

BOUNDED FIRST-ORDER LOGIC

A Logic of Bounded Quantification: Syntax, Semantics, and Metatheory

Working Paper, 2026

Abstract

Standard first-order logic permits unrestricted quantification: a variable may range over the entire domain of a structure without explicit constraint. This paper defines Bounded First-Order Logic (BFOL), a principled restriction of first-order logic in which every quantifier carries an explicit bounding term. Unbounded quantifier forms are not merely restricted in their interpretation; they are absent from the language.

The paper gives the complete syntax, semantics, and proof theory of BFOL in both natural deduction and sequent calculus formulations, and establishes their equivalence. The core metatheory is developed in full: soundness, completeness (via embedding into standard FOL), decidability of truth in finite structures, cut-elimination for the sequent calculus, Craig interpolation (proved directly by Maehara's method on the BFOL sequent calculus), and Beth definability. The relationship between BFOL and standard FOL is made precise: BFOL embeds faithfully into FOL but is strictly less expressive, by design.

A metatheoretic semantic restriction, the Global Boundedness Principle, is introduced to specify BFOL's behaviour when intended models are finite. Under this restriction, compactness and the Löwenheim-Skolem theorems fail, and Trakhtenbrot's theorem establishes that finite-model validity is not recursively enumerable. The contrast between single-model decidability and all-finite-model non-enumerability is a characteristic feature of the logic in this setting. BFOL is suited to any formal theory whose quantification is intended to range over bounded domains.

Keywords: bounded quantification, first-order logic, finite model theory, Craig interpolation, Maehara partition, cut-elimination, bounded foundations, Trakhtenbrot's theorem, Global Boundedness Principle

1. Introduction

First-order logic is the standard logical framework for formal mathematics. Its language is economical, its proof theory well understood, and its semantics (truth defined relative to a structure and a variable assignment) clean and general. In most mathematical contexts, these features make it the natural logical vehicle.

One feature, however, is not neutral in all settings: unrestricted quantification. In standard FOL, the formulas $\forall x \varphi(x)$ and $\exists x \varphi(x)$ are primitive. The variable x ranges over the entire domain of the structure, whatever that domain happens to be. A sentence like $\forall x \varphi(x)$ is grammatically well-formed whether the domain contains three elements, three million, or infinitely many. FOL places no grammatical constraint on the domain.

This generality serves frameworks that accept unbounded or infinite domains. It does not serve settings where quantification is intended to range over explicitly bounded regions. Several such settings arise naturally: bounded arithmetic, where quantifiers are restricted to values below a polynomial bound (Parikh 1971; Buss 1986; Hájek and Pudlák 1993); finite model theory, where domains are finite and quantifier ranges carry structural significance (Ebbinghaus and Flum 1999; Libkin 2004); and foundational programmes that restrict attention to finite or bounded domains as a matter of principle.

In each of these settings, the common practice is to write bounded quantifier forms, $\forall x \leq t \varphi(x)$ and $\exists x \leq t \varphi(x)$, as abbreviations for standard FOL formulas. The abbreviation $\forall x \leq t \varphi(x)$ stands for $\forall x (x \leq 0 \vee t \leq 1 \vee \dots \vee 2(x))$, and $\exists x \leq t \varphi(x)$ stands for $\exists x (x \leq 7 \vee t \leq 8 \vee \dots \vee 9(x))$. This convention is effective and widespread. But it remains a convention: the underlying logic still admits unbounded quantifier forms as well-formed, and the bounded forms are syntactic sugar rather than primitive grammar.

This paper takes the bounded forms as primitive. Bounded First-Order Logic (BFOL) is a logic in which the only quantifier forms are $\forall x \leq t \varphi(x)$ and $\exists x \leq t \varphi(x)$. The unbounded forms $\forall x \varphi(x)$ and $\exists x \varphi(x)$ are not abbreviations for anything; they are not well-formed formulas of the language. The bound is not optional. It is part of the syntax of quantification.

The result is a logic that differs from standard FOL in exactly one structural respect, the primitive quantifier forms, and is identical in all others: term formation, atomic formulas, connectives, equality, and the propositional and equational proof rules. The modification is principled: it aligns the grammar of the logic with settings in which quantification is always over a bounded region. The paper establishes that this alignment preserves full metatheoretic strength (soundness, completeness, interpolation, and definability all survive) while producing a logic with distinctive behaviour when intended models are finite.

The paper proceeds as follows. Section 2 motivates the promotion of bounded quantifiers from abbreviation to primitive status. Section 3 defines the language. Sections 4 and 5 give the terms, formulas, and semantics. Section 6 presents two equivalent proof systems: natural de-

duction and a sequent calculus. Section 7 develops the core metatheory: soundness, decidability of truth in finite structures, completeness, cut-elimination, Craig interpolation via the Maehara partition theorem, and Beth definability. Section 8 introduces the Global Boundedness Principle and establishes the behaviour of BFOL when intended models are finite. Section 9 makes the relationship between BFOL and standard FOL precise. Section 10 concludes. Section 11 identifies future work.

2. Why Bounded Quantifiers as Primitive

A natural objection to the project of this paper is that it is unnecessary. If bounded quantifiers are definable in standard FOL, why promote them to primitive status? The abbreviation works. What is gained by the promotion?

Two things are gained, one syntactic and one semantic.

The syntactic gain is that BFOL enforces a discipline at the level of well-formedness. The question of which quantifier forms should be primitive in a logic has a history: Mostowski (1957) initiated the systematic study of generalized quantifiers, showing that the choice of primitive quantifier forms is not a merely notational decision but one with concrete consequences for expressive power. More broadly, Quine (1970) argued that the logical vocabulary of a theory is not innocent: it carries ontological commitments.

In standard FOL, a formula like $\forall x \varphi(x)$ is well-formed regardless of whether the intended application restricts x to a bounded range. A practitioner may intend bounded quantification, but the grammar permits unbounded quantification, and nothing in the formation rules prevents its accidental or deliberate use. In BFOL, the formation rules themselves require an explicit bound. An unbounded quantification cannot be written, not because it is prohibited by a meta-convention, but because it is not a formula of the language. The discipline is grammatical rather than conventional.

The semantic gain concerns the behaviour of the logic under finite-model restriction. Standard FOL, via the compactness theorem, guarantees that any theory with arbitrarily large finite models also has infinite models. This means that a theory intended to describe only finite structures cannot, in standard FOL, avoid commitment to infinite structures; the logic forces their existence as a consequence of its own metatheory. When bounded quantifiers are primitive and the intended class of models is restricted to finite structures (a restriction this paper formalises as the Global Boundedness Principle), the compactness argument does not apply, and the logic does not force infinite models. This is a structural consequence of the modified grammar. The precise mechanism is established in Section 8.

Neither gain requires rejecting standard FOL. BFOL embeds faithfully into FOL (Section 9), and every BFOL theorem is a theorem of FOL under translation. This paper develops BFOL as a logic suited to settings in which bounded quantification is the intended norm.

3. The Language of BFOL

BFOL is a first-order language; higher-order quantification is not part of the system. A language \mathcal{L} of BFOL consists of the following components.

Variables: x, y, z, \dots (a supply of variables, sufficient to ensure that every formula can be written with distinct bound variables).

Constant symbols: c, d, \dots (names for specific elements).

Function symbols: f, g, h, \dots each with a fixed arity (the number of arguments it takes).

Relation symbols: R, S, P, \dots each with a fixed arity.

The equality symbol: $=$.

A designated bounding relation: written \leq . This symbol plays a structural role in all bounded quantifier forms.

Logical connectives: $\neg, \wedge, \vee, \rightarrow$.

Bounded quantifier forms: $\forall x \leq t$ and $\exists x \leq t$, where t is a term. These are the only quantifier forms in the language. There are no primitive unbounded quantifiers.

Any specific BFOL language has a finite signature: finitely many constant symbols, finitely many function symbols, and finitely many relation symbols, together with equality and the bounding relation.

Remark on the bounding relation. BFOL does not legislate the interpretation of \leq beyond its structural role. It may be interpreted numerically (as the standard ordering on natural numbers), ordinally (as the ordering on finite ordinals), cardinally (as comparison of set sizes), or by any other ordering suited to the application. The specific interpretation is fixed by the theory built on top of BFOL. The logic constrains the quantifier forms; the theory fixes the interpretation of \leq .

4. Terms and Formulas

4.1 Terms

Terms are defined exactly as in ordinary FOL.

Term formation:

- (i) Every variable is a term.
- (ii) Every constant symbol is a term.
- (iii) If f is an n -ary function symbol and t_1, \dots, t_n are terms, then $f(t_1, \dots, t_n)$ is a term.

No modification to FOL's term formation is required. The restriction on quantification applies to formulas, not to terms.

4.2 Formulas

The formulas of BFOL are the class inductively generated by the following clauses.

Atomic formulas:

- (i) $t_1 = t_2$ is an atomic formula (for any terms t_1, t_2).
- (ii) $R(t_1, \dots, t_n)$ is an atomic formula (for any n -ary relation symbol R and terms t_i).

Compound formulas:

- (iii) If φ is a formula, so is $\neg\varphi$.
- (iv) If φ and ψ are formulas, so are $(\varphi \wedge \psi)$, $(\varphi \vee \psi)$, $(\varphi \rightarrow \psi)$.
- (v) If φ is a formula, x is a variable, and t is a term, then $\forall x \leq t \varphi$ and $\exists x \leq t \varphi$ are formulas.

No other formulas are admitted.

In particular: $\forall x \varphi$ and $\exists x \varphi$ (unbounded forms) are not well-formed formulas of BFOL.

Every part of formula formation is identical to standard FOL except the quantifier clause. Where FOL admits $\forall x \varphi$ and $\exists x \varphi$ as primitive, BFOL admits only $\forall x \leq t \varphi$ and $\exists x \leq t \varphi$. The bound t is not optional; it is part of the syntax of quantification. A variable cannot range without a named limit.

Examples of BFOL formulas:

$\forall x \leq c \ R(x)$	Well-formed.
$\exists x \leq f(y) \ (x = y)$	Well-formed. y is free in the bounding term $f(y)$.
$\forall x \leq c \ \exists y \leq x \ R(x, y)$	Well-formed. Nested bounded quantifiers; y is bounded by x , which is itself bounded by c .
$\forall x \ R(x)$	NOT well-formed in BFOL. No bounding term.
$\exists x \ (x = c)$	NOT well-formed in BFOL. No bounding term.

Note that in $\exists x \leq f(y) \ (x = y)$, the bounding term $f(y)$ contains the free variable y . The range of the bounded quantifier therefore depends on the value assigned to y . This is a structural feature of BFOL with no analogue in standard FOL, where quantifier ranges are always the full domain.

4.3 Sentences, free variables, and substitution

A sentence is a formula with no free variable occurrences. The standard definitions of free and bound occurrence, substitution, and variable capture avoidance carry over from FOL without modification, with the following understanding specific to BFOL: in the formula $\forall x \leq t \ \varphi$ (and likewise $\exists x \leq t \ \varphi$), the variable x is bound within φ , but any variables occurring in the bounding term t are free unless themselves bound by an outer quantifier. A consequence is that the free variables of a bounded quantification include those of the bounding term:

$$\begin{aligned} \text{FV}(\forall x \leq t \ \varphi) &= \text{FV}(t) \cup (\text{FV}(\varphi) \setminus \{x\}) \\ \text{FV}(\exists x \leq t \ \varphi) &= \text{FV}(t) \cup (\text{FV}(\varphi) \setminus \{x\}) \end{aligned}$$

Substitution respects this structure: substituting a term u for a free variable y in $\forall x \leq t \ \varphi$ requires substitution in both t and φ , with the usual capture-avoidance conditions on x .

5. Semantics

5.1 Structures and assignments

A BFOL-structure \mathfrak{M} for a language \mathcal{L} consists of:

- a domain D (a set),
- an interpretation of each constant symbol as an element of D ,
- an interpretation of each n -ary function symbol as a total function $D^n \rightarrow D$,
- an interpretation of each n -ary relation symbol as a subset of D^n , and
- an interpretation of \leq as a binary relation on D .

A variable assignment s maps each variable to an element of D .

Remark on the bounding relation in structures. No structural assumptions are imposed on \leq by the logic. It need not be reflexive, transitive, or total. BFOL's semantics is well-defined for any binary relation serving as the bound. Structural assumptions about \leq (such as being a total order) are imposed by the theory built on BFOL, not by the logic itself.

5.2 Term evaluation

The value of a term t in a structure \mathfrak{M} under an assignment s is defined recursively:

$x^{\mathfrak{M}, s}$	=	$s(x)$	for variables
$c^{\mathfrak{M}, s}$	=	$c^{\mathfrak{M}}$	for constants
$f(t_1, \dots, t_n)^{\mathfrak{M}, s}$	=	$f^{\mathfrak{M}}(t_1^{\mathfrak{M}, s}, \dots, t_n^{\mathfrak{M}, s})$	for function application

Term evaluation is unchanged from standard FOL.

5.3 Truth

Truth is defined relative to \mathfrak{M} and s by the following clauses.

Atomic formulas (unchanged from FOL):

$\mathfrak{M}, s \models t_1 = t_2$	iff	$t_1^{\mathfrak{M}, s} = t_2^{\mathfrak{M}, s}$
$\mathfrak{M}, s \models R(t_1, \dots, t_n)$	iff	$(t_1^{\mathfrak{M}, s}, \dots, t_n^{\mathfrak{M}, s}) \in R^{\mathfrak{M}}$

Connectives (unchanged from FOL):

$\mathfrak{M}, s \models \neg\phi$	iff	$\mathfrak{M}, s \not\models \phi$
$\mathfrak{M}, s \models \phi \wedge \psi$	iff	$\mathfrak{M}, s \models \phi$ and $\mathfrak{M}, s \models \psi$
$\mathfrak{M}, s \models \phi \vee \psi$	iff	$\mathfrak{M}, s \models \phi$ or $\mathfrak{M}, s \models \psi$
$\mathfrak{M}, s \models \phi \rightarrow \psi$	iff	$\mathfrak{M}, s \not\models \phi$ or $\mathfrak{M}, s \models \psi$

Bounded quantifiers:

$\mathfrak{M}, s \models \forall x \leq t \ \varphi$ iff for every $a \in D$ such that
 $a \leq^{\mathfrak{M}} t^{\{\mathfrak{M}, s\}}$,
 $\mathfrak{M}, s[x:=a] \models \varphi$

$\mathfrak{M}, s \models \exists x \leq t \ \varphi$ iff there exists $a \in D$ such that
 $a \leq^{\mathfrak{M}} t^{\{\mathfrak{M}, s\}}$ and
 $\mathfrak{M}, s[x:=a] \models \varphi$

Every quantified claim evaluates only over the region of the domain at or below the bounding term. The structure of truth-evaluation is identical to standard FOL except that the quantifier range is explicitly delimited. There is no clause for unbounded quantifiers because unbounded quantifiers do not exist in the language.

Vacuous quantification. If no element of D satisfies $a \leq^{\mathfrak{M}} t^{\{\mathfrak{M}, s\}}$, then $\forall x \leq t \ \varphi$ is vacuously true (the universal quantifier ranges over an empty set) and $\exists x \leq t \ \varphi$ is false (no witness exists). This follows from the standard reading of the clauses above and requires no special case.

5.4 Nonemptiness of the domain

In standard FOL, nonemptiness of the domain is imposed as a convention. In BFOL, it is a consequence of the grammar.

Proposition 0 — BFOL sentences require constants, and constants require domain elements.

- (i) Every BFOL sentence contains at least one constant symbol.

- (ii) Therefore, any structure interpreting a BFOL language in which sentences can be formed has a nonempty domain.

Proof of (i):

A BFOL sentence is a formula with no free variables. Every quantifier in a BFOL formula has the form $\forall x \leq t$ or $\exists x \leq t$, where t is a term. The variable x is bound by the quantifier, but any variables in t are free unless bound by an outer quantifier.

Consider the outermost quantifier in a sentence. Its bounding term t must be a closed term (no free variables, since the formula is a sentence and any variable in t that is not bound by an outer quantifier would be free in the sentence). A closed term is built from constant symbols and function symbols applied to closed terms. By induction on term structure, every closed term contains at least one constant symbol. Therefore the sentence contains at least one constant symbol.

Proof of (ii):

If the language contains a constant symbol c , then any structure \mathfrak{M} for this language must interpret c as an element of D . Therefore D is nonempty.

The domain is nonempty not by stipulation but because the grammar of bounded quantification requires bounding terms, bounding terms in sentences require constants, and constants require domain elements. Nonemptiness follows from the grammar of bounded quantification rather than from a semantic stipulation.

The logic itself does not require domains to be finite; BFOL is well-defined for any structure with a bounding relation. The restriction to finite domains, when desired, is imposed by the Global Boundedness Principle (Section 8).

6. Proof Systems

Two equivalent proof systems are developed: a natural deduction system and a sequent calculus. The natural deduction system is more intuitive for individual proofs. The sequent calculus is the vehicle for cut-elimination and the Maehara partition proof of Craig interpolation.

6.1 Natural deduction for BFOL

The natural deduction system retains all standard propositional and equality rules unchanged. The quantifier rules are replaced by bounded versions.

Bounded universal elimination (\forall -E):

From $\forall x \leq t \ \varphi(x)$ and a proof that $u \leq t$,
infer $\varphi(u)$.

Bounded universal introduction (\forall -I):

If $\varphi(x)$ has been derived under the assumption $x \leq t$,
with x not free in any undischarged assumption other
than $x \leq t$, infer $\forall x \leq t \ \varphi(x)$.

Bounded existential introduction (\exists -I):

From $\varphi(u)$ and a proof that $u \leq t$,
infer $\exists x \leq t \ \varphi(x)$.

Bounded existential elimination (\exists -E):

From $\exists x \leq t \ \varphi(x)$, and a derivation of ψ from the
assumptions $\varphi(x)$ and $x \leq t$ with x fresh and not free
in ψ or in any other undischarged assumption,
infer ψ .

There are no rules for unbounded universal or existential quantifiers because those forms are not in the language. Every proof in BFOL is a proof about explicitly bounded domains.

Example derivation. Suppose Γ contains $\forall x \leq c \ R(x)$ and the assumption $a \leq c$. One step of \forall -E yields $R(a)$. The derivation is:

- | | | |
|----|---------------------------|-------------------------|
| 1. | $\forall x \leq c \ R(x)$ | [from Γ] |
| 2. | $a \leq c$ | [assumption] |
| 3. | $R(a)$ | [\forall -E on 1, 2] |

This illustrates the pattern: bounded universal elimination requires both the universally quantified sentence and a proof that the instantiating term falls within the bound.

Example derivation (\exists -E). Suppose Γ contains $\exists x \leq c R(x)$ and $\forall y \leq c (R(y) \rightarrow S(y))$. We wish to derive $\exists x \leq c S(x)$. The derivation is:

1. $\exists x \leq c R(x)$ [from Γ]
2. $R(a), a \leq c$ [assumption, a fresh]
3. $\forall y \leq c (R(y) \rightarrow S(y))$ [from Γ]
4. $R(a) \rightarrow S(a)$ [\forall -E on 3, using $a \leq c$ from 2]
5. $S(a)$ [\rightarrow -E on 2, 4]
6. $\exists x \leq c S(x)$ [\exists -I on 5, using $a \leq c$ from 2]
7. $\exists x \leq c S(x)$ [\exists -E on 1, discharging 2]

The fresh variable a is introduced and then discharged. The bound $a \leq c$, carried as part of the existential elimination, serves two roles: it licenses the \forall -E instantiation at line 4 and the \exists -I step at line 6. Every quantifier in the derivation is explicitly bounded.

6.2 Sequent calculus for BFOL

The natural deduction system has an equivalent sequent calculus formulation, following the standard Gentzen-style presentation (Gentzen 1935; for modern treatments see Troelstra and Schwichtenberg 2000; Takeuti 1987). Sequents have the form $\Gamma \vdash \Delta$ where Γ and Δ are finite multisets of BFOL formulas. The rules are:

Identity: $A \vdash A$

Structural: Weakening (left and right), Contraction (left and right) – standard.

Propositional: (\wedge -L), (\wedge -R), (\vee -L), (\vee -R), (\rightarrow -L), (\rightarrow -R), (\neg -L), (\neg -R) – standard.

Bounded quantifier rules:

($\forall \leq$ -L): From $\Gamma, \varphi(s) \vdash \Delta$ and a proof that $s \leq t$, infer $\Gamma, \forall x \leq t \varphi(x) \vdash \Delta$.
(s is an arbitrary term; $s \leq t$ is a provable side condition, not a formula in the sequent.)

($\forall \leq$ -R): From $\Gamma, a \leq t \vdash \varphi(a), \Delta$ where a is a fresh eigenvariable not appearing in Γ, t , or Δ , infer $\Gamma \vdash \forall x \leq t \varphi(x), \Delta$.

(The assumption $a \leq t$ appears explicitly in the antecedent of the premise and is discharged by the rule.)

($\exists \leq$ -L): From $\Gamma, a \leq t, \varphi(a) \vdash \Delta$ where a is a fresh eigenvariable not appearing in Γ, t , or Δ , infer $\Gamma, \exists x \leq t \varphi(x) \vdash \Delta$.

(The assumption $a \leq t$ appears explicitly in the antecedent of the premise and is discharged by the rule.)

($\exists \leq$ -R): From $\Gamma \vdash \varphi(s), \Delta$ and a proof that $s \leq t$, infer $\Gamma \vdash \exists x \leq t \varphi(x), \Delta$.

(s is an arbitrary term; $s \leq t$ is a provable side condition.)

Cut: From $\Gamma \vdash A, \Delta$ and $\Gamma, A \vdash \Delta$, infer $\Gamma \vdash \Delta$.

The eigenvariable rules ($\forall \leq$ -R and $\exists \leq$ -L) place the bounding assumption $a \leq t$ in the antecedent of the premise sequent, where it is available for use in the subproof and discharged when the rule is applied. The substitution rules ($\forall \leq$ -L and $\exists \leq$ -R) carry $s \leq t$ as a side condition: a separate obligation that must be established, not a formula within the sequent. This treatment is uniform and makes the role of the bound explicit at every step.

6.3 Equivalence of the two systems

Proposition 1 — The natural deduction system and the sequent calculus derive the same theorems.

For any set of BFOL sentences Γ and BFOL sentence φ :

$\Gamma \vdash_{\text{ND}} \varphi$ if and only if $\Gamma \vdash_{\text{SC}} \varphi$

where \vdash_{ND} denotes derivability in the natural deduction system and \vdash_{SC} denotes derivability in the sequent calculus.

Proof:

The translation between natural deduction and sequent calculus is standard (see Troelstra and Schwichtenberg 2000, Ch. 3). Each natural deduction rule has a sequent calculus counterpart and vice versa. The translation preserves the bounded quantifier structure at every step: bounded quantifier rules in both systems correspond to the same semantic clauses (Section 5.3), and no unbounded quantifier is introduced by any translation step.

7. Metatheory

This section establishes the core metatheoretic properties of BFOL. The results show that BFOL has the standard metatheoretic properties of first-order logic, while possessing distinctive behaviour under finite-model restriction.

7.1 Soundness

Theorem 1 — BFOL Soundness.

If $\Gamma \vdash \varphi$ (that is, φ is derivable from assumptions Γ in the BFOL deductive system) then $\Gamma \models \varphi$: every BFOL-structure \mathfrak{M} and assignment s satisfying all sentences in Γ also satisfies φ .

Proof:

By induction on the length of the derivation. Each rule is verified against the semantics of Section 5.

Propositional rules: unchanged from standard FOL.

Soundness of \neg , \wedge , \vee , \rightarrow rules follows from the truth-table semantics, which are identical to FOL.

Equality rules: unchanged from FOL.

Bounded universal elimination (\forall -E):

Suppose $\mathfrak{M}, s \models \forall x \leq t \varphi(x)$ and $\mathfrak{M}, s \models u \leq t$.

By the semantics of $\forall x \leq t$: for every $a \in D$ with $a \leq^{\mathfrak{M}} t^{\{\mathfrak{M}, s\}}$, $\mathfrak{M}, s[x:=a] \models \varphi(x)$.

Since $u^{\{\mathfrak{M}, s\}} \leq^{\mathfrak{M}} t^{\{\mathfrak{M}, s\}}$, taking $a = u^{\{\mathfrak{M}, s\}}$ gives $\mathfrak{M}, s \models \varphi(u)$.

Bounded universal introduction (\forall -I):

Suppose $\varphi(x)$ has been derived from the assumption $x \leq t$, with x not free in any other undischarged assumption. By the induction hypothesis, every \mathfrak{M}, s satisfying $x \leq t$ also satisfies $\varphi(x)$. That is, for every $a \in D$ with $a \leq^{\mathfrak{M}} t^{\{\mathfrak{M}, s\}}$,

$\mathfrak{M}, s[x:=a] \models \varphi(x)$. This is exactly the semantics of $\forall x \leq t \varphi(x)$.

Bounded existential introduction (\exists -I):

Suppose $\mathfrak{M}, s \models \varphi(u)$ and $\mathfrak{M}, s \models u \leq t$.

Then $a = u^{\mathfrak{M}, s}$ witnesses the existential:

$a \leq^{\mathfrak{M}} t^{\mathfrak{M}, s}$ and $\mathfrak{M}, s[x:=a] \models \varphi(x)$.

So $\mathfrak{M}, s \models \exists x \leq t \varphi(x)$.

Bounded existential elimination (\exists -E):

Suppose $\mathfrak{M}, s \models \exists x \leq t \varphi(x)$, and ψ has been derived from the assumptions $\varphi(x)$ and $x \leq t$ with x fresh.

By the semantics, there exists $a \in D$ with

$a \leq^{\mathfrak{M}} t^{\mathfrak{M}, s}$ and $\mathfrak{M}, s[x:=a] \models \varphi(x)$. Since x is

fresh (not free in ψ or in other undischarged assumptions), the derivation of ψ from $\varphi(x)$ and

$x \leq t$ holds for this witness. By the induction

hypothesis, $\mathfrak{M}, s \models \psi$.

Every rule preserves truth. Therefore every derivable sentence is true in every structure satisfying its assumptions.

The BFOL deductive system never leads from true assumptions to a false conclusion. If something is provable in BFOL, it is true in every BFOL-structure that satisfies the assumptions.

7.2 Decidability of truth in finite structures

Theorem 2 — Decidability of BFOL truth in finite structures.

For any fixed finite BFOL-structure \mathfrak{M} and any BFOL sentence φ , the question " $\mathfrak{M} \models \varphi$?" is decidable.

Proof:

By induction on formula complexity.

Atomic formulas: $t_1 = t_2$ and $R(t_1, \dots, t_n)$ are checked by computing the term values in the finite domain D and performing a lookup. Decidable.

Boolean connectives: \neg , \wedge , \vee , \rightarrow are evaluated by truth tables from the values of their immediate subformulas. Decidable by the induction hypothesis.

Bounded quantifiers: $\forall x \leq t \varphi(x)$ is evaluated by checking $\varphi(a)$ for every $a \in D$ with $a \leq^{\mathfrak{M}} t^{\mathfrak{M}, s}$. This is a finite set, a subset of the finite domain D , so the check is a finite conjunction. $\exists x \leq t \varphi(x)$ is similarly a finite disjunction. Decidable by the induction hypothesis applied to each element of the bounded range.

The total computation terminates in time bounded by a function of $|D|$ and the length of φ .

In any single finite model, truth is checkable by exhaustive computation. Every quantifier ranges over a finite set, every connective is a truth table, and every atomic formula is a finite lookup. This is the property that makes finite models computationally transparent.

7.3 Completeness

Theorem 3 — BFOL Completeness.

If φ is a BFOL sentence true in every BFOL-structure (finite or infinite) satisfying assumptions Γ , then $\Gamma \vdash \varphi$.

Proof:

Every BFOL formula embeds into standard FOL via the translation:

$$\forall x \leq t \varphi(x) \mapsto \forall x (x \leq t \rightarrow \varphi(x))$$

$$\exists x \leq t \varphi(x) \mapsto \exists x (x \leq t \wedge \varphi(x))$$

This embedding preserves truth: a BFOL-structure \mathfrak{M} satisfies a BFOL sentence φ if and only if \mathfrak{M} (viewed as a FOL-structure) satisfies the FOL translation φ^* .

If φ is true in every BFOL-structure satisfying Γ , then φ^* is true in every FOL-structure satisfying Γ^* . By Gödel's completeness theorem for FOL (Gödel 1930), $\Gamma^* \vdash_{\text{FOL}} \varphi^*$. The FOL proof translates back to a BFOL proof, since every FOL rule used on bounded formulas corresponds to a BFOL rule. Therefore $\Gamma \vdash \varphi$.

Corollary 3a — Soundness-completeness biconditional.

$\Gamma \vdash \varphi$ if and only if $\Gamma \models \varphi$.

That is, φ is derivable from Γ in BFOL if and only if φ is true in every BFOL-structure satisfying Γ .

Proof: The left-to-right direction is Theorem 1 (Soundness). The right-to-left direction is Theorem 3 (Completeness).

The deductive system captures exactly the valid entailments of the logic. Nothing true in all structures escapes provability; nothing provable fails in any structure.

Remark on completeness, decidability, and finite models. The completeness theorem above quantifies over all BFOL-structures, finite and infinite. Completeness does not imply decidability: the set of valid sentences of standard FOL is not decidable (Church 1936), and BFOL inherits this limitation via the embedding of Section 9.1. A stronger property would be completeness restricted to finite structures: if φ is true in every finite BFOL-structure, then $\text{BFOL} \vdash \varphi$. This stronger property fails for a deeper reason. Trakhtenbrot’s theorem (Theorem 9, Section 8) establishes that the set of sentences true in all finite structures of a language with a binary relation is not recursively enumerable, while the set of provable sentences in any sound system is. Finite-restricted completeness is therefore impossible for any sound deductive system, including BFOL. The computational classification is strictly higher in the finite-model case: FOL validity is undecidable but recursively enumerable (Church); finite-model validity is not recursively enumerable (Trakhtenbrot).

7.4 Cut-elimination

Theorem 4 — Cut-elimination for BFOL.

Every proof in the BFOL sequent calculus can be transformed into a cut-free proof of the same sequent.

Proof:

The proof follows the standard Gentzen cut-elimination argument (Gentzen 1935; see also Troelstra and Schwichtenberg 2000, Ch. 4; Takeuti 1987, Ch. 1).

Each reduction step (permuting the cut above a rule application, or reducing a cut on a formula introduced on both sides) preserves the bounded quantifier structure because every rule takes BFOL formulas as input and produces BFOL formulas as output. No unbounded quantifier is introduced at any stage. The transformation terminates by the standard measure (cut rank and proof height).

Any proof that uses the cut rule (reasoning via an intermediate result that is then discarded) can be converted into one that proceeds directly, without cut. This is the structural property that makes the Maehara partition construction (Theorem 5a below) possible, and it guarantees that the subformula property holds for cut-free BFOL proofs: every formula appearing in a cut-free proof is a subformula of the conclusion.

7.5 Craig interpolation

Theorem 5 — Bounded Craig Interpolation.

The classical Craig interpolation theorem (Craig 1957) asserts that if φ entails ψ in FOL, there exists an interpolant using only the shared vocabulary. The following is the bounded analogue.

Let φ and ψ be BFOL sentences such that $\varphi \vdash \psi$.
 Let $\text{Var}(\varphi) \cap \text{Var}(\psi)$ be the shared non-logical vocabulary (relation symbols, function symbols, and constant symbols appearing in both φ and ψ ; the logical symbols $=, \leq, \neg, \wedge, \vee, \rightarrow$, and the bounded quantifiers are always shared).

Then there exists a BFOL formula θ such that:

- (i) $\varphi \vdash \theta$
- (ii) $\theta \vdash \psi$
- (iii) The non-logical symbols of θ are among those common to φ and ψ .
- (iv) Every quantifier in θ is bounded.

The proof proceeds by Maehara's method (Maehara 1960), applied directly to the sequent calculus formulation of BFOL. This avoids the round-trip through FOL and ensures the interpolant is a BFOL formula at every stage of the construction.

Theorem 5a — Maehara Partition for BFOL.

The interpolation theorem is derived from a stronger partition theorem on cut-free proofs.

Let π be a cut-free proof of $\Gamma \vdash \Delta$ in the BFOL sequent calculus. Let $(\Gamma_1, \Delta_1 \mid \Gamma_2, \Delta_2)$ be any partition of the end-sequent, meaning $\Gamma = \Gamma_1, \Gamma_2$ and $\Delta = \Delta_1, \Delta_2$.

Then there exists a BFOL formula θ such that:

- (a) $\Gamma_1 \vdash \theta, \Delta_1$ is provable
- (b) $\Gamma_2, \theta \vdash \Delta_2$ is provable
- (c) The non-logical symbols of θ are among those common to (Γ_1, Δ_1) and (Γ_2, Δ_2) .
- (d) Every quantifier in θ is bounded.

Proof:

By induction on the height of the cut-free proof π .

BASE CASE – Identity axiom:

π proves $A \vdash A$ for some formula A .

Case: A is in Γ_1 and A is in Δ_1 .

Take $\theta = \top$. Then $\Gamma_1 \vdash \top$, Δ_1 is immediate, and $\Gamma_2, \top \vdash \Delta_2$ holds since \top is trivially eliminable.

Case: A is in Γ_1 and A is in Δ_2 .

Take $\theta = A$. Then $A \vdash A$ is the identity (giving $\Gamma_1 \vdash \theta$, Δ_1), and $A \vdash A$ gives $\Gamma_2, \theta \vdash \Delta_2$.

A uses only its own symbols, which appear on both sides of the partition.

Case: A is in Γ_2 and A is in Δ_1 .

Take $\theta = \neg A$. Then $\Gamma_1 \vdash \neg A$, A (which is $\Gamma_1 \vdash \theta$, Δ_1) holds by \neg -R from the identity, and $A, \neg A \vdash \Delta_2$ (which is $\Gamma_2, \theta \vdash \Delta_2$) holds by \neg -L.

Case: A is in Γ_2 and A is in Δ_2 .

Take $\theta = \top$ and argue as in the first case.

PROPOSITIONAL CASES:

The propositional rules (\wedge -L), (\wedge -R), (\vee -L), (\vee -R), (\rightarrow -L), (\rightarrow -R), (\neg -L), (\neg -R) are handled exactly as in the standard FOL Maehara proof. In each case:

- If the principal formula is on side 1: apply the rule to the first component, θ passes through unchanged.
- If the principal formula is on side 2: apply the rule to the second component, θ passes through unchanged.
- For two-premise rules (\wedge -R, \rightarrow -L): the IH gives interpolants θ_1 and θ_2 for the two premises; take $\theta = \theta_1 \wedge \theta_2$ or $\theta = \theta_1 \vee \theta_2$ as appropriate.

Since boolean combinations of BFOL formulas are BFOL formulas, conditions (c) and (d) are preserved.

BOUNDED QUANTIFIER CASES – substitution rules:

($\forall \leq$ -L): From $\Gamma, \varphi(s) \vdash \Delta$ with $s \leq t$, infer $\Gamma, \forall x \leq t \varphi(x) \vdash \Delta$. The principal formula $\forall x \leq t \varphi(x)$ is on the LEFT.

If $\forall x \leq t \varphi(x)$ is on side 1: In the premise, $\varphi(s)$ replaces $\forall x \leq t \varphi(x)$ on side 1. By IH, $\exists \theta$ with $\Gamma_1', \varphi(s) \vdash \theta, \Delta_1$ and $\Gamma_2, \theta \vdash \Delta_2$.
Apply ($\forall \leq$ -L) to the first sequent:
 $\Gamma_1', \forall x \leq t \varphi(x) \vdash \theta, \Delta_1$. Second unchanged.
Interpolant θ , unchanged, still BFOL.

If $\forall x \leq t \varphi(x)$ is on side 2: symmetric – apply ($\forall \leq$ -L) to the second component.

($\exists \leq$ -R): From $\Gamma \vdash \varphi(s), \Delta$ with $s \leq t$, infer $\Gamma \vdash \exists x \leq t \varphi(x), \Delta$. Symmetric to ($\forall \leq$ -L).
Apply the rule to whichever side contains the principal formula. θ passes through unchanged.

BOUNDED QUANTIFIER CASES – eigenvariable rules:

These are the critical cases where the interpolant may acquire a new quantifier.

($\forall \leq$ -R): From $\Gamma, a \leq t \vdash \varphi(a), \Delta$ with a fresh, infer $\Gamma \vdash \forall x \leq t \varphi(x), \Delta$. The principal formula $\forall x \leq t \varphi(x)$ is on the RIGHT.

If $\forall x \leq t \varphi(x)$ is on side 1: In the premise, $\varphi(a)$ is on side 1 and $a \leq t$ is in the antecedent.
By IH, $\exists \theta$ with $\Gamma_1, a \leq t \vdash \theta, \varphi(a), \Delta_1'$ and $\Gamma_2, \theta \vdash \Delta_2$.
Since a is fresh (eigenvariable condition), a does

not appear in $\Gamma_1, \Delta_1', \Gamma_2, \Delta_2$. By the subformula property of cut-free proofs, θ is built from subformulas of the sequent; since a appears only in $\varphi(a)$ which is on side 1, a does not appear in θ (θ uses only shared vocabulary, and a is not a non-logical symbol but a fresh variable not in any original formula). Apply $(\forall \leq\text{-R})$ to bind a in the first sequent: $\Gamma_1 \vdash \theta, \forall x \leq t \varphi(x), \Delta_1'$.
 Second unchanged. θ unchanged, still BFOL.

If $\forall x \leq t \varphi(x)$ is on side 2: In the premise, $\varphi(a)$ is on side 2 and $a \leq t$ is in the antecedent. By IH, $\exists \theta$ with $\Gamma_1 \vdash \theta, \Delta_1$ and $\Gamma_2, a \leq t, \theta \vdash \varphi(a), \Delta_2'$.

Sub-case: $a \notin \text{FV}(\theta)$. Then a does not appear in the interpolant. Apply $(\forall \leq\text{-R})$ to the second sequent (a is fresh, not in $\Gamma_2, \theta, \Delta_2'$): $\Gamma_2, \theta \vdash \forall x \leq t \varphi(x), \Delta_2'$.
 θ unchanged, still BFOL.

Sub-case: $a \in \text{FV}(\theta)$. The eigenvariable a has propagated into the interpolant through shared formulas in the proof tree. Since a satisfies $a \leq t$ and the bound term t appears in the principal formula $\forall x \leq t \varphi(x)$ on side 2, and a entered θ only through formulas using shared vocabulary, the symbols of t must be in the shared vocabulary. The contrapositive makes this clear: if any symbol of t were non-shared, it would appear only on side 2; but θ uses only shared vocabulary by the induction hypothesis; therefore a could not have entered θ through a term involving non-shared symbols. So t is expressible in the shared vocabulary, and the bounded quantifier $\exists y \leq t$ can be formed using only

shared symbols.

Construct $\theta' = \exists y \leq t \theta(y/a)$ (existentially bind a with the available bound t , using a fresh variable y).

First sequent: From $\Gamma_1 \vdash \theta(a), \Delta_1$, since $a \leq t$ and $\theta(a)$ holds, apply ($\exists \leq$ -R):
 $\Gamma_1 \vdash \exists y \leq t \theta(y), \Delta_1$.

Second sequent: From $\Gamma_2, \theta(a) \vdash \varphi(a), \Delta_2'$, rename a to a fresh variable b throughout. Apply ($\exists \leq$ -L) on the left: from $\Gamma_2, \theta(b) \vdash \varphi(b), \Delta_2'$ with b fresh, derive $\Gamma_2, \exists y \leq t \theta(y) \vdash \varphi(b), \Delta_2'$. Then apply ($\forall \leq$ -R) on the right: b is fresh and not in $\Gamma_2, \exists y \leq t \theta(y), \Delta_2'$, giving $\Gamma_2, \exists y \leq t \theta(y) \vdash \forall x \leq t \varphi(x), \Delta_2'$.

The interpolant $\theta' = \exists y \leq t \theta(y)$ is a BFOL formula: its quantifier is bounded by t .

($\exists \leq$ -L): From $\Gamma, a \leq t, \varphi(a) \vdash \Delta$ with a fresh, infer $\Gamma, \exists x \leq t \varphi(x) \vdash \Delta$. The principal formula $\exists x \leq t \varphi(x)$ is on the LEFT. Symmetric to ($\forall \leq$ -R):

If on side 1 and $a \in FV(\theta)$: construct $\theta' = \forall y \leq t \theta(y/a)$ (universally bind a). The argument mirrors the $\forall \leq$ -R case with \forall/\exists and left/right exchanged.

If on side 2, or if $a \notin FV(\theta)$: apply the rule to the appropriate component. θ passes through unchanged.

STRUCTURAL CASES – weakening and contraction:

Weakening: the interpolant from the IH is carried

through unchanged.

Contraction: the interpolant from the IH is carried through unchanged (the contracted formula stays on its side).

This completes the induction. In every case, the interpolant θ is a BFOL formula (all quantifiers are bounded) and uses only the shared non-logical vocabulary.

Derivation of Craig Interpolation from the Partition Theorem.

Given $\varphi \vdash \psi$, take the cut-free proof of $\varphi \vdash \psi$ (which exists by cut-elimination, Theorem 4). Apply the Partition Theorem with $\Gamma_1 = \{\varphi\}$, $\Delta_1 = \emptyset$, $\Gamma_2 = \emptyset$, $\Delta_2 = \{\psi\}$.

The theorem gives θ with:

- (a) $\varphi \vdash \theta$
- (b) $\theta \vdash \psi$
- (c) Non-logical symbols of $\theta \subseteq \text{Var}(\varphi) \cap \text{Var}(\psi)$
- (d) All quantifiers in θ are bounded.

This is Theorem 5.

If a BFOL argument leads from hypotheses using vocabulary V_1 to a conclusion using vocabulary V_2 , there is an intermediate BFOL statement using only the shared vocabulary $V_1 \cap V_2$ that bridges the two. The interpolant is constructed directly from the proof tree by Maehara's method: at each node, a bounded formula is built from the formulas at that node using boolean combinations and bounded quantification over eigenvariables. Because the BFOL proof system admits only bounded quantifier rules, the interpolant inherits bounded quantifiers at every stage. This ensures that BFOL proofs can be decomposed into bounded sub-arguments with well-defined shared interfaces.

The eigenvariable cases are the central cases of the proof and the point where BFOL interpolation departs from a translation of the standard FOL result. When an eigenvariable propagates into the interpolant, it must be bound, and the bound must be a BFOL term. The construction uses the bound term t from the quantifier rule that introduced the eigenvariable. This term is available precisely because BFOL quantifiers carry their bounds syntactically. In standard FOL, the analogous step binds the eigenvariable with an unrestricted quantifier; in BFOL, the quantifier is bounded, and the interpolant remains a BFOL formula. The bounded structure of the interpolant is a consequence of the bounded quantifier design, not a coincidence of the proof

technique.

7.6 Beth definability

Theorem 6 — Bounded Beth Definability.

Beth's definability theorem (Beth 1953) establishes that in standard FOL, implicit definability implies explicit definability. The following is the bounded analogue.

Let T be a BFOL theory and R a relation symbol not in T . Suppose T implicitly defines R : for any two models \mathfrak{M} and \mathfrak{M}' of T that agree on all symbols except R ,

$$T \cup T' \models \forall \vec{x} \leq t (R(\vec{x}) \leftrightarrow R'(\vec{x}))$$

where T' is a notational copy of T with R replaced by R' , and t is any bounding term in the shared vocabulary.

Then T explicitly defines R : there exists a BFOL formula $\theta(\vec{x})$ such that

$$T \models \forall \vec{x} \leq t (R(\vec{x}) \leftrightarrow \theta(\vec{x})).$$

Every quantifier in θ is bounded.

Proof: Immediate from Bounded Craig Interpolation (Theorem 5) by the standard derivation of Beth from Craig. The implicit definition is expressed in BFOL, so the interpolant produced by Theorem 5 is a BFOL formula with all quantifiers bounded. This interpolant serves as the explicit definition θ .

If a relation is uniquely determined by a BFOL theory, it can be explicitly expressed by a bounded formula. Implicit definability and explicit definability coincide in BFOL, just as they do in standard FOL.

8. The Global Boundedness Principle and Finite-Model Behaviour

BFOL's syntax and semantics are defined for any structure with a bounding relation. The logic is well-defined over finite and infinite domains alike, and Theorems 1–6 hold without restriction on domain size. This section introduces a metatheoretic semantic restriction, the Global Boundedness Principle, and examines what happens to BFOL's metatheory when it is imposed.

The GBP is not part of BFOL's syntax or proof theory. It is a condition on which structures are admitted as intended models. Its content is as follows:

Global Boundedness Principle (GBP):
Every intended model has a finite domain.

The separation between BFOL as a general logic and GBP as a semantic restriction is deliberate. BFOL is a logic of bounded quantification, applicable in any setting where explicit bounds are desired: bounded arithmetic, database theory, verification, or foundational programmes. GBP is a specific further commitment: that the domains of interest are finite. Other users of BFOL may adopt different semantic restrictions, or none at all.

Under GBP, two classical metatheorems of standard FOL fail. Both the compactness theorem and the Löwenheim-Skolem theorems are foundational results of classical model theory (see Chang and Keisler 1990, Ch. 2–3 for the standard development). In both cases, the failure is a consequence of the finite-domain restriction, not of the bounded quantifier syntax. It is GBP that does the work, not the bounded quantifier syntax; BFOL without GBP embeds into FOL and inherits compactness. It is specifically the restriction to finite intended models that breaks both properties.

8.1 Failure of compactness

Proposition 7 — Failure of Compactness for BFOL under GBP.

There exist sets Γ of BFOL sentences such that every finite subset of Γ has a finite model, but Γ itself has no finite model.

Proof:

Let the language contain a constant symbol c and the bounding relation \leq . For each natural number $n \geq 1$, let σ_n be the BFOL sentence asserting "there exist at least n distinct elements at or below c ":

$$\sigma_1: \exists x_1 \leq c (x_1 = x_1)$$

$$\sigma_2: \exists x_1 \leq c \exists x_2 \leq c (x_1 \neq x_2)$$

$$\sigma_3: \exists x_1 \leq c \exists x_2 \leq c \exists x_3 \leq c$$

$$(x_1 \neq x_2 \wedge x_1 \neq x_3 \wedge x_2 \neq x_3)$$

and so on. Each σ_n is a BFOL sentence: all quantifiers are bounded by c , and c is a constant symbol (satisfying Proposition 0).

Let $\Gamma = \{\sigma_n : n \geq 1\}$. Every finite subset of Γ is satisfied by a sufficiently large finite model (interpret c as an element with at least $\max(n)$ elements at or below it). But any model satisfying all of Γ simultaneously must contain elements of every finite cardinality at or below c , and therefore has an infinite domain.

Under GBP, only finite models are admitted. No finite model satisfies Γ . Therefore compactness fails: finite satisfiability of every subset does not entail satisfiability in the GBP-restricted class.

The compactness argument (“if every finite piece of a theory has a model, the whole theory has a model”) relies on the availability of infinite models as limits of finite approximations. GBP excludes infinite models. With infinite models excluded, a theory can be finitely consistent without being satisfiable. This is the structural reason that a theory formulated in BFOL under GBP can have models of every finite size without being forced into infinite ones.

Remark on the status of Γ . The set $\Gamma = \{\sigma_n : n \geq 1\}$ is an infinite collection of sentences. It is a metatheoretic construction used to demonstrate the failure of compactness, not an object of BFOL or of any theory built on BFOL. The proof is conducted in the metatheory, where infinite collections are available as notational tools. No BFOL theory asserts the existence of Γ as a completed totality.

8.2 Failure of Löwenheim-Skolem

Proposition 8 — Failure of Löwenheim-Skolem for BFOL under GBP.

Neither the downward nor the upward Löwenheim-Skolem theorem holds for BFOL under GBP.

Proof:

Downward: The downward Löwenheim-Skolem theorem asserts that any satisfiable theory in a countable language has a countably infinite model. Under GBP, all intended models are finite. The theorem requires an infinite model to exist as a starting point (from which a countably infinite elementary submodel is extracted). Since GBP excludes infinite models, the hypothesis of the downward theorem is never satisfied in the intended model class.

Upward: The upward Löwenheim-Skolem theorem asserts that any theory with an infinite model has models of every larger cardinality. Under GBP, no intended model is infinite. The hypothesis is never satisfied.

Löwenheim-Skolem guarantees the existence of models at every infinite cardinality. Under GBP, infinite models do not exist. The theorem's conclusions are vacuously blocked, and its machinery, which depends on the availability of infinite structures, has no ground to operate on.

8.3 Decidability versus undecidability: the Trakhtenbrot contrast

Under GBP, a contrast emerges between truth in individual finite models and truth across all finite models.

Theorem 2 established that truth in any single fixed finite structure is decidable: every quantifier ranges over a finite set, and evaluation terminates. A natural question is whether truth in all finite structures is similarly tractable. It is not.

Theorem 9 — Undecidability of finite-model validity (Trakhtenbrot).

For any first-order language containing at least one binary relation symbol, the set of sentences that are finitely valid (true in all finite structures) is not recursively enumerable.

This is Trakhtenbrot's Theorem (Trakhtenbrot 1950; see Libkin 2004, Ch. 9 and Ebbinghaus and Flum 1999, Ch. 7 for modern treatments). The result applies to BFOL under GBP: since any BFOL language may contain a binary relation symbol, the set of BFOL sentences true in all finite models is not recursively enumerable.

The contrast is worth stating as a highlighted observation:

Single-model truth:	Decidable (Theorem 2).
All-finite-model truth:	Not recursively enumerable (Theorem 9).

Truth in any one finite structure can be computed by bounded search. But the question "is this sentence true in every finite structure?" is not merely undecidable; it is not even semi-decidable. No algorithm can enumerate all finitely valid sentences.

This has a direct consequence for the completeness question raised in the remark following Theorem 3. Completeness restricted to finite models (if φ is true in every finite BFOL-structure, then $\text{BFOL} \vdash \varphi$) is impossible for any sound deductive system, because the set of provable sentences is always recursively enumerable, while the set of finitely valid sentences is not. The gap between "provable" and "true in all finite models" is structural and cannot be closed by a stronger proof system.

Remark. The contrast between single-model decidability and all-finite-model undecidability is characteristic of finite model theory generally (see Ebbinghaus and Flum 1999; Libkin 2004; and Fagin 1974 for the foundational connection between finite-model logic and computational complexity). It is not specific to BFOL. What is specific to BFOL is that the contrast acquires foundational significance when BFOL is used as the logical substrate for theories whose intended models are all finite. In that setting, the gap between provability and finite-model truth becomes the precise location of the theory's incompleteness, and the theory must address it. The logic itself does not resolve that question; it identifies it.

9. Relationship to Standard FOL

BFOL differs from standard FOL in exactly one structural respect: the primitive quantifier forms. Everything else is identical: term formation, atomic formulas, connectives, equality, proof rules for propositional and equational reasoning. BFOL is not a rejection of first-order logic. It is a restriction on which quantifier forms are admissible.

Standard FOL:	$\forall x \varphi(x)$ and $\exists x \varphi(x)$ are primitive. Domain quantification is unrestricted.
BFOL:	$\forall x \leq t \varphi(x)$ and $\exists x \leq t \varphi(x)$ are primitive. Domain quantification is always bounded. Unbounded forms are not well-formed.

9.1 The embedding

Any BFOL formula can be embedded into standard FOL. The translation $\varphi \mapsto \varphi^*$ is defined recursively on formula structure:

Atomic formulas:	$(t_1 = t_2)^*$ = $t_1 = t_2$ $R(t_1, \dots, t_n)^*$ = $R(t_1, \dots, t_n)$
Boolean:	$(\neg \varphi)^*$ = $\neg(\varphi^*)$ $(\varphi \wedge \psi)^*$ = $\varphi^* \wedge \psi^*$ $(\varphi \vee \psi)^*$ = $\varphi^* \vee \psi^*$ $(\varphi \rightarrow \psi)^*$ = $\varphi^* \rightarrow \psi^*$
Bounded quantifiers:	$(\forall x \leq t \varphi)^*$ = $\forall x (x \leq t \rightarrow \varphi^*)$ $(\exists x \leq t \varphi)^*$ = $\exists x (x \leq t \wedge \varphi^*)$

The translation replaces each bounded quantifier with its standard FOL expansion and leaves everything else unchanged.

Proposition 2 — Truth preservation of the embedding.

For any BFOL-structure \mathfrak{M} and assignment s , and any BFOL formula φ :

$$\mathfrak{M}, s \models_{\text{BFOL}} \varphi \quad \text{if and only if} \quad \mathfrak{M}, s \models_{\text{FOL}} \varphi^*$$

where \mathfrak{M} is viewed as a FOL-structure in the right-hand evaluation.

Proof: By induction on formula complexity. The atomic and boolean cases are immediate (the translation is the identity). The bounded quantifier cases follow from the semantic equivalence:

$$\begin{aligned} \mathfrak{M}, s \models \forall x \leq t \varphi & \\ \text{iff for every } a \in D \text{ with } a \leq^{\mathfrak{M}} t^{\{\mathfrak{M}, s\}}, & \\ \mathfrak{M}, s[x:=a] \models \varphi & \\ \text{iff for every } a \in D, & \\ a \leq^{\mathfrak{M}} t^{\{\mathfrak{M}, s\}} \text{ implies } \mathfrak{M}, s[x:=a] \models \varphi & \\ \text{iff } \mathfrak{M}, s \models \forall x (x \leq t \rightarrow \varphi^*) & \end{aligned}$$

and similarly for the existential.

In the other direction, not every FOL formula translates into BFOL. Unbounded quantifications $\forall x \varphi(x)$ and $\exists x \varphi(x)$ have no BFOL counterpart. The embedding is one-directional.

9.2 Expressive power

The embedding of §9.1 is one-directional: every BFOL formula has a FOL translation, but not every FOL formula has a BFOL counterpart. BFOL is therefore strictly less expressive than standard FOL. It can say less.

This reduction in expressiveness is the intended consequence of the design. The sentences that BFOL cannot express are exactly those whose semantic evaluation requires ranging over a completed, potentially unbounded domain. In a setting where domains are bounded, those sentences do not have a well-defined intended reading. Their absence from the language is not a limitation on what can be meaningfully said; it is the exclusion of forms whose semantic content exceeds what the intended structures can support. A logic that cannot express claims it cannot evaluate is more precise than one that can.

9.3 Why the embedding does not make BFOL redundant

A natural question: if BFOL embeds into FOL, why not simply use FOL with the convention of writing bounded quantifiers as abbreviations?

The embedding establishes that BFOL and FOL agree on truth-conditions for bounded formulas. It does not establish that the two logics have the same grammar, the same admissible formulas, or the same behaviour under semantic restriction.

Grammatically, FOL admits unbounded quantifier forms as well-formed alongside bounded ones. BFOL does not. A practitioner working in FOL may inadvertently write or derive an unbounded formula; a practitioner working in BFOL cannot. The discipline is enforced by the language, not by convention.

Semantically, the embedding does not preserve the behaviour under GBP. Standard FOL's compactness theorem implies that any theory with arbitrarily large finite models also has infinite models. BFOL under GBP is not subject to this implication; the failure of compactness for finite intended models is established in Proposition 7. This is the structural reason a theory formulated in BFOL under GBP can have models of every finite size without being forced to admit infinite ones. The embedding into FOL does not carry this property, because the embedded theory inherits FOL's compactness.

The two logics are therefore not interchangeable in all settings. For applications that do not restrict to finite models, the embedding suffices and BFOL differs from FOL only in notational discipline. For applications where finite-domain restriction is essential, BFOL under GBP has structural properties that the FOL embedding does not reproduce.

9.4 Summary of the metatheoretic landscape

Property	BFOL (general)	BFOL + GBP
Soundness	Yes (Thm 1)	Yes
Completeness (all models)	Yes (Thm 3)	Yes
Completeness (finite only)	N/A	No (Thm 9)
Single-model decidability	Yes (Thm 2)	Yes
All-model decidability	No (Church 1936)	No (Thm 9)
Cut-elimination	Yes (Thm 4)	Yes
Craig interpolation	Yes (Thm 5)	Yes
Beth definability	Yes (Thm 6)	Yes
Compactness	Yes (via FOL)	No (Prop 7)
Löwenheim-Skolem	Yes (via FOL)	No (Prop 8)

The table makes the two-layer structure visible. BFOL as a general logic inherits, via the

FOL embedding, all classical metatheoretic properties, including the undecidability of validity (Church 1936). BFOL under GBP retains the proof-theoretic properties (soundness, completeness over all models, cut-elimination, interpolation, definability) but loses the model-theoretic properties that depend on the availability of infinite models (compactness, Löwenheim-Skolem, finite-model completeness). The losses are consequences of the finite-domain restriction, not of the bounded quantifier syntax.

10. Conclusion

Bounded First-Order Logic is a principled restriction of first-order logic in which every quantifier carries an explicit bounding term. The modification changes one thing, the primitive quantifier forms, but the consequences of that single change are substantial. The grammar forces sentences to contain constants (Proposition 0), which in turn forces nonempty domains without stipulation. The resulting logic supports the full suite of standard metatheoretic properties: soundness, completeness, cut-elimination, Craig interpolation (with a direct Maehara proof preserving bounded quantifier structure), and Beth definability. Reduced expressiveness relative to FOL is an intended consequence of the design: BFOL cannot express claims whose evaluation requires ranging over unbounded domains, and in settings where domains are bounded, those claims have no well-defined intended reading.

The logic acquires distinctive character under the Global Boundedness Principle, which restricts intended models to finite domains. In this setting, compactness and Löwenheim-Skolem fail. FOL validity is undecidable but recursively enumerable (Church 1936); finite-model validity is not recursively enumerable (Trakhtenbrot 1950). The contrast between single-model decidability and all-finite-model non-enumerability is a structural feature of the logic under GBP, and it identifies the precise location where any theory built on BFOL with finite intended models must confront its own incompleteness.

BFOL is suited to serve as a logical substrate for any formal theory whose quantification is intended to range over bounded domains. The theories that may be built on it are not the subject of this paper. This paper establishes only the logic: its syntax, its semantics, its proof theory, and the metatheoretic properties that establish it as a formal system rather than a notational convention.

The logic is defined. It works. What is built on it is a separate question.

11. Future Work

The metatheoretic results in this paper are established using classical frameworks that assume infinite objects. Completeness (Theorem 3) is proved via the FOL embedding and Godel's completeness theorem, which constructs models from consistent theories using inherently infinite methods. The compactness failure (Proposition 7) uses an infinite set of sentences. Trakhtenbrot's theorem (Theorem 9) quantifies over all finite structures, an infinite collection. Cut-elimination, interpolation, and definability (Theorems 4, 5, 6) use induction principles formulated in classical arithmetic.

These results are correct as mathematical theorems. They establish BFOL's position within the classical metatheoretic landscape. However, a logic whose grammar enforces bounded quantification is proving its own properties using unbounded reasoning. This tension does not invalidate the results, but it leaves open the question of whether they can be established within bounded frameworks.

Three directions for future work:

Bounded completeness. A completeness result for BFOL under GBP that does not pass through the FOL embedding or Godel's completeness theorem. Such a result would establish that BFOL's proof system is complete for finite models using only finitary reasoning.

Bounded Trakhtenbrot. Theorem 9 establishes undecidability for truth across all finite structures. Theories built on BFOL (such as Bounded Set Theory) define their models as specific families of finite structures, not all finite structures. A reformulation of Trakhtenbrot for specific families of finite models would close this gap and provide direct undecidability results for such theories.

Bounded metatheory. The induction principles used in Theorems 4, 5, and 6 are formulated in classical arithmetic. Bounded analogues using only bounded induction would align the metatheory with the logic's own design principles.

References

- Beth, E.W. (1953). "On Padoa's Method in the Theory of Definition." *Indagationes Mathematicae*, 15, 330–339.
- Buss, S. (1986). *Bounded Arithmetic*. Bibliopolis.
- Chang, C.C. and Keisler, H.J. (1990). *Model Theory*. 3rd edition. North-Holland.
- Church, A. (1936). "A Note on the Entscheidungsproblem." *Journal of Symbolic Logic*, 1(1), 40–41. Correction: 1(3), 101–102.
- Craig, W. (1957). "Three Uses of the Herbrand-Gentzen Theorem in Relating Model Theory and Proof Theory." *Journal of Symbolic Logic*, 22(3), 269–285.
- Ebbinghaus, H.-D. and Flum, J. (1999). *Finite Model Theory*. 2nd edition. Springer.
- Fagin, R. (1974). "Generalized First-Order Spectra and Polynomial-Time Recognizable Sets." In R. Karp (ed.), *Complexity of Computation*, SIAM-AMS Proceedings, vol. 7, 43–73.
- Gentzen, G. (1935). "Untersuchungen über das logische Schließen." *Mathematische Zeitschrift*, 39, 176–210, 405–431. English translation: "Investigations into Logical Deduction," in M.E. Szabo (ed.), *The Collected Papers of Gerhard Gentzen*, North-Holland, 1969.
- Gödel, K. (1930). "Die Vollständigkeit der Axiome des logischen Funktionenkalküls." *Monatshefte für Mathematik und Physik*, 37, 349–360. English translation: "The Completeness of the Axioms of the Functional Calculus of Logic," in J. van Heijenoort (ed.), *From Frege to Gödel*, Harvard University Press, 1967.
- Hájek, P. and Pudlák, P. (1993). *Metamathematics of First-Order Arithmetic*. Springer.
- Libkin, L. (2004). *Elements of Finite Model Theory*. Springer.
- Maehara, S. (1960). "On the Interpolation Theorem of Craig." *Sūgaku*, 12, 235–237. (In Japanese.)
- Mostowski, A. (1957). "On a Generalization of Quantifiers." *Fundamenta Mathematicae*, 44, 12–36.
- Parikh, R. (1971). "Existence and Feasibility in Arithmetic." *Journal of Symbolic Logic*, 36(3), 494–508.
- Quine, W.V.O. (1970). *Philosophy of Logic*. Prentice-Hall.
- Takeuti, G. (1987). *Proof Theory*. 2nd edition. North-Holland.
- Trakhtenbrot, B. (1950). "Невозможность алгоритма для проблемы разрешимости в конечных классах." *Доклады Академии Наук СССР*, 70, 569–572. English summary: "The Impossibility of an Algorithm for the Decision Problem in Finite Classes."
- Troelstra, A.S. and Schwichtenberg, H. (2000). *Basic Proof Theory*. 2nd edition. Cambridge University Press.