Key Applications of Model Theoretic Semantics for FOL: Soundness, Completeness, Compactness, *etc.* [1]

(1) **A Summary of What We've Done So Far for PL**

a. We've given a purely *syntactic* characterization of 'valid inference' in FOL
$S \vdash \psi$

b. We've given a formal (model-theoretic) semantics for FOL notation, and used it to provide a (proper) semantic definition of 'valid inference' in FOL.
$S \vDash \psi$

In these notes, I'll review some key applications of our model-theoretic semantics for FOL. These results will provide further motivation for using 'models' as a mathematical characterization of 'interpretation' for FOL…

---

1. Proving Soundness and Completeness for First Order Logic

(2) **The Theorems We Wish to Prove**

Soundness of FOL:
If ψ can be derived from S in our natural deduction system for FOL, then S entails ψ

- If $S \vdash \psi$, then $S \vDash \psi$
- If $S \vdash \psi$, then if $\mathcal{M}$ is a model for S, $[[\psi]]^M = T$

Completeness of FOL:
If S entails ψ, then ψ can be derived from S in our natural deduction system.

- If $S \vDash \psi$, then $S \vdash \psi$
- If every model $\mathcal{M}$ for S is also a model for ψ, then $S \vdash \psi$

(3) **On the Proof of Soundness for FOL**
The proof of soundness for FOL is not importantly different from the proof of soundness for PL.

- We simply extend the soundness proof of PL so that the induction step also considers the cases where:

(i) ψ is derived by E∃ (ii) ψ is derived by I∃
(iii) ψ is derived by E∀ (iv) ψ is derived by I∀

[1] These notes are based upon material in the following required readings: Gamut (1991), Chapter 4 pp. 148-155; Crossley *et al.* (1972), Chapter 2; Partee *et al.* (1993) Chapter 8 pp. 198-201.

(4) **Illustration of the Additional Steps for Proving Soundness of FOL**
The following subcase can be added to the induction step in (9b) of the handout "Proving the Soundness and Completeness of Propositional Logic"

13. *Deriving ψ by I∃*
Suppose that S ⊢ ψ with a proof consisting of n lines, where the final line has 'I∃' as the justification.

- By definition of 'I∃', ψ = ∃xφ and S ⊢ [α/x]φ with a proof of length m < n.
- By the induction assumption, then S⊨[α/x]φ.
- Now, let $\mathcal{M}$ be any model <D,I> for S, and g be any variable assignment based on $\mathcal{M}$.
- It follows that $\mathcal{M}$ is a model for [α/x]φ.
- Therefore $[[\ [\alpha/x]\varphi\ ]]^{M,g} = 1$
- **Let I(α) = a. It follows that $[[\ \varphi\ ]]^{M,g(x/a)} = 1$.**
- Therefore, there is an a ∈ D such that $[[\ \varphi\ ]]^{M,g(x/a)} = 1$
- Therefore, $[[\ \exists x\varphi\ ]]^{M,g} = 1$.
- Since g was arbitrary, it follows that $[[\ \exists x\varphi\ ]]^{M} = 1$, and so $\mathcal{M}$ is also a model for ∃xφ.
- Since $\mathcal{M}$ was arbitrary, it follow that any model of S is also a model of ∃xφ, and so S ⊨ ∃xφ,

For reasons of time, I won't walk through the other four additional cases…
Nevertheless, we can still see that this proof wouldn't even get off the ground without our formal, model-theoretic semantics for FOL…

(4) **On the Proof of Completeness for FOL**
The proof of completeness for FOL has the same general structure as the proof of completeness for PL.

a. <u>Key Lemma:</u>
Let S be a set of sentences in FOL. S ∪ {ψ} is inconsistent *iff* S ⊢ ~ψ

b. <u>Consistency Theorem:</u>
If S is a consistent set of sentences in FOL, then there is a model $\mathcal{M}$ of S.

c. <u>Consistency Theorem Entails Completeness:</u>
- If S ⊨ ψ, then S ∪ {~ψ} has ***no model.***
- Therefore, consistency theorem (4b) entails that S ∪ {~ψ} is inconsistent.
- Therefore, key lemma (4a) entails that S ⊢ ~~ψ, and so S ⊢ ψ

Again, the most arduous step of the completeness proof is showing proving the Consistency Theorem (4b)…

And the most difficult part of proving (4b) is proving the lemma in (5)…
Also, proving this lemma is <u>much more involved</u> for FOL than for PL.

(5) **Lindenbaum's Lemma**
Let S be a consistent set of sentences in FOL. There is a consistent set S* such that S ⊆ S* and S* has the following key 'closure properties'.[2]

For any formulae φ and ψ of PL:

a. φ ∈ S* *iff* ~φ ∉ S*
b. (φ & ψ) ∈ S* *iff* φ ∈ S* and ψ ∈ S*
c. (φ ∨ ψ) ∈ S* *iff* φ ∈ S* or ψ ∈ S*
d. (φ → ψ) ∈ S* *iff* φ ∉ S* or ψ ∈ S*
e. ∃xφ ∈ S* *iff* there is an individual constant b such that [b/x]φ ∈ S*
f. ∀xφ ∈ S* *iff* for every individual constant b, [b/x]φ ∈ S*

Again, now that we have this huge set S, we're home free!...*

(6) **The Cool Central Insight of Henkin's Proof**
You can take a set S* with the properties in (5), and *directly build a model for S* from the formulas in S* itself!*

(7) **The Model Existence Lemma**
If a consistent set S* of FOL sentences has the closure properties in (5), then there is a model $\mathcal{M}$ for S*.

a. <u>The Model for S*</u>
Let $\mathcal{M}$ be the model <D, I> where:
(i) D is the set of all the <u>individual constants</u> in the language.

(ii) The function I satisfies the conditions below.

1. If α is an individual constant, then I(α) = α

2. If Φ is a predicate letter of arity n, then I(Φ) is such that:
$<\alpha_1, \ldots, \alpha_n> \in I(\Phi)$ *iff* $\Phi\alpha_1 \ldots \alpha_n \in S^*$

b. <u>Key Claim:</u> The model $\mathcal{M}$ defined in (7a) is a model for S*

[2] Those who are intimately familiar with the completeness proof for FOL will know that I'm 'fudging' here on the statement of Lindenbaum's Lemma for FOL (since I'm not mentioning extending our FOL language by adding infinitely many constants).

To prove the key claim in (7), and thus the Model Existence Lemma, we again do (strong) induction on the number of logical operators in an FOL sentence…

- The proof is basically the same as that for PL; the only new and interesting steps come with the atomic formulae and the quantificational formulae

(8) **Proof of the Model Existence Lemma**

Claim:
Let $\mathcal{M}$ be the model defined in (7a). For any natural number n, if φ is a sentence of FOL with n logical operators, then φ ∈ S* iff $[[\varphi]]^M = 1$.

Proof (by Induction on Complexity of Formulae):

a. *Base Step: Atomic Sentences* (n = 0)
Suppose that φ is an atomic formula $\Phi\alpha_1 \ldots \alpha_n$.
- $\Phi\alpha_1 \ldots \alpha_n \in S^*$ *iff* (by condition 2 in definition of $\mathcal{M}$)
- $<\alpha_1, \ldots, \alpha_n> \in I(\Phi)$ *iff* (by condition 1 in definition of $\mathcal{M}$)
- $<I(\alpha_1), \ldots, I(\alpha_n)> \in I(\Phi)$ *iff* (by definition of a 'model')
- $[[\Phi\alpha_1 \ldots \alpha_n]]^M = 1$

b. *Induction Step:*
Let n be such that for all m < n, if φ is a sentence of FOL with m logical operators, then φ ∈ S* iff $[[\varphi]]^M = 1$. We will now show that if φ is a sentence of FOL with n logical operators, then φ ∈ S* iff $[[\varphi]]^M = 1$.

There are six cases to consider: ~, &, ∨, →, ∃, ∀

1. φ = ~ψ *Same as in completeness proof for PL*
2. φ = (ψ & χ) *Same as in completeness proof for PL*
3. φ = (ψ ∨ χ) *Same as in completeness proof for PL*
4. φ = (ψ → χ) *Same as in completeness proof for PL*
5. **φ = ∃xψ**

- ∃xψ ∈ S* *iff* (by closure property (5e))
- There is an individual constant b such that [b/x]ψ ∈ S*
  iff (by the induction assumption)
- $[[\ [b/x]\psi\ ]]^M = 1$ *iff*
- $[[\ \psi\ ]]^{M,g(x/b)} = 1$, for an arbitrary variable assignment g *iff*
- $[[\exists x\psi\ ]]^M = 1$

6. **φ = ∀xψ** ***Proof similar to the one in 5.***

(9) **Putting it All Together**

a. Given the lemma in (5), any consistent set S of FOL sentences can be 'expanded' into a larger consistent set S* with the properties in (5a-f).

b. Given (7)-(8), any such set S* with the properties in (5a-f) has a model $\mathcal{M}$.

c. Since $\mathcal{M}$ is a model for S*, and S $\subseteq$ S*, it follows that $\mathcal{M}$ is a model for S.
Thus, any consistent set of FOL sentences has a model $\mathcal{M}$. QED (4b).

(10) **Really Cool Thing to Notice**

In the Henkin proof of FOL's completeness, we *construct* a model structure directly from the set of FOL sentences S*.

- The domain is the set of individual constants
- Each constant is interpreted as itself
- Each predicate letter Φ is interpreted as the relation holding of α_1, …, α_n iff the sentence '$\Phi\alpha_1 \ldots \alpha_n$' is in S*

Thus, the very nature of models themselves – objects of the form <D, I> - factor into the central core step of the proof.

- This could be viewed as giving 'additional motivation' for this particular formalization of the notion of 'FOL interpretation.'

(11) **Taking Stock of What We've Done**

- In our last set of notes, we developed a mathematically rigorous characterization of what it means for a sentence of FOL to be 'true under an interpretation' (model).

- We've just seen how this notion has allowed us to prove that our syntactic proof system for FOL *<u>is a perfect syntactic characterization of validity in FOL</u>*

- **For the first time in human history, we've shown that we can indeed give a perfect, purely syntactic characterization of what it means for an inference to be valid (in a specified language)**

2. Some Other Important Results of our Model Theoretic Semantics

(12) **Demonstrating Consistency**
Suppose you want to know whether some set S of FOL formulae are *consistent* or not; that is, you want to know whether $S \vdash \bot$ (inconsistent) or $S \nvdash \bot$ (consistent).

a. The Challenge:
If all we have is the natural deduction system, if $S \vdash \bot$, then we can eventually show that (we'll have the proof). But if $S \nvdash \bot$, there's no way to conclusively show this (in finite time) with just the natural deduction system.

b. The Solution:
- Given soundness, if $S \vdash \bot$, then $S \vDash \bot$, and so there is no model $\mathcal{M}$ for S.
- **Thus, with our model theoretic semantics, we can show that S is consistent ($S \nvdash \bot$) by devising a model for S!**

(13) **Demonstrating Independence**
Suppose you want to know whether some FOL sentence φ can be derived from S or not. That is, you want to know whether $S \vdash \varphi$ or $S \nvdash \varphi$.
- In the latter case, we say that φ is 'independent' of S.

a. The Challenge:
Again, if all we have is the natural deduction system, there's no way to conclusively show (in finite time) that $S \nvdash \varphi$.

b. The Solution:
Given soundness again, if $S \nvDash \varphi$ then $S \nvdash \varphi$. Therefore, if we can show that that there is a model $\mathcal{M}$ of S such that $[[\varphi]]^M = 0$, we've shown that $S \nvdash \varphi$.

(14) **Historical Relevance of (12) and (13)**

For centuries, mathematicians struggled to show whether the fifth axiom of Euclid's geometry was 'independent' of the other axioms (or not).

- People tried to derive contradictions from the negation of the fifth axiom, but couldn't succeed.

- Finally, however, mathematicians succeeded in showing that there were (informal) models of geometric systems where (only) the fifth axiom is negated.

  - Such (informal) models show that non-Euclidean geometries are consistent (12) and that the fifth axiom is indeed independent (13).

(14) **Compactness Theorem**

a. Claim:
A set of FOL sentences S has a model *iff* every finite subset of S has a model.

b. Proof:
(i) If S has a model, then of course every finite subset does (duh).
(ii) Suppose S *doesn't* have a model.

- Thus, $S \vDash \bot$ and so by completeness, $S \vdash \bot$.
- Since derivations are finite, it follows that there is some finite subset $S' \subseteq S$ such that $S' \vdash \bot$.
- Thus, by soundness, $S' \vDash \bot$, and so S' doesn't have a model.
- Thus, *not* every finite subset $S' \subseteq S$ has a model.

The theorem in (14) is a powerful tool in advanced meta-logic and model-theory, since it allows for easier proofs that certain infinite sets of sentences are consistent…

(15) **The Löwenheim-Skolem Theorems**

a. The 'Downward' Löwenheim-Skolem Theorem
If there is a model $\mathcal{M}$ (= <D,I>) for S, then there is a model $\mathcal{M}'$(= <D',I'>) for S such that D' is countable (finite or countably infinite).

b. The 'Upward' Löwenheim-Skolem Theorem
If there is a model $\mathcal{M}$ (= <D,I>) for S whose domain D is countably infinite, then there is a model $\mathcal{M}'$(= <D',I'>) for S such that D' is *uncountable*.

The theorems in (15) entail that – although sentences of FOL can 'say' that there are infinitely many things – they cannot say whether that infinity is countable or not…

- Some philosophers (Putnam) have also tried to connect (15) with philosophical problems relating to the nature of reference and meaning…

(16) **Lindstrom's Theorem (Informally Put)**
Any logic that satisfies conditions C (unnamed here), and satisfies compactness (14), and also satisfies the Downard Löwenheim-Skolem Theorem *just is* FOL.

Lindstrom's Theorem provides a powerful tool for showing that a given logical system – no matter how superficially different from FOL it is – is ultimately just a notational variant of FOL.

3. Models Beyond First Order Logic

As shown above, the introduction of 'models' as abstract characterizations of 'interpretation' for FOL has been an extremely important and fertile development in logic and mathematics…

- Throughout the 50s and 60s, logicians developed model-theoretic semantics for logical systems beyond FOL.

(17) **Modal Logic**
One of the first huge advances was Saul Kripke's development of a model theoretic semantics for modal logic, along with the first completeness proofs of such logics.

a. Syntax: PL/FOL + the unary sentence connectives ◊ 'it is possible that' and □ 'it is necessary that'

◊(p & □q) 'it's possible that *p* and it's necessary that *q*'.

b. Semantics: We add to the model structure a set of possible worlds (and a relation over those worlds, a.k.a. 'the modal base')

(18) **Temporal Logic**

a. Syntax: PL/FOL + the unary sentence connectives P 'it was the case that' and F 'it will be the case that'

F(p & Pq) 'it will be the case that *p* and it was the case that *q*.

b. Semantics: We add to the model structure a set of times (and a linear ordering over those times)

(19) **Second Order Logic**

a. Syntax: FOL + variables over n-ary relations (for every natural number n)

∀*P* (*P*ab → *P*bc)
'every relation that holds between a & b also holds between c & b'

b. Semantics: We extend variable assignments so that they can accommodate these new variables.

When we combine these notations and model-theoretic structures, we seem to be approaching the expressive capacity of a significant sub-part of human language!

'It was once possible that Dave would have something in common with Mary'
P◊F∃*P* (*P*d & *P*m)

(20) **The Exciting Possibility (*ca.* 1950s and 1960s)**
Could we develop a model theoretic semantics for a natural language (or at least a significant part of one)?

(21) **Why Would We Want a Model-Theoretic Semantics for English?**

a. <u>A Compositional Semantics for Natural Language</u>
A model theoretic semantics *is* a compositional semantics.

- The definition of 'interpretation w.r.t. a model $\mathcal{M}$' specifies how the semantic value of a complex expression is computed in terms of (i) the semantic value of its component expressions, and (ii) how the complex expression is syntactically constructed.

- Why would we want a compositional semantics for a natural language? See Linguistics 610….

b. <u>A Theory of How Language Connects with 'Reality'</u>
Interpreting a language w.r.t. a model gives us a theory of how the language 'hooks up' with language-external reality.

- A model is a little mathematical representation of the universe (or 'multiverse', if we have possible worlds).

  - The domain D are the entities that exist.
  - The range of I gives us properties and relations between those entities

- The function I in the model maps linguistic expressions to 'things' in the model of reality:
  - Individual constants are mapped to entities
  - Predicates are mapped to properties and relations holding of them

- Why would we want a semantic theory that maps language to language/mind-external reality? (See Intro to Philosophy of Language)

c. <u>Natural Languages are Not Fundamentally Different from Artificial Languages</u>

- For most of the 20^{th} century, it was commonly held that there is a fundamental difference between natural languages and artificial languages (like FOL).

- This difference in type was held to entail that the mathematical tools for analyzing artificial languages could not be applied to natural languages.

- **But, if we could give a model theoretic semantics for NL, that would weaken the idea that there is such a fundamental division (Montague)**

(22) **Another, Partisan Comment**

- When we semanticists lay out a model theoretic semantics for a natural language, **we thereby specify what (some of) the valid inferences in that language are.**

    - Under any model theoretic semantics, some sentences will end up being predicted to be true whenever some other (set of) sentences are…

- Moreover, in as much as the theory is truly rigorous and formal, **these predictions can be calculated without any understanding of the (object) language itself.**

Thus, we semanticists are the true inheritors of the original Artisotelian logical program: providing a formal/predictive characterization of valid inference for (natural) language.

(23) **Immediate Challenges For Developing a Model-Theoretic Semantics for English**

Alas, in the world of the late 50s and early 60s, the following were some obvious roadblocks to developing such a program for natural language semantics…

a. <u>Semantic Ambiguity of (Seemingly) Syntactically Unambiguous Sentences</u>

- Sentences like the following are semantically ambiguous, but they don't seem (at first glance) to be *syntactically* ambiguous.

    (i) Some girl loves every man.
    (ii) John wants to buy an ugly car.

- A model-theoretic semantics, though, will map every sentence structure to exactly one interpretation in the model…

b. <u>The Ungodly Complexity of Natural Language Syntax</u>

- A model-theoretic semantics builds upon a recursive characterization of the sentences of a language (e.g. FOL).

- Thus, to give a model-theoretic semantics for English, *we need some kind of recursive characterization of a 'well-formed' English sentence…*

- *But who in the world understands English grammar well enough to give a recursive formal syntax for English?*

c. The Ungodly Complexity of the Syntax/Semantics Interface in NL

The main reason why people were so pessimistic about a formal semantics for natural language is that there just seemed to be so many *puzzles* about the syntax/semantics interface in natural language.

(i) John wants to find a unicorn and eat it.

- This sentence doesn't commit us to there being a unicorn…
- So what the heck does 'it' refer to?
- Could 'it' function as a bound pronoun? *How*?

(ii) Every boy sings or dances.

- This sentence doesn't mean the same thing as 'every boy sings or every boy dances'.
- Thus, this isn't a case of two conjoined sentences with ellipsis.
- So, 'or' must join together the two VPs directly…
- *BUT HOW IS THAT POSSIBLE, IF IT ALSO JOINS TOGETHER TWO SENTENCES?*

(iii) The temperature is 90 degrees and rising.

- The predicate "is 90 degrees" seems to want the *extension* of "the temperature". (~ the extension of "the temperature" at this time is 90).

- The predicate "is rising" seems to want the *intension* of "the temperature".

- But, in this sentence, the NP "the temperature" is argument to *both* those predicate simultaneously… *HOW?!*

(iv) A man who is smoking walked in.

- The truth-conditions of this sentence seem to be:
  $\exists x((Mx \;\&\; Sx) \;\&\; Wx)$
  'there is an x s.t. x is a man and x is smoking and x walked in.'

- Clearly, 'Mx' is contributed by 'man' and 'Wx' by walked in…
  - But how does 'who is smoking' contribute '(… & Sx)'?
  - *HOW DOES THAT WORK??*

Of course, Montague (1974) showed how the issues in (23a-c) can be solved…
To see how he did this, though, we need to start with ALGEBRA…