Model Comparison Games for Generalized Quantifiers

Joint work with Antti Kuusisto & Matias Selin

From logical separability → games → formula complexity → minor quantifiers

The question behind the paper

- First-order logic cannot express every structural distinction.
- Generalized quantifiers increase expressive power — but how do we measure the new distinctions?
- Model-comparison games turn “there exists a separating formula” into a concrete winning strategy.

Central problem

Given structures A and B, characterize when some FO(Q)-formula is true in A and false in B.

Why games?

They expose the resources a formula needs: depth, size, or a more general cost.

Article: Abstract & §1

TALK STRUCTURE

Roadmap

FO(Q): first-order syntax and semantics

$$\varphi ::= x_1 = x_2 \;\;|\;\; R(x_1,\ldots,x_k) \;\;|\;\; Q\, x_1,\ldots,x_k\, (\varphi_1,\ldots,\varphi_k)$$

Let A be a τ-model and f an assignment over A whose domain includes all free variables of the formula being evaluated. The semantics of FO(Q) are as follows:

1. $A, f \vDash x_1 = x_2 \;\Leftrightarrow\; f(x_1) = f(x_2)$.
2. $A, f \vDash R(x_1, \ldots, x_k) \;\Leftrightarrow\; (f(x_1), \ldots, f(x_k)) \in R^A$.
3. $A, f \vDash Q\, x_1, \ldots, x_k\, (\varphi_1, \ldots, \varphi_k) \;\Leftrightarrow$
   $(A, \|\varphi_1\|_{x_1}^{A,f}, \ldots, \|\varphi_k\|_{x_k}^{A,f}) \in Q$,

where $\|\varphi\|_x^{A,f} := \{a \mid A, f[a/x] \vDash \varphi\}$

is the extension of φ relative to (A, f) and x.

Article: §2.1, especially the grammar and semantics of FO(Q)

Quantifier depth and d-types

Quantifier depth

The quantifier depth of an FO(Q)-formula is the maximum number of nested quantifiers in the formula.

$FO(Q)^d$

denotes the formulas of quantifier depth at most d.

d-type

For $a \in A$, relative to (A,f) and x, the d-type of a is its equivalence class under $FO(Q)^d$:

$$[a]^d_x = \{ b \in A \mid (A,f[a/x]) \equiv^d_{FO(Q)} (A,f[b/x]) \}$$

Thus two points have the same d-type when no formula of depth $\le d$ distinguishes them.

Article: §2.2 (Types)

Three quantifiers to keep in mind

$\exists x\ \varphi$

Width 1, type (1).
Checks whether the extension of φ is nonempty.

Härtig $I\ x,y(\varphi,\psi)$

Width 2, type (1,1).
Checks whether the two extensions have the same cardinality.

Ham $(xy)\ \varphi$

Width 1, type (2).
Checks whether the binary relation defined by φ contains a Hamiltonian path.

Generalized quantifiers can encode genuinely global properties.

Article: Example 2.1

Separability is the target notion

$A,f \models \varphi$ $B,g \not\models \varphi$

If such a φ exists, the two pointed structures are separable in FO(Q).

Article: §2.1, definition of separability

TECHNICAL BRIDGE

Depth-d types make definable sets manageable

- $FO(Q)^d$ contains formulas of quantifier depth at most d.
- Two points have the same d-type when no depth-d formula distinguishes them.
- Key result: every d-type is definable by a depth-d formula (finite vocabulary).
- Therefore every union of d-types is also depth-d definable.

Why this matters in the game

Witness sets that survive contestation must be unions of d-types — hence they correspond to actual formulas.

Article: Lemma 2.3 and Corollary 2.4

Generalized EF game: replace points by sets

Classical EF intuition

Generalized-quantifier version

The set X is meant to be the extension of the subformula under Q.

Article: §3 opening intuition

GAME 1

One EF{Q} round

1. Attacker chooses Q and variable x
2. Attacker chooses witness X in one model + spillover P in the other
3. Defender chooses a Q-witness X′ containing P
4. Attacker challenges membership/non-membership and picks points
5. Check partial isomorphism; if it survives, continue

Article: Definition 3.1

Step 4 tests membership in the witness sets

After the defender chooses X', the attacker chooses one of two challenges.

A. Challenge non-membership

The attacker chooses

$v' \in N \setminus X'$ **and** $v \in X$.

The players swap roles, then continue from

$(M, N, h[x \mapsto v], h'[x \mapsto v'])$.

A point outside X' is matched with one inside X.

B. Challenge membership

The attacker chooses $v' \in X'$. The defender replies:

- $v \in X \;\rightarrow\;$ continue with models M and N;
- $v \in P \;\rightarrow\;$ continue inside N on both sides.

$(M, N, h[x \mapsto v], h'[x \mapsto v'])$

or $(N, N, h'[x \mapsto v], h'[x \mapsto v'])$.

Why this step matters: It tests whether membership and non-membership are preserved by the assignments.

Article: Definition 3.1, step 4 (p. 6)

Why a spillover set?

Problem without P

A witness set X in model A only represents d-types realized in A. Model B may contain additional d-types that also satisfy the intended subformula.

Role of P

P forces the defender's witness X′ to include relevant satisfying types in B that are not realized in A. It prevents an artificially "too small" response.

P handles satisfying types that exist only on the other side.

Article: intuition following Figure 1

Contestation prevents “cheating” with undefinable sets

Witness set contested

If X splits a d-type…

The opponent picks one point inside X and an equivalent point outside X. If they are depth-d equivalent, the set could not be the extension of a depth-d formula.

Spillover contested

If P contains a type already realized…

The opponent pairs that spillover point with a matching point in the first model and reduces the claim to the lower-depth game.

Contestation enforces semantic legitimacy locally.

Article: Definition 3.1 contestation rules; Lemma 3.3 intuition

Worked example: “there exist exactly three”

Separating formula: $\exists^{=3} x\,(B(x) \vee R(x))$

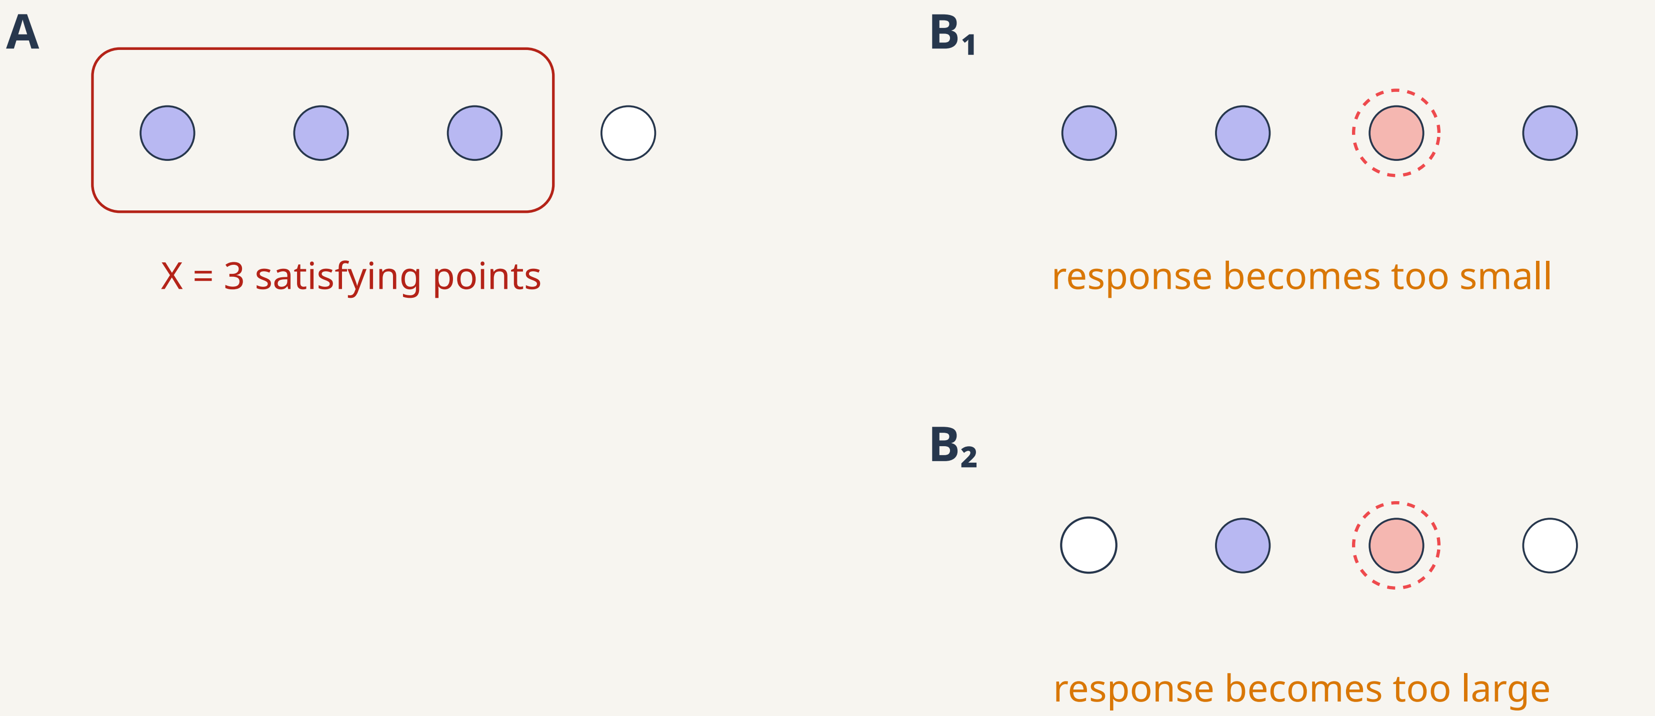

Game idea

Choose X as the extension in A. Any legal Q-witness X′ in B_1/B_2 disagrees with the intended extension in a way the attacker can expose.

Article: Figure 1 and accompanying explanation

GAME 1

The depth theorem

Player I wins $EF^d\{Q\}(A,B,f,g)$ ⇔ **(A,f) and (B,g) are separated by an FO(Q) formula of depth ≤ d**

- Base case: atomic disagreement = depth 0.
- Induction: a winning set move yields definable d-type unions; a separating Qxφ yields the attacker's witness/spillover strategy.
- Unbounded version: Player I wins iff A and B are separable in FO(Q).

Article: Lemmas 3.2–3.4 and Theorem 3.5

Proof skeleton: game ↔️ formula

Winning strategy ⇒ formula

1. Contestation forces witness sets to be unions of d-types.
2. Corollary 2.4 turns those unions into formulas θ, ψ.
3. Combine them under Q to build a depth-(d+1) separator.

Formula ⇒ winning strategy

1. Start from a separator Qxφ.
2. Choose X = extension of φ in A.
3. Put into P the satisfying B-types absent from A.
4. Any defender response exposes a lower-depth difference.

The induction mirrors the syntax tree of the separating formula.

Article: proofs of Lemmas 3.3 and 3.4

From depth to size: another measure

Size of a formula

$s(\text{atomic}) = 1$

$s(\neg\varphi) = s(\varphi)+1$

$s(\varphi \wedge \psi) = s(\varphi)+s(\psi)+1$

$s(Qx\varphi) = s(\varphi)+1$

What changes?

Depth asks for the longest nesting chain. Size charges for every syntactic contribution, including negation. So branching matters, not only nesting.

A budgeted game should spend exactly what the formula spends.

Article: §4 opening and size definition

Formula-size game between classes

- Position: (budget s, left class A, right class B).
- Immediate win if an atomic formula separates the classes.
- Player I's move follows the outermost connective/quantifier of a candidate separator.

Negation

Spend 1 and swap the two classes.

Conjunction

Split the budget and cover the right class by two subproblems.

Quantifier

Spend 1, choose per-model subsets, and continue on inside/outside assignments.

Article: Definition 4.1

Conjunction = split the adversary, split the budget

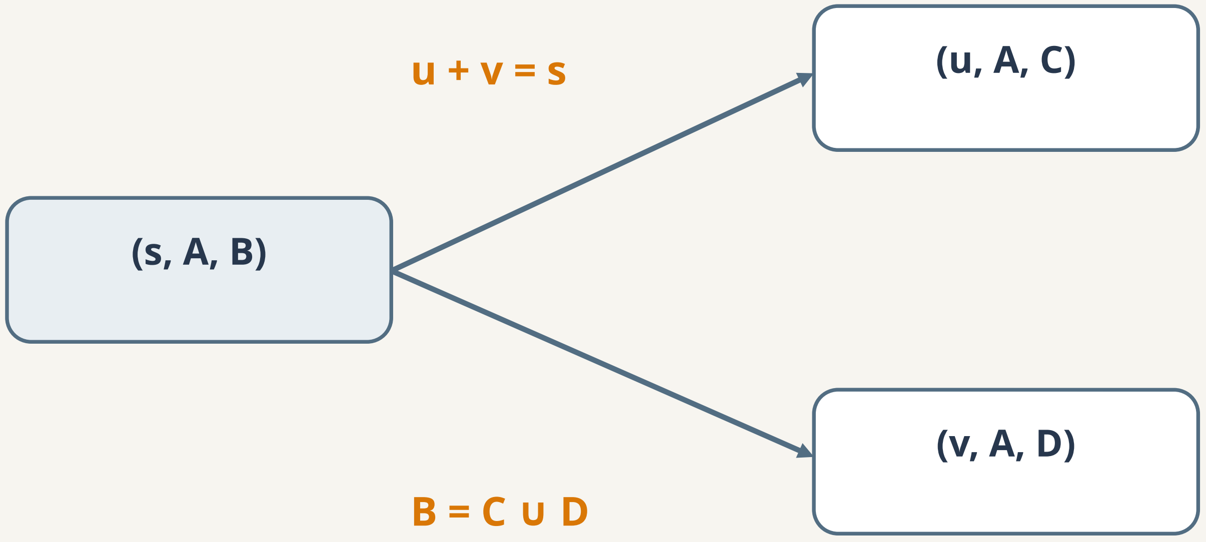

Player II chooses which branch must be defended.

Logical meaning

If $\varphi \wedge \psi$ separates A from B, every B-model falsifies φ or ψ. The cover records exactly that.

Article: Definition 4.1(b) and Theorem 4.2 proof

Quantifier move = choose the intended extension

For each model C, Player I chooses $P(C,h) \subseteq C$

Left class

Require $(A, P(A,f)) \in Q$.
Assignments using points inside $P(C,h)$ go to C^+.

Right class

Require $(B, P(B,g)) \notin Q$.
Assignments using points outside $P(C,h)$ go to C^-.

The game continues from $(s - 1, C^+, C^-)$.

A lower-budget formula must separate C^+ from C^-.

Article: Definition 4.1(c) and intuition after it

GAME 2

The size theorem

Player I wins $FS_s\{Q\}(A,B)$ ⇔ **A and B are separated by an FO(Q) formula of size ≤ s**

Proof (idea)

Induction on the budget.

Negation swaps classes, conjunction partitions the countermodels and budget, and a quantifier move identifies the extension of a smaller separating formula.

Article: Theorem 4.2

Strong vs weak separability

Strong separability

One uniform formula φ separates every member of class A from every member of class B.

Weak separability

For each pair (A,f) ∈ A and (B,g) ∈ B, there may be a different separating formula φ(A,B).

Strong ⇒ weak. For infinite classes, the converse can fail.

Article: §4.2

Infinite classes: weak need not imply strong

Example with unary P

A_n has exactly n elements satisfying P

Odd n

$\{A_n \mid n \text{ odd}\}$

Even n

$\{A_n \mid n \text{ even}\}$

Each odd/even pair can be separated by a suitable finite counting formula, but no single first-order formula separates all odd sizes from all even sizes.

Article: Example 4.3

For finite classes, weak = strong

Theorem 4.6

If A and B are finite classes, they are separable iff they are weakly separable.

Finite conjunctions collect all opponents for one A; the finite disjunction then covers all A.

Article: Theorem 4.6

A more direct game between two models

- The class-based game is powerful but gives Player II limited agency.
- The paper also develops a direct model-vs-model formula-size game, better suited to lower bounds.
- However, naively letting Player I choose arbitrary witness subsets can make the budget undercount the formula needed to define those subsets.

Cautionary phenomenon

Player I can win with a small budget by choosing a set that itself requires a larger formula to define.

Article: discussion before §5 and example immediately preceding §5

Minor quantifiers: witnesses need not be full extensions

Generalized Q asks for the entire extension P.

Full extension

P = { tuples satisfying φ }
Q decides from P.

Minor M records sufficient positive/negative evidence.

Witness / falsifier sets

$P^+ \subseteq$ extension(φ)
$P^- \subseteq$ complement(extension(φ))
Enough evidence certifies Q or its complement.

Article: §5.1 and Lemma 5.1

Definition of a minor quantifier

When C witnesses Q

Let Q have width 1 and type (1). An isomorphism-closed class C of triples (D, P^+, P^-) witnesses Q if:

1. $D \neq \emptyset$, $P^+ \cup P^- \subseteq D$, and $P^+ \cap P^- = \emptyset$.

2. For every $(D,P) \in Q$, some $(D,P^+,P^-) \in C$ satisfies $P^+ \subseteq P$ and $P^- \subseteq D \setminus P$.

3. For every triple in C, there is no (D,P') in the complement of Q with $P^+ \subseteq P'$ and $P^- \subseteq D \setminus P'$.

A minor of Q

The complement is
$Q^C = \{(D,P) : (D,P) \notin Q\}$.

A pair $M = (C,D)$, where C witnesses Q and D witnesses Q^C, is a minor of Q.

Semantics

$A,f \models Mx\varphi$ iff some $(A,P^+,P^-) \in C$ satisfies $P^+ \subseteq \mathrm{Ext}(\varphi; A,f,x)$ and $P^- \subseteq A \setminus \mathrm{Ext}(\varphi; A,f,x)$.

Article: §5.1, pp. 16–17 (definition immediately before Lemma 5.1)

Connectives become quantifiers too

Negation

A width-1, type-(0) quantifier can test whether the extension of φ is empty.

Conjunction

A width-2, type-(0,0) quantifier can encode the truth values of both subformulas.

$\top$ / $\bot$

Width-0 quantifiers can encode constants and domain properties.

Once operators are quantifiers, one generic move can represent the whole syntax.

Article: Examples 2.2 and 5.3; §5 motivation

Conjunction as a minor quantifier

Conjunction has width 2 and type (0,0). For a 0-ary relation, $\{\emptyset\}$ represents true and $\emptyset$ represents false.

Canonical minor $M_\wedge = (C_\wedge, D_\wedge)$

Truth witness — both conjuncts are true:

$C_\wedge = \{ (D, \{\emptyset\}, \{\emptyset\}, \emptyset, \emptyset) \}$

Falsifying witnesses — at least one conjunct is false:

$D_\wedge = \{$

$(D, \{\emptyset\}, \emptyset, \emptyset, \{\emptyset\}),$ [true, false]

$(D, \emptyset, \{\emptyset\}, \{\emptyset\}, \emptyset),$ [false, true]

$(D, \emptyset, \emptyset, \{\emptyset\}, \{\emptyset\})$ [false, false]

$\}$

A stricter minor

Keep the same $C_\wedge$, but use:

$D'_\wedge = \{$

$(D, \emptyset, \emptyset, \emptyset, \{\emptyset\}),$

$(D, \emptyset, \emptyset, \{\emptyset\}, \emptyset)$

$\}$

These are the two prime implicants of the complement of conjunction: one certifies that the second conjunct is false; the other, that the first is false.

Article: Example 5.3, p. 18

Formula-cost game: one framework, customizable complexity

Assign a cost to each relation and minor quantifier.

Step 1

If an affordable atomic/width-0 formula separates, Player I wins.

Step 2

Choose one minor quantifier plus witness/falsifier sets; spend its cost.

Step 3

Player II chooses a branch; continue with the allocated sub-budget.

Article: Definition 5.4

The formula-cost characterization

Theorem 5.8

Player I has a winning strategy in the $FC_s\{M\}(A, B)$-game if and only if A and B are separable by an FO(M)-formula of cost ≤ s.

Article: Theorem 5.8, p. 21

The same framework also subsumes EF-style rank

Choose costs relations = 0 · $\neg$ = 0 · $\wedge$ = 0 · $\forall,\exists$ = 1

Formula-size style

Budgets are split across branches because size is additive.

Quantifier-rank style

Copy the remaining budget to every branch because rank is a maximum, not a sum.

Article: §5.3

The quantifier-rank characterization

Theorem 5.10

Player I has a winning strategy in the $EF_s\{M\}(A, B)$-game if and only if A and B are separable by an FO(M)-formula of quantifier rank ≤ s.

Article: Theorem 5.10, p. 22

SYNTHESIS

What the paper contributes

1	**Generalized EF game**	Characterizes FO(Q) separability by quantifier depth.
2	**Formula-size game**	Characterizes separation of model classes by formula size, with negation counted explicitly.
3	**Weak separability analysis**	Shows weak and strong separability coincide for finite classes.
4	**Minor-quantifier framework**	Unifies formula-cost and EF-style games and supports flexible cost measures.

Article: Abstract, §1, Theorems 3.5, 4.2, 4.6, 5.8, 5.10

Thank you

Model Comparison Games for Generalized Quantifiers

Antti Kuusisto · Miguel Moreno · Matias Selin