Part III - Model Theory
The Witness Property and Henkin Constructions

These notes discuss the fact stated during the proof of the Omitting Types Theorem. Recall that we started with a countable language $\mathcal{L}$, and defined $\mathcal{L}^* = \mathcal{L} \cup C$, where C is a countably infinite set of new constant symbols. However, the restriction on the cardinality of $\mathcal{L}$ and C was only important for later parts of the proof of omitting types (in particular, enumerating all $\mathcal{L}^*$-sentences). In the present context, we don't need to assume $\mathcal{L}$ or C is countable. Thus, to simplify things, we can just start with an arbitrary language $\mathcal{L}$ and *assume* $\mathcal{L}$ already has enough constants for witnesses. So we will actually consider the following slightly stronger version of the fact stated during lecture.

Fact 1. *Let T be a complete satisfiable $\mathcal{L}$-theory with the **witness property**, i.e., for any $\mathcal{L}$-formula $\varphi(x)$ there is a constant symbol c in $\mathcal{L}$ such that $T \models \exists x\varphi(x) \to \varphi(c)$. Then the "Henkin model" is a model of T.*

Proof. Let's first recall what is meant by the Henkin model. Let C be the set of constant symbols in $\mathcal{L}$. Define an equivalence relation $\sim$ on C such that $c \sim d$ if and only if $T \models c = d$. Let M be the set $C/\sim$ of equivalence classes (which we denote by $[c]$ for $c \in C$). We define an $\mathcal{L}$-structure $\mathcal{M}$ with universe M as follows:

(*i*) Given a constant symbol $c \in C$, let $c^{\mathcal{M}} = [c]$.

(*ii*) Given an n-ary function symbol f and $c_1, \ldots, c_n, d \in C$, we let $f^{\mathcal{M}}([c_1], \ldots, [c_n]) = [d]$ if and only if $T \models f(c_1, \ldots, c_n) = d$.

(*iii*) Given an n-ary relation symbol R, let $R^{\mathcal{M}} = \{([c_1], \ldots, [c_n]) \in M^n : T \models R(c_1, \ldots, c_n)\}$.

We need to check that $\mathcal{M}$ is well-defined. There is no issue in (*i*), but we need to analyze (*ii*) and (*iii*).

Suppose f is an n-ary function symbol. We need to show that $f^{\mathcal{M}}$ is a well-defined function. So fix elements $a_1, \ldots, a_n \in M$. We may choose constant symbols $c_1, \ldots, c_n \in C$ such that $a_i = [c_i]$. Now let $\varphi(x)$ be the $\mathcal{L}$-formula $x = f(c_1, \ldots, c_n)$. By the witness property, there is a constant symbol $d \in C$ such that $T \models \exists x\phi(x) \to \phi(d)$. Since $T \models \exists x\phi(x)$ (in any model of T the interpretation of $f(c_1, \ldots, c_n)$ exists), we have $T \models \phi(d)$. Therefore we at least have some d such that $T \models f(c_1, \ldots, c_n) = d$, and so our definition of $f^{\mathcal{M}}$ does assign a value $[d]$ to $f^{\mathcal{M}}(a_1, \ldots, a_n)$. But we need to make sure that this value doesn't depend on the choice of representatives $c_1, \ldots, c_n$. So suppose $c'_1, \ldots, c'_i, d'$ are constant symbols such that $c'_i \sim c_i$ and $T \models f(c'_1, \ldots, c'_n) = d'$. We need to show that $d \sim d'$. But this is clear from the definition of $\sim$. Indeed, $T \models c'_i = c_i$ for all i, and so $T \models f(c_1, \ldots, c_n) = f(c'_1, \ldots, c'_n)$, and thus $T \models d = d'$.

Finally, suppose R is an n-ary relation symbol. We need to show that the definition of $R^{\mathcal{M}}$ does not depend on the choice of $\sim$-representative. But this is again immediate from the definition of $\sim$. In particular, if $c_1, \ldots, c_n, c'_1, \ldots, c'_n$ are constant symbols, and $T \models c_i = c'_i$ for all i, then $T \models R(c_1, \ldots, c_n) \leftrightarrow R(c'_1, \ldots, c'_n)$.

Now that we know $\mathcal{M}$ is well-defined, we can prove that $\mathcal{M}$ is a model of T. In particular, we show by induction on formulas that, for any $\mathcal{L}$-formula $\varphi(x_1, \ldots, x_n)$ and any constant symbols $c_1, \ldots, c_n \in C$,

$$\mathcal{M} \models \varphi([c_1], \ldots, [c_n]) \;\Leftrightarrow\; T \models \varphi(c_1, \ldots, c_n). \tag{$*$}$$

As usual, this requires a lemma on terms. In particular, for any $\mathcal{L}$-term $t(x_1, \ldots, x_n)$ and constant symbols $c_1, \ldots, c_n, d \in C$, we claim that $t^{\mathcal{M}}([c_1], \ldots, [c_n]) = [d]$ if and only if $T \models t(c_1, \ldots, c_n) = d$. This can be proved by induction on terms. If t is a constant symbol c, then the claim is $[c] = [d]$ if and only if $T \models c = d$, which is by definition of $\sim$. If t is a variable then, in light of how M is defined, the claim reduces to what we just observed for constant symbols. Finally, for the inductive step, assume the result for terms $t_m, \ldots, t_m$, say each of arity n, and fix an m-ary function symbol f. Fix $c_1, \ldots, c_n, d$ and let d_i be a constant symbol such that $t_i^{\mathcal{M}}([c_1], \ldots, [c_n]) = [d_i]$. Let t be the new term $f(t_1, \ldots, t_m)$. Then

$$t^{\mathcal{M}}([c_1], \ldots, [c_n]) = [d] \;\Leftrightarrow\; f^{\mathcal{M}}([d_1], \ldots [d_m]) = [d] \;\Leftrightarrow\; T \models f(d_1, \ldots, d_n) = d \;\Leftrightarrow\; T \models t(c_1, \ldots, c_n) = d,$$

where the first equivalence uses the definition of t, the second uses the definition of $f^{\mathcal{M}}$, and the third uses the induction hypothesis (which tells us $T^* \models t_i(c_1, \ldots, d_n) = d_i$ for all i).

Now we can return to proving the main claim $(*)$ by induction on φ. Suppose φ is atomic, i.e., of the form $R(t_1, \ldots, t_m)$ for some m-ary relation symbol R and terms $t_1, \ldots, t_m$, each of arity n. Fix constant symbols $c_1, \ldots, c_n$, and let $d_1, \ldots, d_m$ be constant symbols such that $t_i^{\mathcal{M}}([c_1], \ldots, [c_n]) = [d_i]$. So $T \models t_i(c_1, \ldots, c_n) = d_i$ by the lemma for terms. Therefore

$$\mathcal{M} \models \varphi([c_1], \ldots, [c_n]) \;\Leftrightarrow\; ([d_1], \ldots, [d_m]) \in R^{\mathcal{M}} \;\Leftrightarrow\; T \models R(d_1, \ldots, d_m) \;\Leftrightarrow\; T \models \varphi(c_1, \ldots, c_n).$$

The previous argument works the same way if R is equality.

Next we do the induction steps. The step for conjunctions is trivial, and boils down to the fact that, for any sentences ϕ and ψ, we have $T \models \varphi \wedge \psi$ if and only if $T \models \varphi$ and $T \models \psi$. Similarly, the negation step boils down to the fact that $T \models \neg\phi$ if and only if $T \not\models \phi$. But we need to justify why this is the case. On one hand, if $T \not\models \phi$, then $T \models \neg\phi$ since T is complete. On the other hand, if $T \models \neg\phi$ then $T \not\models \phi$ since T is satisfiable.

So all we have left is the quantifier step, and we'll work with existential quantifiers. Assume $(*)$ holds for $\varphi(x_1, \ldots, x_n, y)$ and consider $\exists y \varphi(\bar{x}, y)$. Fix constant symbols $c_1, \ldots, c_n$.

First assume $\mathcal{M} \models \exists y \varphi([c_1], \ldots, [c_n], y)$. So there is a constant symbol d such that $\mathcal{M} \models \varphi([c_1], \ldots, [c_n], [d])$. By induction, $T \models \varphi(c_1, \ldots, c_n, d)$, which implies $T \models \exists y \varphi(c_1, \ldots, c_n, d)$.

Finally, assume $T \models \exists y \varphi(c_1, \ldots, c_n, d)$. By the witness property, there is a constant symbol d such that $T \models \varphi(c_1, \ldots, c_n, d)$. By induction $\mathcal{M} \models \varphi([c_1], \ldots, [c_n], [d])$, and so $\mathcal{M} \models \exists y \varphi([c_1], \ldots, [c_n])$. □

Direct proof of Compactness

One can prove the Compactness Theorem directly (i.e., without going through the proof system), using the same ideas as above. Call an $\mathcal{L}$-theory T **maximal** if for any $\mathcal{L}$-sentence φ, either $\varphi \in T$ or $\neg\varphi \in T$. So, in particular, maximality is a strong form of completeness. We also say that an $\mathcal{L}$-theory T has the **strong witness property** if for any $\mathcal{L}$-formula $\phi(x)$ there is a constant symbol c in $\mathcal{L}$ such that the sentence $\exists x\phi(x) \to \phi(c)$ is in T.

Lemma 2. *Suppose T is a finitely satisfiable $\mathcal{L}$-theory. Then there is an expansion $\mathcal{L}^*$ of $\mathcal{L}$ by constants, and an $\mathcal{L}^*$-theory T^* such that T^* is maximal, finitely satisfiable, has the strong witness property, and contains T.*

Proof. We first build $\mathcal{L}^*$ and a finitely satisfiable $\mathcal{L}^*$-theory $T' \supseteq T$ with the strong witness property. This is very much like the proof of #9 on Examples Sheet 1. In particular, start with $\mathcal{L}_0 = \mathcal{L}$ and $T_0 = T$. Given $\mathcal{L}_n$ and a finitely satisfiable $\mathcal{L}_n$-theory $T_n \supseteq T$, let $\mathcal{L}_{n+1}$ be obtained from $\mathcal{L}_n$ by adding a new constant symbol c_φ for every $\mathcal{L}_n$-formula $\varphi(x)$ in one free variable. Then let T_{n+1} be obtained from T by adding the sentence $\exists x\varphi(x) \to \varphi(c_\varphi)$ for every $\mathcal{L}_n$-formula $\varphi(x)$. Let H_n denote this new set of sentences added to T_n (so $T_{n+1} = T_n \cup H_n$). We need to check that T_{n+1} is still finitely satisfiable. Note that any finite subset of T_{n+1} is contained in $\Delta \cup H_n$ for some finite subset $\Delta \subseteq T_n$. So it suffices to fix a finite set $\Delta \subseteq T_n$, and show that $\Delta \cup H_n$ is satisfiable. Since T_n is finitely satisfiable, there is an $\mathcal{L}_n$-structure $\mathcal{M} \models \Delta$. Expand $\mathcal{M}$ to an $\mathcal{L}_{n+1}$ structure by interpreting each c_φ as a solution to $\varphi(x)$, if one exists in $\mathcal{M}$. Then $\mathcal{M} \models \Delta \cup H_n$.

Now let $\mathcal{L}^* = \bigcup_{n\geq 0} \mathcal{L}_n$ and $T' = \bigcup_{n\geq 0} T_n$. Then T' is finitely satisfiable (since any finite subset is contained in T_n for some n), has the strong witness property (since any $\mathcal{L}^*$-formula is a $\mathcal{L}_n$-formula for some n), and contains T. All we are missing is maximality and, for this, we need Zorn's Lemma.

Consider the set X of *all* $\mathcal{L}^*$-theories that are finitely satisfiable, have the strong witness property, and contain T. So we have just shown that X is nonempty. Consider X as a partial order under the subset relation. Let $\mathcal{C} \subseteq X$ be a chain with respect to this partial order, and let $T_\mathcal{C}$ be the union of all elements (theories) in the chain. We claim that $T_\mathcal{C} \in X$. It is clear that $T_\mathcal{C}$ contains T and has the strong witness property. Moreover, $T_\mathcal{C}$ is finitely satisfiable since $\mathcal{C}$ is a *chain*, and so any finite subset of $T_\mathcal{C}$ is contained in some theory in $\mathcal{C}$.

So we can apply Zorn's Lemma to X and obtain a theory T^* that is order-maximal in X. We claim that this corresponds to maximality the way we defined it for theories above. So fix an $\mathcal{L}^*$-sentence φ and suppose, for a contradiction, that $\varphi, \neg\varphi \notin T^*$. Let $T_1 = T^* \cup \{\varphi\}$ and $T_2 = T^* \cup \{\neg\varphi\}$. Both of these theories properly extend T^*, and so neither can be in X since T^* is order-maximal in X. Since both theories have the strong witness property and contain T, it must then be the case that neither is finitely satisfiable. So there are finite sets $\Delta_1, \Delta_2 \subseteq T^*$ such that both $\Delta_1 \cup \{\varphi\}$ and $\Delta_2 \cup \{\neg\varphi\}$ are unsatisfiable. It follows that $\Delta := \Delta_1 \cup \Delta_2$ is unsatisfiable, since any model of Δ would have to satisfy either φ or $\neg\varphi$. Since Δ is a finite subset of T^*, this contradicts finitely satisfiability of T^*. □

Theorem 3 (Compactness)**.** *Any finitely satisfiable theory is satisfiable.*

Proof. Let T be a finitely satisfiable theory. By Lemma 2, we may assume without loss of generality that T is maximal and has the strong witness property. The idea now is to use a variation of Fact 1 to build a model of T. Note that the only assumption from Fact 1 we are missing is satisfiability, which is of course what we are trying to prove. So we re-formulate Fact 1 as follows.

Fact 1*. Let T be a maximal finitely satisfiable $\mathcal{L}$-theory with the strong witness property, and let C be the set of constant symbols in $\mathcal{L}$. Define $\sim$ on C such that $c \sim d$ if and only if $c = d$ is in T. Define an $\mathcal{L}$-structure $\mathcal{M}$ on $M = C/\!\sim$ exactly as in the proof of Fact 1, but replace all occurrences of "$T \models \varphi$" with "$\varphi \in T$". Then $\mathcal{M}$ is a model of T. In particular, for any $\mathcal{L}$-formula $\phi(x_1, \ldots, x_n)$ and constant symbols $c_1, \ldots, c_n$ in $\mathcal{L}$, we have

$$\mathcal{M} \models \phi([c_1], \ldots, [c_n]) \;\Leftrightarrow\; \phi(c_1, \ldots, c_n) \in T. \qquad (**)$$

The proof of this variation is nearly identical. The main subtlety is that all of the uses of satisfiability in the proof of Fact 1 only relied on a finite amount of information. By working with a maximal theory, and replacing the notion "$T \models \varphi$" with "$\varphi \in T$", we are able to recover everything with finite satisfiability of T. Let me just illustrate two examples of this, which hopefully will convince you that the same proof idea works (and you can check the details if you are really motivated).

For one example, let's see why the interpretation of function symbols is still well-defined. Let f be an n-ary function symbol, and fix $c_1, \ldots, c_n, d, c'_1, \ldots, c'_n, d' \in C$ such that $c_i \sim c'_i$ for all $1 \leq i \leq n$, and T contains the sentences $f(c_1, \ldots, c_n) = d$ and $f(c'_1, \ldots, c'_n) = d'$. Then we need to show $d \sim d'$ in order for our choice of $f^{\mathcal{M}}([c_1], \ldots, [c_n])$ to be uniquely defined. So toward a contradiction, suppose $d' \not\sim d$, i.e., $d = d'$ is not in T. By maximality, $d \neq d'$ is in T. But now we have the following finite subset of T:

$$\{c_1 = c'_1, \ldots, c_n = c'_n, f(c_1, \ldots, c_n) = d, f(c'_1, \ldots, c'_n) = d', d \neq d'\}.$$

This set is not satisfiable, contradicting finite satisfiability of T.

For another example, let's verify the negation step in the induction argument for $(**)$. Similar to the proof of Fact 1, and in light of how $(**)$ has been modified compared to $(*)$, this boils down to showing that for any $\mathcal{L}$-sentence φ, we have $\neg\varphi \in T$ if and only if $\varphi \notin T$. For one direction, if $\varphi \notin T$ then $\neg\varphi \in T$ by maximality. For the other direction, if $\neg\varphi \in T$ then $\varphi \notin T$ since otherwise $\{\varphi, \neg\varphi\}$ would contradict finite satisfiability of T.

The rest of the modifications to the proof go very much the same way, and in the end we build a model $\mathcal{M}$ of T, which is what we needed to do. □

Finally, let's point out how the previous proof gives us the Downward Löwenheim-Skolem Theorem (as stated in Lecture 2).

Corollary 4 (DLST)**.** *If T is a finitely satisfiable $\mathcal{L}$-theory, then T has a model of size at most $|\mathcal{L}| + \aleph_0$.*

Proof. Note that in the previous proof of compactness, we have $|M| \leq |C| \leq |\mathcal{L}|$. However, in the second line of the proof, we used Lemma 2 and replaced $\mathcal{L}$ by a bigger language. So what we really need to check is that in Lemma 2, we have $|\mathcal{L}^*| \leq |\mathcal{L}| + \aleph_0$. This is a straightforward cardinality exercise, very much like the ones we've done before. The key point is that for any language $\mathcal{L}$, the number of $\mathcal{L}$-formulas is $|\mathcal{L}| + \aleph_0$. □