game 5.1 on $(\mathcal{A}, \mathcal{B})$. We also note that when $k = r$, both Games 5.1 and 5.5 are identical to ordinary MS games. Furthermore, it is straightforward to see that Duplicator's oblivious strategy is optimal for both these variants.

We observe that the instance from Lemma 5.2 does not work as a counterexample for Game 5.5, as Spoiler does win the 3-round, 2-color hereditary MS game with repegbling on $(\{\text{LO}(3)\}, \{\text{LO}(2)\})$. We leave this to the reader to verify.

However, it turns out that this stronger candidate game also fails to simultaneously capture the number of quantifiers and number of variables for FO-distinguishability. This will follow from Proposition 5.8, but we will first need to set up stronger methods in order to prove such a result.

We now introduce a game that *does* simultaneously capture these two parameters.

5.2. The Quantifier-Variable Tree Game. We define the (r, k) -*quantifier-variable tree (QVT) game*, denoted the $\mathcal{QVT}(r, k)$ game, on $(\mathcal{A}, \mathcal{B})$ as follows.

Two players, Spoiler and Duplicator, play by growing a *game tree* \mathcal{T} , starting from a single root node X_{root} . Once again, we have a set \mathcal{C} of *pebble colors*, with $|\mathcal{C}| = k$, and arbitrarily many pebbles available of each color. Throughout the rest of this section, name the colors in \mathcal{C} as x_1, \dots, x_k . The color x_i will turn out to correspond to the variable x_i , and so we use this ambiguous notation rather deliberately.

We consider the leaf nodes in \mathcal{T} to be *open* or *closed*, with the root X_{root} being considered an open leaf at the start of the game. We say \mathcal{T} is *closed* if all its leaves are closed.

Each node of \mathcal{T} consists of a *left* side, a *right* side, and a *counter* $r' \in \mathbb{N}$. Each side consists of a set of pebbled τ -structures. We denote a node in \mathcal{T} as a tuple $\langle (\text{left}, \text{right}) \mid r' \rangle$.

At the root node $X_{\text{root}} \in V(\mathcal{T})$,

- the left side consists of \mathcal{A} , viewed as pebbled τ -structures (with no pebbles placed yet).
- the right side consists of \mathcal{B} , viewed as pebbled τ -structures (with no pebbles placed yet).
- the counter r' is set to r .

The root node, therefore, is denoted $X_{\text{root}} = \langle (\mathcal{A}, \mathcal{B}) \mid r \rangle$.

The contents $\langle (\text{left}, \text{right}) \mid r' \rangle$ of the node $X \in V(\mathcal{T})$ will correspond to a *configuration* of the QVT game. When the context is clear, we will often identify a \mathcal{T} -node by its configuration. Throughout the tree \mathcal{T} , we will maintain the invariant that on every node $X \in V(\mathcal{T})$, a pebble color in \mathcal{C} appears in one pebbled structure iff it appears in every pebbled structure throughout the configuration.

Spoiler on his turn can perform any of the following moves:

- **Pebble-Left :**
 - (a) Spoiler chooses an open leaf node $X = \langle (\mathcal{A}', \mathcal{B}') \mid r' \rangle$ with $r' \geq 1$, and a pebble color $x_i \in \mathcal{C}$. If all structures in X contain the pebble color x_i , Spoiler removes all of these pebbles.
 - (b) For each pebbled structure $\mathbf{A} \in \mathcal{A}'$, Spoiler places a pebble colored x_i on an element in the universe of \mathbf{A} . Call this new set of pebbled structures \mathcal{A}'' .
 - (c) For each pebbled structure $\mathbf{B} \in \mathcal{B}'$, Duplicator *may* make any number of copies of \mathbf{B} , then *must* place a pebble colored x_i on an element in the universe of \mathbf{B} and an element in the universe of each copy. Call this new set of pebbled structures \mathcal{B}'' .
 - (d) Spoiler makes a new open leaf $X' = \langle (\mathcal{A}'', \mathcal{B}'') \mid r' - 1 \rangle$ in \mathcal{T} , with parent X . Note that X is no longer a leaf in \mathcal{T} .
- **Pebble-Right:** This move is dual to **Pebble-Left**; Spoiler plays on \mathcal{B} , and Duplicator responds on \mathcal{A} .
- **Split-Left:**
 - (a) Spoiler chooses an open leaf node $X = \langle (\mathcal{A}', \mathcal{B}') \mid r' \rangle$ and $r'_1, r'_2 \in \mathbb{N}$ such that $r' = r'_1 + r'_2$.
 - (b) Spoiler partitions \mathcal{A}' so that $\mathcal{A}' = \mathcal{A}'_1 \cup \mathcal{A}'_2$.
 - (c) Spoiler makes two new open leaf nodes $X_1 = \langle (\mathcal{A}'_1, \mathcal{B}') \mid r'_1 \rangle$ and $X_2 = \langle (\mathcal{A}'_2, \mathcal{B}') \mid r'_2 \rangle$ in \mathcal{T} . Both new nodes have parent X .
- **Split-Right:** This move is dual to **Split-Left**; Spoiler partitions \mathcal{B}' .
- **Swap:** Spoiler chooses an open leaf node $X = \langle (\mathcal{A}', \mathcal{B}') \mid r' \rangle$ in \mathcal{T} . He makes a new open leaf node $X' = \langle (\mathcal{B}', \mathcal{A}') \mid r' \rangle$ in \mathcal{T} , with parent X .

- **Close:** Spoiler chooses an open leaf node X in \mathcal{T} , say $\langle (\mathcal{A}', \mathcal{B}') \mid r' \rangle$. Suppose there is an atomic formula $\varphi \in \text{FO}^k$ with variables $\{x_{i_1}, \dots, x_{i_m}\}$, such that:
 - (a) Every pebbled structure $\langle \mathbf{A} \mid a_1, \dots, a_\ell \rangle$ in \mathcal{A} has elements $a_{i_1}, \dots, a_{i_m} \in A$ pebbled by colors x_{i_1}, \dots, x_{i_m} respectively, and $\mathbf{A} \models \varphi(x_{i_1}/a_{i_1}, \dots, x_{i_m}/a_{i_m})$.
 - (b) Every pebbled structure $\langle \mathbf{B} \mid b_1, \dots, b_\ell \rangle$ in \mathcal{B} has elements $b_{i_1}, \dots, b_{i_m} \in B$ pebbled by colors x_{i_1}, \dots, x_{i_m} respectively, and $\mathbf{B} \models \neg\varphi(x_{i_1}/b_{i_1}, \dots, x_{i_m}/b_{i_m})$.
 Then, Spoiler can mark X *closed*.

We say that \mathcal{T} is *closed* if there are no open leaf nodes. Spoiler wins the game if he closes \mathcal{T} , and Duplicator wins otherwise.

The following theorem characterizes the expressive power of the QVT game, showing that it does indeed capture simultaneous bounds on number of quantifiers and number of variables. The key intuition behind the proof is that a closed game tree \mathcal{T} is by design isomorphic to the parse tree for a separating sentence for $(\mathcal{A}, \mathcal{B})$, with the moves **Pebble-Left**, **Pebble-Right**, **Split-Left**, **Split-Right**, and **Swap** in the game tree corresponding respectively to $\exists, \forall, \vee, \wedge$, and \neg in the separating sentence. We omit the proof, as it will follow immediately from the more general Theorem 6.7 and Example 6.2 in Section 6.

Theorem 5.6 (Equivalence theorem for QVT). *Spoiler has a winning strategy for the QVT(r, k) game on $(\mathcal{A}, \mathcal{B})$ iff there is an r -quantifier, k -variable sentence over τ that is separating for $(\mathcal{A}, \mathcal{B})$.*

For ease of arguments, we will henceforth refer to **Pebble-Left** or **Pebble-Right** moves as *pebble moves*, and **Split-Left** or **Split-Right** moves as *split moves*. We will say an internal vertex of \mathcal{T} is *closed* when all \mathcal{T} -leaves descended from it are closed. Note that the value of the counter r' at each \mathcal{T} -node effectively represents the “budget” on the number of pebble moves available to Spoiler to close that node.

It is not entirely obvious that Duplicator’s optimal strategy in the QVT game is also to play the oblivious strategy. The following proposition asserts that this is still true. The proof, once again, is omitted, as it will follow from the more general Proposition 6.8.

Proposition 5.7. *If Duplicator has a winning strategy in the QVT game, then the oblivious strategy is a winning strategy.*

Figure 7 depicts a complete game tree \mathcal{T} for the (3, 2)-QVT game on $(\text{LO}(3), \text{LO}(2))$. Duplicator plays her oblivious strategy, but Spoiler is able to close the tree. Note that this is the same instance as both Figure 1 and Lemma 5.2.

5.3. The hereditary MS game with repebbling does not capture number of variables. We are now ready to prove that the r -round, k -color hereditary MS game with repebbling (Game 5.5) does *not* capture the number of variables and quantifiers simultaneously, using Theorem 5.6.

Proposition 5.8. *The following are both true:*

- (a) *There is no 3-quantifier, 2-variable separating sentence for $(\text{LO}(4), \text{LO}(3))$.*
- (b) *Spoiler has a winning strategy in the 3-round, 2-color hereditary MS game with repebbling on $\{\{\text{LO}(4)\}, \{\text{LO}(3)\}\}$.*

Proof. We first show (a), i.e., there is no 3-quantifier, 2-variable separating sentence for $(\text{LO}(4), \text{LO}(3))$. By Theorem 5.6, it suffices to show that Duplicator has a winning strategy

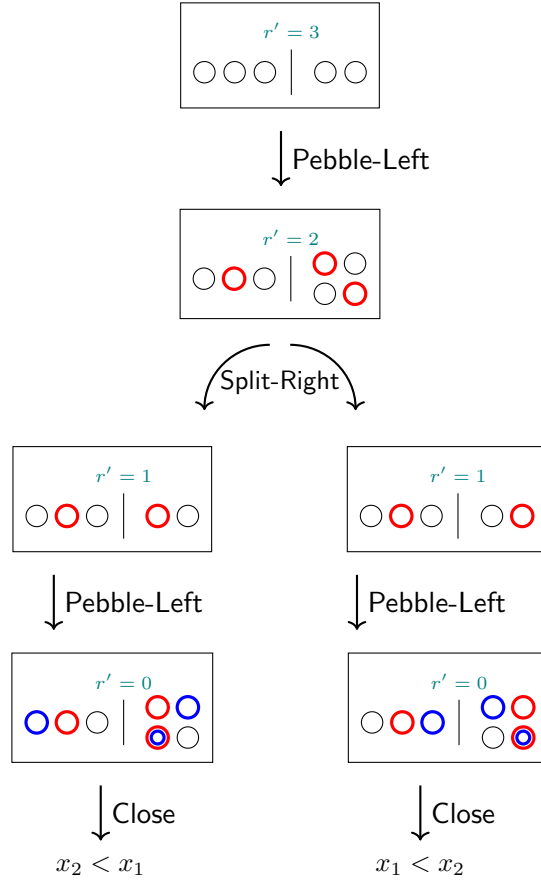


Figure 7: A complete play of the $QVT(3, 2)$ game on $(\{LO(3)\}, \{LO(2)\})$. Spoiler closes the game tree \mathcal{T} using only three pebble moves, and two pebble colors x_1 (red) and x_2 (blue). The atomic separating sentences are shown below, and the resulting separating sentence is $\exists x_1(\exists x_2(x_2 < x_1) \wedge \exists x_2(x_1 < x_2))$, which can be read off \mathcal{T} .

on the game $QVT(3, 2)$ game on $(\{LO(4)\}, \{LO(3)\})$. Since we are starting with singleton sets, Spoiler cannot play a split move, and gains nothing from a **Swap** move at the root node of \mathcal{T} . Let us consider each of the possible pebble moves one by one. Note that the counter will decrement after this move from 3 to 2.

It is straightforward to check that, at the root node of \mathcal{T} , a **Pebble-Right** play on L1 or L3 is met easily with a Duplicator response on B1 or B4 respectively. This results in a Duplicator victory in the 3-round MS game, and therefore in $QVT(3, 3)$ game on $(\{LO(4)\}, \{LO(3)\})$, and therefore also in $QVT(3, 2)$ game on $(\{LO(4)\}, \{LO(3)\})$. So assume Spoiler plays **Pebble-Right** on L2 at the root of \mathcal{T} . By Proposition 5.7, we can assume Duplicator responds obliviously, and we are again at the situation depicted in Figure 10. At this point, a Spoiler split move can be ruled out, as the \mathcal{T} -node resulting from the split containing the second structure on the left will require two more pebble moves to close, while the other \mathcal{T} -node from the split will require at least one more pebble move to close. Spoiler therefore has to

use up his second pebble move at this point, and we examine in turn the two possible pebble moves at this \mathcal{T} -node. After this move, the counter will decrement to 1.

If for his second pebble move, Spoiler plays **Pebble-Right** on L1, then Duplicator can respond on B1 on the second board on the left, and survive another pebble move on just this pair of structures. A symmetric argument applies when Spoiler plays **Pebble-Right** on L3. Finally, it is easy to see that Spoiler playing **Pebble-Right** on L2 achieves nothing. So Spoiler must play **Pebble-Left** for his second pebble move.

Let us now examine this second pebble move by Spoiler on the second structure on the left. Where does Spoiler place a pebble? If he plays on B1 or B2, Duplicator responds with L1 or L2 respectively, and easily survives another pebble move. So, Spoiler must play either B3 or B4, which must be met with a Duplicator response on L3. In fact, for Spoiler to win on this pair of pebbled structures with one more pebble move, he must play on B3 (not B4) for his second pebble move. A symmetric argument shows that on the third structure on the left, Spoiler must play on B2 in his second pebble move. We can now look at Figure 8, where we have only depicted four relevant pebbled structures after two pebble moves.




Figure 8: Configuration (partial) after two pebble moves, depicting just two pairs of pebbled structures that maintain isomorphisms.

At this point, splitting is not an option, as the counter is at 1, and so splitting would cause one of the two resultant \mathcal{T} -nodes to have the counter value at 0, and Spoiler cannot close that node. But Spoiler cannot win without splitting either, as moving any pebble on any of the four structures in Figure 8 can be met with a response on at least one of the two structures on the other side that maintains an isomorphism. It follows that Spoiler cannot win with a **Pebble-Right** move at the root of \mathcal{T} .

It is straightforward to check that, at the root of \mathcal{T} , a **Pebble-Left** play on B1 or B4 is met easily with a Duplicator response on L1 or L3 respectively. So assume Spoiler plays **Pebble-Left** on B2 for his first pebble move (B3 will be symmetric). By Proposition 5.7, we can assume Duplicator responds obliviously, as depicted in Figure 9. The counter is at 2.




Figure 9: Configuration of the game $QVT(3, 2)$ game on $(\{LO(4)\}, \{LO(3)\})$ after the first pebble move, where Spoiler plays **Pebble-Left** on B2 and Duplicator responds obliviously.

By a similar argument as before, we can conclude that a Spoiler split move can be ruled out. We can now argue that in order to win in three pebble moves, Spoiler must play

Pebble-Right for his second pebble move. Indeed, if instead, Spoiler plays Pebble-Left, then a play of

- B1 is met with Duplicator playing on L1 on the middle structure on the right and surviving one more pebble move on that pair;
- B3 is met with Duplicator playing L2 on the top structure on the right and L3 on the middle structure on the right, and surviving one more pebble move on one of these structures;
- B4 is met with Duplicator playing L3 on both of the top two structures on the right, and surviving one more pebble move on one of these structures.

We now claim that for Spoiler’s second pebble move with the Pebble-Right, he must play on L3 on the middle structure on the right. Otherwise, a play on L1 or L2 is met by a Duplicator move on B1 or B2 respectively on the structure on the left, where she easily survives one more pebble move just on this pair. However, a Spoiler play of L3 is met with a Duplicator play on B4 on the left, and she can survive one more pebble move on this pair. This concludes the proof of (a).

We next show (b), i.e., Spoiler wins the 3-round, 2-color hereditary MS game with repegging on $\{\{LO(4)\}, \{LO(3)\}\}$. Spoiler’s first round Play-Right move, with Duplicator’s oblivious responses on the left side, is shown in Figure 10.




Figure 10: Configuration after round 1 of the 3-round, 2-color hereditary MS game with repegging on $\{\{LO(4)\}, \{LO(3)\}\}$.

Spoiler will play Play-Left for the two remaining rounds. His second round plays are indicated in blue in Figure 11.




Figure 11: Configuration (partial) after Spoiler’s second round Play-Left move (given in blue) for the 3-round, 2-color hereditary MS game with repegging on $\{\{LO(4)\}, \{LO(3)\}\}$.

In response, let us WLOG only consider Duplicator’s plays on L1 and L3 (the L2 move achieves nothing). Responding with L1 on the right keeps an isomorphism going with pebbled structures 3 and 4 on the left, and responding with L3 on the right keeps an isomorphism going with pebbled structures 1 and 2 on the left. In response, in the third round, Spoiler plays the red pebble on B3, B4, B1, and B2 on pebbled structures 1, 2, 3, and 4 respectively, in all cases breaking all hereditary isomorphisms with pebbled structures on the right. This concludes the proof. \square

It is worth noting that we could *not* have proven Proposition 5.8 via a straightforward pebble game argument with 3 rounds and 2 pebbles. Indeed, Spoiler easily wins such a game beginning with a play on B2, B3 or L2. This implies that there is a rank-3, 2-variable separating sentence for $(\{\text{LO}(4)\}, \{\text{LO}(3)\})$, namely, $\exists x((\exists y(y < x)) \wedge \exists y(x < y \wedge \exists x(y < x)))$.

6. THE SYNTACTIC GAME

Consider a node $X = \langle (\mathcal{A}', \mathcal{B}') \mid r' \rangle$ in the game tree \mathcal{T} for a given QVT game. The counter $r' \in \mathbb{N}$ in X is the number of pebble moves Spoiler may use to close the subtree of \mathcal{T} rooted at X . The proof of Theorem 5.6 will follow from the fact that this is exactly the number of quantifiers needed in some sense to “distinguish” between \mathcal{A}' and \mathcal{B}' . In fact, as stated earlier, we will prove a significantly stronger generalization. We first need the formal notion of a *syntactic measure*. Such measures were introduced to derive combinatorial games for “reasonable” complexity measure of infinitary formulas in the logic $\mathcal{L}_{\omega_1, \omega}$ [VW13, Definition 5.1]. In this section, we will derive games precisely capturing many familiar complexity measures, including: quantifier count, quantifier rank, and formula size. The main idea in all of these games would be to have specific types of *pebble moves* corresponding to the inductive ways to build up FO formulas; Spoiler will use these moves to explicitly simulate a separating formula, and will try to do this and close the game tree without running out of a “budget” on the syntactic measure, which we keep track of by means of a counter. In these moves, the number of pebble colors captures the number of variables. For each logical connective occurring in the language in question, there is a move and a corresponding cost of using that operator. For example, to count the quantifier number, we add 1 for each quantifier used, as in the QVT game. Thus, these games capture the descriptive cost function under consideration. They are the natural tool for proving descriptive lower bounds; sometimes they are the best choice for proving upper bounds as well.

6.1. Compositional syntactic measures. Recall that $\text{FO}^k(\tau)$ is the set of well-formed FO formulas over the schema τ , that use only the variables $\{x_1, \dots, x_k\}$. Note that $\text{FO}^k(\tau)$ is defined inductively, where each formula is the application of a logical symbol to some subformulas. A *compositional syntactic measure* uses such an inductive definition to measure formula complexity.

Recall that $\text{Pred}(\tau)$ is the set of predicate symbols of τ , assumed to be finite.

Definition 6.1. A *compositional syntactic measure* is a function $f : \text{FO}^k(\tau) \rightarrow \mathbb{N}$ defined in terms of helper functions $h_{\neg}, h_{\exists}, h_{\forall} : \mathbb{N} \rightarrow \mathbb{N}$, $h_{\vee}, h_{\wedge} : \mathbb{N}^2 \rightarrow \mathbb{N}$, $h_{\text{atomic}} : \text{Pred}(\tau) \rightarrow \mathbb{N}$, as follows:

- if φ is atomic, say $\varphi = P(t_1, \dots, t_r)$ for some $P \in \text{Pred}(\tau)$, and terms t_1, \dots, t_r , where r is the arity of P , then $f(\varphi) = h_{\text{atomic}}(P)$.

- if $\varphi = \neg\psi$, then $f(\varphi) = h_{\neg}(f(\psi))$.
- if $\varphi = (\psi \vee \gamma)$, then $f(\varphi) = h_{\vee}(f(\psi), f(\gamma))$.
- if $\varphi = (\psi \wedge \gamma)$, then $f(\varphi) = h_{\wedge}(f(\psi), f(\gamma))$.
- if $\varphi = \exists x(\psi)$, then $f(\varphi) = h_{\exists}(f(\psi))$.
- if $\varphi = \forall x(\psi)$, then $f(\varphi) = h_{\forall}(f(\psi))$.

We call such a measure “compositional” because the value of f is determined by the value of f on its subformulas. Compositional syntactic measures are a large and natural class of syntactic measures, which includes quantifier rank, quantifier number, formula size, and number of atomic formulas. However, we remark that the number of distinct variables is *not* a compositional measure, as the value of this measure on, say, a disjunction depends not only on the values on the subformulas, but actually on the subformulas themselves. We also note that in Definition 6.1, we treat equality as an ordinary binary predicate.

We could have stated Definition 6.1 purely in terms of well-formed formulas over τ , instead of the slice $\text{FO}^k(\tau)$. However, we chose our convention to evoke the idea of fixing k variables at the outset, which will be relevant in the rest of this section.

The following examples show that quantifier number, quantifier rank, and formula size are indeed compositional syntactic measures.

Example 6.2. The *quantifier count* f_q satisfies:

- if φ is atomic, $f_q(\varphi) = 0$. In other words, $h_{\text{atomic}}(\varphi) = 0$.
- if $\varphi = \neg\psi$, $f_q(\varphi) = f_q(\psi)$. In other words, $h_{\neg}(n) = n$.
- if $\varphi = (\psi \wedge \gamma)$ or $\varphi = (\psi \vee \gamma)$, $f_q(\varphi) = f_q(\psi) + f_q(\gamma)$. In other words, $h_{\vee}(m, n) = h_{\wedge}(m, n) = m + n$.
- if $\varphi = \exists x\psi$ or $\varphi = \forall x\psi$, $f_q(\varphi) = 1 + f_q(\psi)$. In other words, $h_{\exists}(n) = h_{\forall}(n) = n + 1$.

Example 6.3. The *quantifier rank* f_r satisfies:

- if φ is atomic, $f_r(\varphi) = 0$. In other words, $h_{\text{atomic}}(\varphi) = 0$.
- if $\varphi = \neg\psi$, $f_r(\varphi) = f_r(\psi)$. In other words, $h_{\neg}(n) = n$.
- if $\varphi = (\psi \wedge \gamma)$ or $\varphi = (\psi \vee \gamma)$, $f_r(\varphi) = \max(f_r(\psi), f_r(\gamma))$. In other words, $h_{\vee}(m, n) = h_{\wedge}(m, n) = \max(m, n)$.
- if $\varphi = \exists x\psi$ or $\varphi = \forall x\psi$, $f_r(\varphi) = 1 + f_r(\psi)$. In other words, $h_{\exists}(n) = h_{\forall}(n) = n + 1$.

Example 6.4. The *formula size* f_s satisfies:

- if φ is atomic, $f_s(\varphi) = 1$. In other words, $h_{\text{atomic}}(\varphi) = 1$.
- if $\varphi = \neg\psi$, $f_s(\varphi) = 1 + f_s(\psi)$. In other words, $h_{\neg}(n) = n + 1$.
- if $\varphi = (\psi \wedge \gamma)$ or $\varphi = (\psi \vee \gamma)$, $f_s(\varphi) = 1 + f_s(\psi) + f_s(\gamma)$. In other words, $h_{\vee}(m, n) = h_{\wedge}(m, n) = 1 + m + n$.
- if $\varphi = \exists x\psi$ or $\varphi = \forall x\psi$, $f_s(\varphi) = 1 + f_s(\psi)$. In other words, $h_{\exists}(n) = h_{\forall}(n) = n + 1$.

We are now ready to define the syntactic game.

6.2. The syntactic game. Fix a τ -structure \mathbf{A} , and let $\alpha_{\mathbf{A}}$ be a *partial assignment*, i.e., a function that maps a set of variables to elements in A . We call the ordered pair $(\mathbf{A}, \alpha_{\mathbf{A}})$ a *structure-assignment pair*. Note that if we identify pebble colors with variables, a pebbled τ -structure $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ has a one-to-one correspondence with the pair $(\mathbf{A}, \alpha_{\mathbf{A}})$, where $\alpha_{\mathbf{A}}$ maps the variable corresponding to a pebble color to the element in A that it is placed on. In particular, a τ -structure \mathbf{A} has a one-to-one correspondence with the pair (\mathbf{A}, \emptyset) , where \emptyset is the empty assignment.

Just as a τ -structure \mathbf{A} satisfies a sentence over τ , we can extend this notion to a structure-assignment pair $(\mathbf{A}, \alpha_{\mathbf{A}})$ satisfying a *formula*, which is allowed to have free variables as long as all such variables are in the domain² of $\alpha_{\mathbf{A}}$. Formally, let $(\mathbf{A}, \alpha_{\mathbf{A}})$ be a structure-assignment pair, and ψ be a formula such that $\text{Free}(\psi) = \{x_{i_1}, \dots, x_{i_m}\} \subseteq \text{dom}(\alpha_{\mathbf{A}})$. Suppose $\alpha_{\mathbf{A}}(x_{i_j}) = a_{i_j} \in A$ for all $1 \leq j \leq m$. We say that $(\mathbf{A}, \alpha_{\mathbf{A}})$ *satisfies* ψ (or $(\mathbf{A}, \alpha_{\mathbf{A}}) \models \psi$) if $\mathbf{A} \models \psi(x_{i_1}/a_{i_1}, \dots, x_{i_m}/a_{i_m})$.

Of course, for every structure-assignment pair $(\mathbf{A}, \alpha_{\mathbf{A}})$ and formula ψ with $\text{Free}(\psi) \subseteq \text{dom}(\alpha_{\mathbf{A}})$, exactly one of $(\mathbf{A}, \alpha_{\mathbf{A}}) \models \psi$ and $(\mathbf{A}, \alpha_{\mathbf{A}}) \models \neg\psi$ is true.

For the rest of this paper, by convention, we will denote sets of structure-assignment pairs with script typeface \mathcal{A}, \mathcal{B} , and so on. Note that the τ -structures among the pairs in any such \mathcal{A} can in general be different. We now state a definition for convenience.

Definition 6.5. Two sets \mathcal{A} and \mathcal{B} of structure-assignment pairs are called *domain-consistent* if there is a set W of variables such that for every structure-assignment pair $(\mathbf{X}, \alpha_{\mathbf{X}}) \in \mathcal{A} \cup \mathcal{B}$, we have that $\text{dom}(\alpha_{\mathbf{X}}) = W$. We denote this common domain W by $\text{dom}(\mathcal{A})$ (or, equivalently, $\text{dom}(\mathcal{B})$).

In other words, every partial assignment across all pairs in the two sets have the same domain. In particular, every structure-assignment pair within the set \mathcal{A} (or, equivalently, \mathcal{B}) also has the same domain for its partial assignment.

We can now formalize what it means for a formula with free variables to be a separating formula, thereby generalizing Definition 2.2.

Definition 6.6. Let \mathcal{A} and \mathcal{B} be domain-consistent sets of structure-assignment pairs with common domain $\text{dom}(\mathcal{A})$. A *separating formula for* $(\mathcal{A}, \mathcal{B})$ is a FO-formula ψ with $\text{Free}(\psi) \subseteq \text{dom}(\mathcal{A})$, such that every $(\mathbf{A}, \alpha_{\mathbf{A}}) \in \mathcal{A}$ satisfies $(\mathbf{A}, \alpha_{\mathbf{A}}) \models \psi$, and every $(\mathbf{B}, \alpha_{\mathbf{B}}) \in \mathcal{B}$ satisfies $(\mathbf{B}, \alpha_{\mathbf{B}}) \models \neg\psi$.

Note that this does generalize Definition 2.2, using the correspondence between \mathbf{A} and (\mathbf{A}, \emptyset) . A separating sentence is just a separating formula with no free variables.

We are now ready to define the syntactic game. Fix two domain-consistent sets \mathcal{A} and \mathcal{B} of structure-assignment pairs. For a compositional syntactic measure f , define the (r, k) -*syntactic game on* f , denoted $\mathcal{SG}(r, k; f)$, on $(\mathcal{A}, \mathcal{B})$ as follows.

Two players, Spoiler and Duplicator, play by growing a *game tree* \mathcal{T} , starting from a single root node X_{root} . As in the QVT game, we have a set $\mathcal{C} = \{x_1, \dots, x_k\}$ of pebble colors, with arbitrarily many pebbles available of each color. Once again, the color x_i will correspond to the variable x_i . The leaf nodes in \mathcal{T} are open or closed, with the root X_{root} considered to be an open leaf at the start of the game. As before, we define \mathcal{T} to be *closed* if all its leaves are closed. The notion of an internal vertex being closed also carries over from Section 5. The nodes $X \in V(\mathcal{T})$ will correspond to a tuple $\langle (\text{left}, \text{right}) \mid c \rangle$ as before. We define a configuration of the game as before, and maintain the same invariant as in the QVT game. Recall that a pebbled τ -structure $\langle \mathbf{A} \mid a_1, \dots, a_t \rangle$ now has a one-to-one correspondence with the pair $(\mathbf{A}, \alpha_{\mathbf{A}})$, where $\alpha_{\mathbf{A}}$ is the assignment that maps a variable corresponding to a pebble color to the element in A it is placed on. To avoid notational clutter, we will view \mathcal{A} or \mathcal{B} as either a set of pebbled τ -structures or as a set of structure-assignment pairs interchangeably, where the precise one will be clear from context.

At the root node $X_{\text{root}} \in V(\mathcal{T})$,

²Note that throughout this section, the term *domain* will always refer to the domain of the assignment functions, and never to the universe of τ -structures.

- the left side consists of \mathcal{A} , viewed as pebbled τ -structures (with the pebbling inherited from the assignment functions).
- the right side consists of \mathcal{B} , viewed as pebbled τ -structures (with the pebbling inherited from the assignment functions).
- the counter r' is set to r .

The root node, therefore, is denoted $X_{\text{root}} = \langle (\mathcal{A}, \mathcal{B}) \mid r \rangle$. Note that since \mathcal{A} and \mathcal{B} are domain-consistent, every pebbled τ -structure at the root has the same set of pebbles on it. We will maintain this invariant throughout every node of \mathcal{T} (i.e., every configuration will correspond to a pair of domain-consistent sets \mathcal{A}' and \mathcal{B}'). As before, the root is considered an open leaf node at the start.

Once again, Spoiler on his turn can perform any of the following moves on an open leaf node $X = \langle (\mathcal{A}', \mathcal{B}') \mid r' \rangle$, as long as the precondition for that move is met:

- **Pebble-Left** : precondition: $h_{\exists}^{-1}(r') \neq \emptyset$
 - (a) Spoiler chooses a pebble color $x_i \in \mathcal{C}$. If all structures in X contain the pebble color x_i , Spoiler removes all of these pebbles.
 - (b) For each pebbled structure $(\mathbf{A}, \alpha_{\mathbf{A}}) \in \mathcal{A}'$, Spoiler places a pebble colored x_i on an element in the universe of \mathbf{A} . Call this new set of pebbled structures \mathcal{A}'' .
 - (c) For each pebbled structures $(\mathbf{B}, \alpha_{\mathbf{B}}) \in \mathcal{B}'$, Duplicator *may* make any number of copies of $(\mathbf{B}, \alpha_{\mathbf{B}})$, then *must* place a pebble colored x_i on an element in the universe of \mathbf{B} and an element in the universe of each copy. Call this new set of pebbled structures \mathcal{B}'' .
 - (d) Spoiler chooses some $r'' \in \mathbb{N}$ with $h_{\exists}(r'') = r'$. Note that this is guaranteed to exist by the precondition.
 - (e) Spoiler makes a new open leaf $X' = \langle (\mathcal{A}'', \mathcal{B}'') \mid r'' \rangle$ in \mathcal{T} , with parent X . Note that X is no longer a leaf in \mathcal{T} .
- **Pebble-Right** : precondition: $h_{\forall}^{-1}(r') \neq \emptyset$

This move is dual to **Pebble-Left**; Spoiler plays on \mathcal{B}' , Duplicator responds on \mathcal{A}' , and Spoiler chooses some $r'' \in \mathbb{N}$ with $h_{\forall}(r'') = r'$.
- **Split-Left**: precondition: $h_{\vee}^{-1}(r') \neq \emptyset$
 - (a) Spoiler partitions \mathcal{A}' so that $\mathcal{A}' = \mathcal{A}'_1 \cup \mathcal{A}'_2$.
 - (b) Spoiler chooses $r'_1, r'_2 \in \mathbb{N}$ such that $h_{\vee}(r'_1, r'_2) = r'$. Note that these are guaranteed to exist by the precondition.
 - (c) Spoiler makes two new open leaf nodes $X_1 = \langle (\mathcal{A}'_1, \mathcal{B}') \mid r'_1 \rangle$ and $X_2 = \langle (\mathcal{A}'_2, \mathcal{B}') \mid r'_2 \rangle$ in \mathcal{T} . Both new nodes have parent X .
- **Split-Right**: precondition: $h_{\wedge}^{-1}(r') \neq \emptyset$

This move is dual to **Split-Left**; Spoiler partitions \mathcal{B}' and chooses $r'_1, r'_2 \in \mathbb{N}$ such that $h_{\wedge}(r'_1, r'_2) = r'$.
- **Swap**: precondition: $h_{\neg}^{-1}(r') \neq \emptyset$.
 - (a) Spoiler chooses $r'' \in \mathbb{N}$ such that $h_{\neg}(r'') = r'$.
 - (b) Spoiler makes a new open leaf node $X' = \langle (\mathcal{B}', \mathcal{A}') \mid r'' \rangle$ in \mathcal{T} , with parent X .
- **Close**: If there is an atomic formula φ over τ that is separating (Definition 6.6) for $(\mathcal{A}', \mathcal{B}')$, such that $h_{\text{atomic}}(\varphi) = r'$, Spoiler can mark X *closed*.

We say that \mathcal{T} is *closed* if there are no open leaf nodes. Spoiler wins the game if he closes \mathcal{T} , and Duplicator wins otherwise.

If \mathcal{A} and \mathcal{B} are sets of τ -structures, we will abuse notation slightly and talk about the syntactic game on $(\mathcal{A}, \mathcal{B})$, where these will be viewed as sets of structure-assignment pairs with the empty assignment throughout.

We are now ready to characterize the expressive power of the syntactic game. The key to the proof is the observation that a closed game tree \mathcal{T} is isomorphic to the parse tree of a formula separating $(\mathcal{A}, \mathcal{B})$.

Theorem 6.7 (Syntactic Game Fundamental Theorem). *Fix a compositional syntactic measure f , and domain-consistent sets \mathcal{A}, \mathcal{B} of structure-assignment pairs, with $\text{dom}(\mathcal{A}) \subseteq \{x_1, \dots, x_k\}$. Spoiler has a winning strategy for the game $\mathcal{SG}(r, k; f)$ on $(\mathcal{A}, \mathcal{B})$ if and only if there is a formula $\varphi \in \text{FO}^k(\tau)$ with $f(\varphi) = r$, that is separating for $(\mathcal{A}, \mathcal{B})$.*

Proof.

(\implies .) Suppose Spoiler closes \mathcal{T} even with optimal play from Duplicator. We use induction on the closed game tree \mathcal{T} to prove that there is a separating formula. If \mathcal{T} consists of a single node, then Spoiler closed the root node for his first move, and this corresponds to an atomic formula with the specified counter value by definition of the Close move. Otherwise, consider a closed tree \mathcal{T} , and consider Spoiler's first move on the root. Inductively, each of the subtrees corresponds to a separating formula. We consider each possible first move by Spoiler. Let the root node be $X_{\text{root}} = \langle (\mathcal{A}, \mathcal{B}) \mid r \rangle$.

- Suppose Spoiler plays **Swap**, creating $X = \langle (\mathcal{B}, \mathcal{A}) \mid r' \rangle$. By induction, $(\mathcal{B}, \mathcal{A})$ is separable by a formula ψ such that $f(\psi) = r'$. Consider the formula $\varphi = \neg\psi$, and observe that φ is a separating formula for $(\mathcal{A}, \mathcal{B})$. Furthermore, $f(\varphi) = h_{\neg}(f(\psi)) = h_{\neg}(r') = r$.
- Suppose Spoiler plays **Split-Left**, creating $X_1 = \langle (\mathcal{A}_1, \mathcal{B}) \mid r_1 \rangle$ and $X_2 = \langle (\mathcal{A}_2, \mathcal{B}) \mid r_2 \rangle$. By induction, $(\mathcal{A}_1, \mathcal{B})$ is separable by a formula φ_1 with $f(\varphi_1) = r_1$, and $(\mathcal{A}_2, \mathcal{B})$ is separable by a formula φ_2 with $f(\varphi_2) = r_2$. Then, $\varphi = \varphi_1 \vee \varphi_2$ is a separating formula for $(\mathcal{A}, \mathcal{B})$, and $f(\varphi) = h_{\vee}(r_1, r_2) = r$.
- If Spoiler plays **Split-Right**, the analysis is similar.
- Suppose Spoiler plays **Pebble-Left** using the pebble color x_i , creating $X = \langle (\mathcal{A}', \mathcal{B}') \mid r' \rangle$. By induction, $(\mathcal{A}', \mathcal{B}')$ has a separating formula ψ with $f(\psi) = r'$. Let $\varphi = \exists x_i \psi$, and consider one particular structure-assignment pair $(\mathbf{A}, \alpha_{\mathbf{A}}) \in \mathcal{A}$. We know this pair became a pair $(\mathbf{A}, \alpha'_{\mathbf{A}}) \in \mathcal{A}'$, such that $(\mathbf{A}, \alpha'_{\mathbf{A}}) \models \psi$. Furthermore, the only difference between the two pairs is the mapping of the domain element x_i . Regardless of whether or not $x_i \in \text{dom}(\alpha_{\mathbf{A}})$, we have $(\mathbf{A}, \alpha_{\mathbf{A}}) \models \exists x_i \psi$, just by reverting the mapping of the element x_i (or arbitrarily mapping it, if it was not in the domain). On the other hand, if some $(\mathbf{B}, \alpha_{\mathbf{B}}) \in \mathcal{B}$ satisfies $(\mathbf{B}, \alpha_{\mathbf{B}}) \models \exists x_i \psi$, then Duplicator can respond to Spoiler's Pebble-Left move by (re-)mapping the variable x_i to a witness for ψ , and the resulting pair $(\mathbf{B}, \alpha'_{\mathbf{B}}) \in \mathcal{B}'$ would satisfy $(\mathbf{B}, \alpha'_{\mathbf{B}}) \models \psi$. This is a contradiction, so every $(\mathbf{B}, \alpha_{\mathbf{B}}) \in \mathcal{B}$ satisfies $(\mathbf{B}, \alpha_{\mathbf{B}}) \models \neg \exists x_i \psi \equiv \neg \varphi$. Therefore, φ is a separating formula for $(\mathcal{A}, \mathcal{B})$. Of course, $f(\varphi) = f(\exists x_i \psi) = h_{\exists}(f(\psi)) = h_{\exists}(r') = r$.
- If Spoiler plays **Pebble-Right**, the analysis is similar.

(\impliedby .) We use induction on the separating formula. Suppose there is a separating formula $\varphi \in \text{FO}^k(\tau)$ for $(\mathcal{A}, \mathcal{B})$ with $f(\varphi) = r$. We will show inductively that Spoiler has a winning strategy in the game $\mathcal{SG}(r, k; f)$ on $(\mathcal{A}, \mathcal{B})$.

If φ is atomic, Spoiler can just use φ to close the root node $X_{\text{root}} = \langle (\mathcal{A}, \mathcal{B}) \mid r \rangle$. Otherwise, inductively, we have the following cases:

- If $\varphi = \neg\psi$, then Spoiler plays **Swap**, creating child $X_1 = \langle (\mathcal{B}, \mathcal{A}) \mid r' \rangle$ where $f(\psi) = r'$ and $r = h_{\neg}(r')$. Note that the precondition is met, so Spoiler can play this move. Note that $\psi \in \text{FO}^k(\tau)$ separates $(\mathcal{B}, \mathcal{A})$, so by induction, Spoiler wins the remaining game.
- If $\varphi = \psi \vee \gamma$, then every $(\mathbf{A}, \alpha_{\mathbf{A}}) \in \mathcal{A}$ satisfies $(\mathbf{A}, \alpha_{\mathbf{A}}) \models \psi \vee \gamma$, whereas every $(\mathbf{B}, \alpha_{\mathbf{B}}) \in \mathcal{B}$ satisfies $(\mathbf{B}, \alpha_{\mathbf{B}}) \models \neg(\psi \vee \gamma) \equiv \neg\psi \wedge \neg\gamma$. Let \mathcal{A}_1 be the subset of \mathcal{A} satisfying ψ , and let $\mathcal{A}_2 = \mathcal{A} - \mathcal{A}_1$. Let $f(\psi) = r_1$ and $f(\gamma) = r_2$, so that $r = h_{\vee}(r_1, r_2)$. Now, Spoiler plays **Split-Left** with nodes $\langle (\mathcal{A}_1, \mathcal{B}) \mid r_1 \rangle$ and $\langle (\mathcal{A}_2, \mathcal{B}) \mid r_2 \rangle$. Note that the precondition is met, so Spoiler can play this move. The formula $\psi \in \text{FO}^k(\tau)$ separates $(\mathcal{A}_1, \mathcal{B})$ and satisfies $f(\psi) = r_1$. Similarly, the formula $\gamma \in \text{FO}^k(\tau)$ separates $(\mathcal{A}_2, \mathcal{B})$ and satisfies $f(\gamma) = r_2$. By induction, Spoiler wins the remaining games.
- If $\varphi = \psi \wedge \gamma$, Spoiler plays **Split-Right** analogously.
- Suppose $\varphi = \exists x_i \psi$. If $f(\psi) = r'$, then note that $r = h_{\exists}(r')$. Note that the precondition is met, so Spoiler can play **Pebble-Left**. By assumption, every $(\mathbf{A}, \alpha_{\mathbf{A}}) \in \mathcal{A}$ has $(\mathbf{A}, \alpha_{\mathbf{A}}) \models \exists x_i \psi$. So for some $a \in A$, $(\mathbf{A}, \alpha'_{\mathbf{A}}) \models \psi$, where $\alpha'_{\mathbf{A}}$ is identical to $\alpha_{\mathbf{A}}$ except that x_i has been (re)mapped to the element a . Spoiler plays **Pebble-Left**, placing pebble x_i on the witness a on each \mathbf{A} . Note that this creates exactly $(\mathbf{A}, \alpha'_{\mathbf{A}})$. On the other hand, every $(\mathbf{B}, \alpha_{\mathbf{B}}) \in \mathcal{B}$ satisfies $(\mathbf{B}, \alpha_{\mathbf{B}}) \models \neg \exists x_i \psi \equiv \forall x_i \neg \psi$, and so every $b \in B$ is a witness for $\neg \psi$, and so even after making copies, Duplicator can never play on a witness for ψ in any \mathbf{B} . So for each pair $(\mathbf{B}, \alpha'_{\mathbf{B}})$ created by Duplicator, we have $(\mathbf{B}, \alpha'_{\mathbf{B}}) \models \neg \psi$. At the end, therefore, we have the new tree node $X' = \langle (\mathcal{A}', \mathcal{B}') \mid r' \rangle$, where $\psi \in \text{FO}^k(\tau)$ is a separating formula for $(\mathcal{A}', \mathcal{B}')$ with $f(\psi) = r'$, and so by induction, Spoiler wins the remaining game.
- If $\varphi = \forall x_i \psi$, Spoiler plays **Pebble-Right** analogously.

This proves both directions of the theorem. \square

We remark here that Theorem 6.7 is applicable directly to several straightforward compositional syntactic measures, including the ones in Examples 6.2, 6.3, and 6.4. In particular, an examination of Example 6.2 shows that the (r, k) -QVT game is identical to $\text{SG}(r, k; f_q)$, where f_q is the quantifier-number measure; this proves Theorem 5.6 immediately.

We now complete this subsection with a result showing that a familiar tool for analyzing these games carries over to this general framework as well. The following proposition shows that Duplicator's oblivious strategy is optimal in the syntactic game for *every* compositional syntactic measure.

Proposition 6.8. *Let f be a compositional syntactic measure. If Duplicator has a winning strategy in the $\text{SG}(r, k; f)$ on $(\mathcal{A}, \mathcal{B})$, then the oblivious strategy is winning.*

Proof. It suffices to consider the case where Spoiler plays **Pebble-Left** on some node $X = \langle (\mathcal{A}', \mathcal{B}') \mid r' \rangle$ of \mathcal{T} , creating the new node $X' = \langle (\mathcal{A}'', \mathcal{B}'') \mid r'' \rangle$. Note that $h_{\exists}(r'') = r'$. The **Pebble-Right** move will be analogous, and Duplicator has no agency elsewhere.

Suppose there is no separating formula $\varphi \in \text{FO}^k(\tau)$ for $(\mathcal{A}', \mathcal{B}')$ with $f(\varphi) = r'$. It suffices to show that regardless of Spoiler's move, as long as Duplicator responds obliviously, there is no separating formula $\psi \in \text{FO}^k(\tau)$ for $(\mathcal{A}'', \mathcal{B}'')$ with $f(\psi) = r''$.

Suppose WLOG Spoiler plays the pebble color x_i on his move. Consider all pairs $(\mathbf{A}, \alpha'_{\mathbf{A}}) \in \mathcal{A}''$, and take any arbitrary formula $\psi \in \text{FO}^k(\tau)$ with $f(\psi) = r''$ that is true for *all* these pairs. In other words, for every pair $(\mathbf{A}, \alpha'_{\mathbf{A}}) \in \mathcal{A}''$, we have $(\mathbf{A}, \alpha'_{\mathbf{A}}) \models \psi$. An arbitrary pair $(\mathbf{A}, \alpha'_{\mathbf{A}}) \in \mathcal{A}''$ arose from a pair $(\mathbf{A}, \alpha_{\mathbf{A}}) \in \mathcal{A}'$, with Spoiler changing the

assignment $\alpha_{\mathbf{A}}$ by just (re)assigning the variable x_i . Furthermore, every pair $(\mathbf{A}, \alpha_{\mathbf{A}}) \in \mathcal{A}'$ gave rise to a new pair in this way.

We claim that this means the formula $\sigma = \exists x_i \psi$ is true of all pairs $(\mathbf{A}, \alpha_{\mathbf{A}}) \in \mathcal{A}'$. This follows by simply reversing the assignment of the variable x_i in each of the pairs. It follows that for each pair $(\mathbf{A}, \alpha_{\mathbf{A}}) \in \mathcal{A}'$, we have $(\mathbf{A}, \alpha_{\mathbf{A}}) \models \sigma$. Note that $f(\sigma) = f(\exists x_i \psi) = h_{\exists}(f(\psi)) = h_{\exists}(r'') = r'$. By our assumption, therefore, there must be some pair $(\mathbf{B}, \alpha_{\mathbf{B}}) \in \mathcal{B}'$ satisfying $(\mathbf{B}, \alpha_{\mathbf{B}}) \models \sigma$. In other words, $(\mathbf{B}, \alpha_{\mathbf{B}}) \models \exists x_i \psi$.

Now, when Duplicator responds obliviously, she (re)assigns the variable x_i in all possible ways in each of the pairs in \mathcal{B}' . In particular, therefore, in one of the copies of the pair $(\mathbf{B}, \alpha_{\mathbf{B}})$ defined above, she plays on a witness $b \in B$ for ψ , creating the pair $(\mathbf{B}, \alpha'_{\mathbf{B}}) \in \mathcal{B}''$. But then, $(\mathbf{B}, \alpha'_{\mathbf{B}}) \models \psi$, and so ψ cannot be a separating formula for $(\mathcal{A}'', \mathcal{B}'')$.

Since ψ was arbitrary, we are done. \square

Of course, Proposition 5.7 follows immediately as a corollary.

It is worth remarking that the proof of Proposition 6.8 also captures why it is critical for Duplicator to use her copying ability. Duplicator needs a pair $(\mathbf{B}, \alpha_{\mathbf{B}}) \in \mathcal{B}'$ to be a witness for the formula ψ . However, the same pair can be a witnessing structure for two different such formulas. In particular, it could be the case that $(\mathbf{B}, \alpha_{\mathbf{B}}) \models \exists x_i \psi_1$ and $(\mathbf{B}, \alpha_{\mathbf{B}}) \models \exists x_i \psi_2$, but the elements in B that witness ψ_1 and ψ_2 are *different*. Making copies of the pair $(\mathbf{B}, \alpha_{\mathbf{B}})$ in order to have two new assignments of x_i , enabling different elements to act as witnesses simultaneously in the same move, achieves this.

6.3. Properties of some measures. Having defined the syntactic game, it is natural to ask for the properties of specific syntactic measures that can enable us to apply existing tools of analyzing similar games, in order to have the results carry over to this realm.

Consider the quantifier rank syntactic measure, denoted by f_r in Example 6.3. Consider the game $\mathcal{SG}(r, k; f_r)$, and observe that it is very similar to the QVT game. The only difference is that in the Split-Left and Split-Right moves on a tree node $X = \langle (\mathcal{A}', \mathcal{B}') \mid r' \rangle$, Spoiler labels the new nodes in \mathcal{T} with two values whose maximum is r' , instead of two values that sum to r' .

The following proposition considers when we could look at games on each of the singleton pairs, rather than entire classes of τ -structures.

Proposition 6.9. *Consider the syntactic game for any compositional syntactic measure f satisfying:*

1. *there are finitely many equivalence classes of the equivalence relation “satisfy the same FO-sentences with f -measure r and k variables”,*
2. *each such class is definable by a sentence with the parameters defining the class,*
3. *the set of sentences with these parameters is closed under disjunction.*

Then, for all nonempty sets \mathcal{A} and \mathcal{B} of τ -structures, and for all $r, k \in \mathbb{N}$, the following statements are equivalent:

- (a) *Spoiler wins the $\mathcal{SG}(r, k; f)$ on $(\mathcal{A}, \mathcal{B})$.*
- (b) *For every $\mathbf{A} \in \mathcal{A}$ and every $\mathbf{B} \in \mathcal{B}$, there is a separating k -variable sentence φ for $(\{\mathbf{A}\}, \{\mathbf{B}\})$ with $f(\varphi) = r$.*

Proof.

(\implies): Take a separating k -variable sentence φ for $(\mathcal{A}, \mathcal{B})$ with $f(\varphi) = r$. This is separating for every pair $(\{\mathbf{A}\}, \{\mathbf{B}\})$ with $\mathbf{A} \in \mathcal{A}$ and $\mathbf{B} \in \mathcal{B}$.

(\Leftarrow):) By property 1, there are finitely many equivalence classes of the equivalence relation “satisfy the same FO-sentences with measure r and k variables”, and by property 2, each such class is definable by a FO-sentence with these parameters. Let $\theta_1, \dots, \theta_m$ be the list of all such sentences arising from structures in \mathcal{A} , and consider the sentence $\psi = \theta_1 \vee \dots \vee \theta_m$. By property 3, ψ is a k -variable formula with $f(\psi) = r$. We claim that ψ is a separating sentence for $(\mathcal{A}, \mathcal{B})$. Indeed, clearly, every structure in \mathcal{A} satisfies one of the θ_j 's. We claim that no structure $\mathbf{B} \in \mathcal{B}$ does. Otherwise, take $\mathbf{B} \in \mathcal{B}$ satisfying θ_j for some j . Then \mathbf{B} satisfies the same sentences with measure r and k variables as some structure $\mathbf{A} \in \mathcal{A}$. But then $(\{\mathbf{A}\}, \{\mathbf{B}\})$ is inseparable, which is a contradiction. \square

Proposition 6.9 immediately helps us relate the syntactic game on the quantifier rank measure f_r to the other canonical technique we have for quantifier rank — the EF games.

Corollary 6.10. *Let f_r be the quantifier rank syntactic measure. For all nonempty sets \mathcal{A} and \mathcal{B} of τ -structures, and for all $r, k \in \mathbb{N}$, the following statements are equivalent:*

- (a) *Spoiler wins the game $\mathcal{SG}(r, k; f_r)$ on $(\mathcal{A}, \mathcal{B})$.*
- (b) *For every $\mathbf{A} \in \mathcal{A}$ and every $\mathbf{B} \in \mathcal{B}$, Spoiler wins the $EF(r, k)$ game on (\mathbf{A}, \mathbf{B}) .*

We next remark on the formula size measure, denoted by f_s in Example 6.4. Consider the $\mathcal{SG}(r, k; f_s)$ game, and observe that it would be very similar to the QVT game, with the following differences on a node $X = \langle (\mathcal{A}', \mathcal{B}') \mid r' \rangle$:

- For a **Pebble-Left**, **Pebble-Right**, or **Swap** move, Spoiler would decrement the value of the counter r' to $r' - 1$.
- For a **Split-Left** or **Split-Right** move on a node with counter value r' , Spoiler would choose r'_1 and r'_2 such that $r' = r'_1 + r'_2 + 1$.
- Spoiler would only close nodes with counter value of 1.

Note that this last game is very similar to the ones in [AI01] and [HV15]. In fact, as in [HV15], we could have changed the rules of the syntactic game and allowed Duplicator to choose one of the branches from a split move to continue along as well, without affecting the eventual winner. But then we would lose the property that a closed game tree is isomorphic to a parse tree of a separating sentence.

Finally, we remark briefly on the measure f_t that is always zero. This is clearly a valid compositional syntactic measure. It should be clear that $\mathcal{SG}(0, k; f_t)$ is equivalent to the k -pebble game [Bar77, Imm82]. Formally, we claim that Spoiler wins the k -pebble game on (\mathbf{A}, \mathbf{B}) if and only if he wins $\mathcal{SG}(0, k; f_t)$ on $(\{\mathbf{A}\}, \{\mathbf{B}\})$. Indeed, if Spoiler has a winning strategy on the k -pebble game, he can make the same move at the root node of \mathcal{T} , and then once Duplicator responds, he can split the resulting node into singletons and win inductively on each of them. Conversely, if Duplicator wins the k -pebble game, she just plays the same moves on the syntactic game without making any copies, proving the result.

7. OPEN PROBLEMS & FUTURE DIRECTIONS

Our results suggest several open problems about multi-structural games and their variants. To begin with, we saw that the hereditary MS game with repegging (Game 5.5) does not simultaneously capture the number of quantifiers and the number of variables. What is the fragment of FO captured by this game? Considered from a slightly different perspective, how is the “hereditary” property reflected by FO semantics?

We also investigated the variant of the MS game in which Spoiler wins without ever playing on top. In some cases, Spoiler wins the MS game on $(\mathcal{A}, \mathcal{B})$ without playing on top on the left side (but may play on top on the right), while in other cases, this is reversed. Is it true that Spoiler wins the r -round MS game on $(\mathcal{A}, \mathcal{B})$ without playing on top if and only if Spoiler wins both the r -round MS game on $(\mathcal{A}, \mathcal{B})$ without playing on top on the left side, and the r -round MS game on $(\mathcal{A}, \mathcal{B})$ without playing on top on the right side?

We investigated compositional syntactic measures and their associated games for first-order logic (introduced for the case of infinitary formulas by [VW13]). These definitions allowed us to relate structural properties of such complexity measures to properties of game strategies. What other relationships are possible? Can we unify and better understand logical inexpressibility results by studying syntactic measures?

Beyond these and other related problems, the main challenge is to use syntactic games to prove lower bounds on any reasonable complexity measure, first on unordered structures and then, hopefully, on ordered structures — the latter will be very difficult, because of connections to longstanding open problems in computational complexity (see Section 1). A few exceptions notwithstanding (e.g., [Sch94, Ruh99, Sch01]), there are relatively few inexpressibility results obtained using EF games on ordered structures or on structures with built-in arithmetic predicates. Can we go beyond these applications of the EF games? What about other syntactic measures on ordered structures, such as simultaneous bounds on number of quantifiers and number of variables?

ACKNOWLEDGMENTS

We would like to thank Ryan Williams (MIT) for many helpful discussions. Rik Sengupta acknowledges NSF grant CCF-1934846. We also would like to thank the anonymous referees for an incredibly careful reading of the paper, and for many helpful suggestions and edits.

REFERENCES

- [AI01] Micah Adler and Neil Immerman. An $n!$ Lower Bound on Formula Size. In *16th Annual IEEE Symposium on Logic in Computer Science, Boston, Massachusetts, USA, June 16-19, 2001, Proceedings*, pages 197–206. IEEE Computer Society, 2001. doi:10.1109/LICS.2001.932497.
- [Bar77] Jon Barwise. On Moschovakis Closure Ordinals. *J. Symb. Log.*, 42(2):292–296, 1977. doi:10.2307/2272133.
- [CFI⁺24] Marco Carmosino, Ronald Fagin, Neil Immerman, Phokion G. Kolaitis, Jonathan Lenchner, and Rik Sengupta. On the Number of Quantifiers Needed to Define Boolean Functions. In Rastislav Kráľovič and Antonín Kučera, editors, *49th International Symposium on Mathematical Foundations of Computer Science (MFCS 2024)*, volume 306 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 34:1–34:16, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.MFCS.2024.34>, doi:10.4230/LIPIcs.MFCS.2024.34.
- [dR84] Michel de Rougemont. Uniform Definability on Finite Structures with Successor. In Richard A. DeMillo, editor, *Proceedings of the 16th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1984, Washington, DC, USA*, pages 409–417. ACM, 1984. doi:10.1145/800057.808707.
- [Ehr61] Andrzej Ehrenfeucht. An Application of Games to the Completeness Problem for Formalized Theories. *Fundamenta Mathematicae*, 49:129–141, 1961. doi:10.4064/fm-49-2-129-141.
- [Fag75] Ronald Fagin. Monadic Generalized Spectra. *Math. Log. Q.*, 21(1):89–96, 1975. doi:10.1002/malq.19750210112.

- [FLRV21] Ronald Fagin, Jonathan Lenchner, Kenneth W. Regan, and Nikhil Vyas. Multi-Structural Games and Number of Quantifiers. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13, 2021. doi:10.1109/LICS52264.2021.9470756.
- [FLVW22] Ronald Fagin, Jonathan Lenchner, Nikhil Vyas, and Ryan Williams. On the Number of Quantifiers as a Complexity Measure. In Stefan Szeider, Robert Ganian, and Alexandra Silva, editors, *47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022)*, volume 241 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 48:1–48:14, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: <https://drops.dagstuhl.de/opus/volltexte/2022/16846>, doi:10.4230/LIPIcs.MFCS.2022.48.
- [Fra54] Roland Fraïssé. Sur Quelques Classifications des Systèmes de Relations. *Université d'Alger, Publications Scientifiques, Série A*, 1:35–182, 1954. doi:10.2307/2963939.
- [FSV95] Ronald Fagin, Larry J. Stockmeyer, and Moshe Y. Vardi. On Monadic NP vs. Monadic co-NP. *Inf. Comput.*, 120(1):78–92, 1995. doi:10.1006/inco.1995.1100.
- [GS05] Martin Grohe and Nicole Schweikardt. The Succinctness of First-Order Logic on Linear Orders. *Log. Methods Comput. Sci.*, 1(1), 2005. doi:10.2168/LMCS-1(1:6)2005.
- [HV15] Lauri Hella and Jouko Väänänen. The Size of a Formula as a Measure of Complexity. In Åsa Hirvonen, Juha Kontinen, Roman Kossak, and Andrés Villaveces, editors, *Logic Without Borders: Essays on Set Theory, Model Theory, Philosophical Logic and Philosophy of Mathematics*, pages 193–214, Berlin, München, Boston, 2015. De Gruyter. doi:doi:10.1515/9781614516873.193.
- [Imm81] Neil Immerman. Number of Quantifiers is Better than Number of Tape Cells. *Journal of Computer and System Sciences*, 22(3):384–406, 1981. URL: <https://www.sciencedirect.com/science/article/pii/0022000081900398>, doi:10.1016/0022-0000(81)90039-8.
- [Imm82] Neil Immerman. Upper and Lower Bounds for First Order Expressibility. *Journal of Computer and System Sciences*, 25(1):76–98, 1982. URL: <https://www.sciencedirect.com/science/article/pii/0022000082900113>, doi:10.1016/0022-0000(82)90011-3.
- [Imm99] Neil Immerman. *Descriptive Complexity*. Graduate Texts in Computer Science. Springer, 1999. URL: <https://doi.org/10.1007/978-1-4612-0539-5>.
- [KV92] Phokion G. Kolaitis and Moshe Y. Vardi. Infinitary Logics and 0-1 Laws. *Inf. Comput.*, 98(2):258–294, 1992. doi:10.1016/0890-5401(92)90021-7.
- [Ruh99] Matthias Ruhl. Counting and Addition Cannot Express Deterministic Transitive Closure. In *14th Annual IEEE Symposium on Logic in Computer Science, Trento, Italy, July 2-5, 1999*, pages 326–334. IEEE Computer Society, 1999. doi:10.1109/LICS.1999.782627.
- [Sch94] Thomas Schwentick. Graph Connectivity and Monadic NP. In *35th Annual Symposium on Foundations of Computer Science, Santa Fe, New Mexico, USA, 20-22 November 1994*, pages 614–622. IEEE Computer Society, 1994. doi:10.1109/SFCS.1994.365730.
- [Sch01] Nicole Schweikardt. *On the expressive power of first-order logic with built-in predicates*. Universität Mainz, 2001.
- [VW13] Jouko Väänänen and Tong Wang. An Ehrenfeucht-Fraïssé game for $L_{\omega_1\omega}$. *Math. Log. Q.*, 59(4-5):357–370, 2013. doi:10.1002/malq.201200104.