Ultrafilters as Propositional Theories

Jason Turner

November 20, 2021

Open up a (sufficiently advanced) mathematics textbook to the right page, and you can find a definition like the following:

Definition.

Let I be a non-empty set. A **filter over I** is a set $\mathscr{F} \subseteq \mathscr{P}(I)$ that does not contain $\emptyset$ and which meets the following conditions:

(i) If $S, T \in \mathscr{F}$, then $S \cap T \in \mathscr{F}$, and
(ii) If S in $\mathscr{F}$ and $S \subseteq T \subseteq I$, then $T \in \mathscr{F}$.

Furthermore, if $\mathscr{F}$ is also complement-complete, then it is an **ultrafilter**.

If you're at all like me, you spend a good couple hours trying to get a sense of how this definition works and why it allows us to prove interesting things before eventually giving up and seeing what's new on Netflix.

No surprise; mathematics texts are a genre unto themselves, and one most philosophers haven't been trained to read. But it is also unfortunate. Ultrafilters are a powerful technical tool that philosophers can put to good use. Philosophically accessible introductions to other technical tools aren't hard to find, but somehow filters and ultrafilters have been overlooked.[1] As a result, ultrafilters remain a common gap in the technically-minded philosoper's toolkit.

My goal here is to fill that gap, presenting ultrafilters and theorems about them using philosophically familiar ideas. While some philosophers might encounter ultrafilters while studying topology, I'll present them as used in another of their natural habitats, model theory. Łos's theorem — the central model-theoretic result about filters and ultrafilters — says that, if a set of propositions comprises an ultrafilter, some model will make true exactly the sentences expressing propositions in the filter. It is a powerful model-theoretic tool for showing that certain models exist. But if we don't understand the tool, we can't use it.

What philosophically familiar ideas are we going to use here? Those of *possible worlds* and *propositions*. In particular, we can think of ultrafilters as a special kind of set of propositions, or a special kind of 'propositional theory'. A filter is a set of propositions that meet certain conditions, and an ultrafilter is a filter that meets an additional condition. In particular, a filter is a set of propositions that is closed under (finite) conjunction and implication; and an ultrafilter is a filter that is also negation-complete.

[1]They go unmentioned, for instance, in popular textbooks such as Boolos et al. 2002, Gamut 1991a,b, and Sider 2010. Enderton 2001: 142 mentions them briefly only to set them aside. Discussions in other books popular in philosophy graduate programs (e.g. Mendelson 1997: 129–136) aren't much more accessible than those in standard mathematics texts, such as Chang and Keisler 1990: 211–219.

Starting from this basic idea, I will run through two main results — the central theorem on ultrafilters and Łos's Theorem — and then show how we use them to prove the compactness theorem (which says that if every finite subset of Γ has a model, then Γ does, too.) None of the results are novel, but my hope is that, by seeing them worked through in propositional clothing, they will become accessible to a larger audience.

1 Propositional Theories

1.1 Background

Suppose we have a set of worlds W. They may be possible worlds, but they don't have to be. The only things that matters (for our purposes) about them are that (i) they make sentences true or false, and (ii) every world is 'classically logical' — that is, for any world w, there is a model $\mathcal{M}$ that agrees with w. More precisely, for every world w, there is a model $\mathcal{M}_w$ where ϕ is true at w if and only if ϕ is true on $\mathcal{M}_w$. (We will make this even more precise later, but this is enough to be getting on with for now.)

With this in hand, here are a few standard definitions.

Definition.

- A ***W*-proposition** is any subset of W.
  - If P and Q are W-propositions, their **conjunction** is $P \cap Q$.
  - If P is a W-proposition, its **negation** is $W - P$, that is, the set of worlds in W that are not in P.
  - If P and Q are W-propositions, then ***P* implies *Q*** iff $P \subseteq Q$.
  - $\varnothing$ is the **impossible propositions**.
  - W is the ***W*-necessary** proposition.
  - If ϕ is a sentence, the ***W*-intension of ϕ**, written $[\![\phi]\!]_W$, is the set of worlds where ϕ is true.

When we know what the set W is, we drop all the 'W's in the above and just call them propositions, intensions, and so on. We also write the negation of P as just $-P$. That the definitions for propositions are pretty good is shown by the fact that — given our condition (ii) on W — we have:

Proposition 1.

- $[\![\top]\!] = W$
- $[\![\bot]\!] = \varnothing$
- $[\![\sim\phi]\!] = -[\![\phi]\!]$

- $[\![\phi \wedge \psi]\!] = [\![\phi]\!] \cap [\![\psi]\!]$
- $[\![\phi]\!]$ implies $[\![\psi]\!]$ iff $[\![\phi \to \psi]\!] = [\![\top]\!]$.

1.2 *Filters*

1.2.1 As Finitely Consistent Theories

Just as we often call a set of sentences a 'theory', we can call a set of propositions a 'theory' as well. More precisely, if T is any set of W-propositions, we will call it a a ***W*-propositional theory**.

From this perspective, a filter is simply a special kind of propositional theory:

Definition.

If T is a W-propositional theory, then T is a **filter over *W*** iff it

(i) is closed under pairwise conjunction;

(ii) is closed under implication; and

(iii) does not contain the impossible proposition.

By working through our definitions of 'implication' and the like, you can see that this definition is the same as the one at the beginning of this paper, with W replacing I. But philosophers tend to have a much better developed sense about how to think about propositions, conjunction, and implication than about arbitrary sets, powersets, intersection, and the like, making our current definition easier to digest.

So a filter is a kind of propositional theory. Why should we care about it? Well, because filters are something like *consistent* propositional theories. Here's the intuitive idea. Suppose a propositional theory contained contradictory premises. Then those premises have, as an implication, the impossible proposition. But since a filter is closed under conjunction and implication, if a filter had contradictory premises, it would contain the impossible proposition — which it doesn't.

This intuitive idea is close, but not quite, right. Filters are closed under pairwise conjunction: if p and q are in an ultrafilter, so is $p \cap q$. As a result, they will also be closed under finite conjunction, since any finite conjunction can be gotten by a finite number of pairwise conjunctions. (p-and-q-and-r is just p-and-(q-and-r), and so on.) But a filter could still, in principle, contain an *essentially infinite* contradiction. For instance, it could contain, for each n, the proposition that there are at least n F's, but then also include the proposition that there are only finitely many Fs. No finite conjunction of those propositions are inconsistent; but of course the conjunction of all infinitely many of them is.

This is a big part of why filters are interesting. Closure under finite conjunction is much easier to achieve than closure under arbitrary conjunction. If we can get a propositional theory that has the weaker property, we can use various results about filters to learn more about it. One example, that we will come to at the end of this paper, is the first-order *compactness theorem*. This theorem says that if every finite subset of a set of formulas Γ has a model, then so does Γ. In other words, when it comes to first-order logic, any set of formulas that is finitely consistent is also infinitely consistent. The reason we can prove this is that we can use the finite consistency of all of Γ's subsets to come up with a finitely consistent propositional theory, and then use facts about that propositional theory to construct a model that makes *all* of Γ true.

1.2.2 Alternative Definitions

Sometimes texbooks provide alternative definitions of filters. A reasonably common variant drops the third condition from the definition — the one that excludes the impossible proposition — and goes on to distinguish 'proper filters', which exclude $[\![\bot]\!]$, from 'improper filters'.[2] This is mere teminological deviation. If we wanted, we could call propositional theories closed under implication and conjunction 'prefilters'. Then our 'prefilters' would correspond to their 'filters' and our 'filters' would correspond to their 'proper filters'. (Their 'improper filters' would be our 'prefilters that are not filters.')

With a fixed background set of worlds W, a different definition for 'filter' (in our 'proper' sense) is also possible:

Definition.

If T is a W-propositional theory, then T is a **filter* over W** iff

(i*) it is conjunction complete: a W-conjunction is in T iff both of its conjuncts are; and

(ii*) it does not contain all of the W-propositions.

It is not terribly difficult to see that the two definitions are equivalent.

Proof.

First, suppose T is a filter*. We need to show it meets conditions (i)–(iii). The conjunction of every p with $[\![\bot]\!]$ is $[\![\bot]\!]$, so if $[\![\bot]\!]$ was in T, every W-proposition p would be in T, violatiing condition (ii*) and making T not a filter*. Since it is a filter*, T meets condition (iii). Also, T meets condition (i) by the right-to-left direction of (i*). For (ii), note that if p implies q, then $p = p \cap q$. So if $p \in T$ and p implies q, by the left-to-right direction of (i*), $q \in T$. So T is closed under implication, meeting (ii).

For the other direction, suppose T is a filter. Since by (iii) it doesn't contain $[\![\bot]\!]$, it meets condition (ii*). By (i), if $p, q \in$T, then $p \cap q \in T$, so T meets one direction of (i*). For the other direction, note that a conjunction implies each of its conjuncts, so closure under implication tells us that if $p \cap q \in T$, both p and q are, too, meeting the other direction of (i*). ∴

[2]E.g. Chang and Keisler (1990: 221–222).

1.3 *Ultrafilters*

Suppose that T is a filter over W. Then T might (or might not) meet any of the following three conditions.

Definition.

Maximality: T is the only filter over W that includes T.
That is, if $T \subseteq T'$ and T' is a filter over W, then $T = T'$.

Prime: T is disjunction-complete.
That is, for any W-propositions p and q, $p \cup q \in T$ iff either $p \in T$ or $q \in T$.

Ultra: T is negation-complete.
That is, for any W-proposition p, either p or its negation is in T.

Conceptually, these are all different. Maximality tells us that T is a 'biggest' filter: there's no way to add propositions to the filter without making it a non-filter. Primeness tells us something about how the filter deals with disjunctions, and ultra-ness tells us something about how it interacts with negation.

Despite the conceptual differences, though, the three are equivalent properties over filters:

Theorem 2.

If T is a W-filter, then it is maximal iff it is prime iff it is an ultrafilter.

We'll prove it here, leaving the 'W' indexes off for readability.

Proof.

Suppose first that T is an ultrafilter. Take any proposition p. Since T is ultra, either p or $-p$ is in T. Suppose we extend T by adding p. If $p \in T$, then our addition doesn't extend it. If instead $-p \in T$, then if our extension is closed under conjunctions, $p \cap -p = [\![\bot]\!]$ will be in this extension and so it will not be a filter. Thus it is maximal.

Next, suppose that T is maximal. If p is in T, then so is $p \cup q$, since p implies $p \cup q$ and filters are closed under implication. Conversely, suppose $p \cup q$ is in T but neither p nor q are. If the conjunction of each of p and q with $p \cup q$ were impossible, then $p \cup q$ itself would be impossible. So at least one of p or q could be added to T without making it not a filter, contradicting its maximality. So either $p \in T$ or $q \in T$, making it prime.

Finally, suppose T is prime. Since $[\![\top]\!]$ is in T (being implied by every proposition), and for any proposition p, $p \cup -p = [\![\top]\!]$, by T's primeness, either $p \in T$ or $-p \in T$. So T is ultra. ∴

The proof shows we can go in a circle from ultra-ness, through maximality and primeness, and back to ultra-ness, which means all three properties are equivalent.

2 The Central Theorem on Ultrafilters

2.1 *The Theorem and What it Says*

Ultrafilters are important for a number of reasons. One happy fact, though, is

Theorem 3. The Central Theorem on Ultrafilters

Every filter can be extended to an ultrafilter.

This is known as the **Central Theorem on Ultrafilters**.

In first-order logic, we have Lindenbaumm's Lemma, which tells us that a (proof-theoretically) consistent theory (that is, set of sentences) can be expanded to a consistent, negation-complete theory. For that lemma, we start with a consistent theory Γ, and then take a well-ordering of the sentences of the language. We construct a maximal theory by going through those sentences in order, adding each to our theory if the result would be consistent, and leaving it out otherwise. By the 'end' of this (infinite) process we will have considered every sentence, and thanks to a few facts about how negation and consistency work together, for each sentence we will have thrown either it or its negation (when the latter's turn came around) into our theory.

The idea behind the proof of the Central Theorem is essentially the same. We well-order all the W-propositions. Then we take our starting filter and go through each of these W-propositions in order. If the result of adding a W-proposition to it and closing the result up under conjunction and implication is also a filter, we do that. Otherwise, we don't. By the 'end' of this (infinite) process, we will have considered every W-proposition, and thanks to a few facts about how set theory works, for each W-proposition we will have added either it or its negation (when the latter's turn came around).

In order to do this, we have to make sense of 'the result of adding a W-proposition to a filter and closing the result up under conjunction and implication'. We do this with **additions**: if T is a W-filter and p a W-proposition, then $T + p$ is, intuitively, the result of adding p to T and closing the result up under conjunction and implication.

The way we do this is twofold. First, we conjoin p with each proposition in T. Then we take all those propositions and add in any other proposition they imply. The second step makes the result closed under implication, and since T was conjunction-closed, the first step makes our new theory conjunction-closed, too. Furthermore, since $p \cap q$ always implies q, every proposition in T will be in our new theory. And since $p \cap W = p$, p will be in our new theory as well.

When we roll the two steps together we get:

Definition.

Addition: $T + p = \{q \subseteq W : \text{for some } a \in T, a \cap p \text{ implies } q\}$.

(Note that if $a \in T$, then $a \cap p \in T + p$ because $a \cap p$ implies itself.)

This is then the notion that we use in proving the Central Theorem. When we get to each W-proposition p in our well-ordering, we trade in our previous filter T for the new one $T + p$ if that is, in fact, a filter. Happily, it's possible to show that, if T is a filer, then $T + p$ is a filter if and only if $-p$ is not already in T. This means that, by the time we're done, for each W-proposition, we either added it or refrained from adding it because its negation was already in.

One final point is worth mentioning. When it comes to Lindenbaumm's lemma, if the language of our theory is uncountable, the axiom of choice must be assumed; otherwise, there's no guarantee that the sentences of the language can be well-ordered. Similarly, when it comes to the Central Theorem, if W is infinite, the axiom of choice must be assumed so that the W-propositions can be well-ordered. (Since the set of all W-propositions is the powerset of W, even if W is only countably infinite, there will uncountably many W-propositions.) The Central Theorem — and indeed, any result that follows from it — should be understood as requiring choice.

2.2 *Proving the Theorem*

Before we prove the theorem itself, we should start with verifying our claims about additions. We have:

Lemma 4.

Suppose T is a filter over W and p is a W-proposition. Then

(i) $T \subseteq T + p$ and $p \in T + p$.

(ii) $T + p$ is closed under conjunction and implication.

(iii) $T + p$ is a filter iff $-p \notin T$.

Proof.

For (i): Since $p = p \cap W$ and $p \cap W \subseteq W$, this follows from the definition of $W + p$.

For (ii), we have two parts. First, that $T + p$ is closed under implication. Suppose $q \in T + p$ and q implies r. Since $q \in T + p$, there must be some $a \in T$ where $a \cap p$ implies q. In this case, though, $a \cap p$ implies r (since implication is transitive), so r is also in T.

For (ii) and conjunction closure: If q_1 and q_2 are in $T + p$, then there are a_1 and a_2 in T where $a_1 \cap p$ implies q_1 and $a_2 \cap p$ implies q_2. Since T is a filter, $a_1 \cap a_2$ is also in T. But the conjunction of $a_1 \cap p$ and $a_2 \cap p$ is the same as the conjunction of $(a_1 \cap a_2)$ with p, which will imply the conjunction of p with $q_1 \cap q_2$. Thus $q_1 \cap q_2$ will be in $T + p$, and so it is closed under conjunctions.

For (iii): If $-p \in T$, then by part (i), $-p \in T + p$ and $p \in T + p$. Since the conjunction of p and $-p$ is $[\![\bot]\!]$, this will be in $T + p$ by part (ii), making $T + p$ not a filter.

Conversely, suppose that $T + p$ is not a filter. By part (ii), this can only happen if $[\![\bot]\!] \in T + p$. By the definition of $T + p$, there must be a a in T where $a \cap p$ implies $[\![\bot]\!]$. This means that a implies $-p$; since T is closed under implication, it contains $-p$. ∴

Now we can prove theorem 3. To do this, we rely on a version of the axiom of choice that says that every set can be indexed by ordinals, and we rely on transfinite induction on the ordinals. There are only three kinds of ordinals: zero, 'successor' ordinals which are some ordinal α plus one, and 'limit' ordinals, which are the union of all ordinals smaller than them. We will use an ordinal indexing of all the W-propositions to construct an ordinal-indexed series of filters, and then show that the union of all these theories is an ultrafilter.

Proof.

Let T be a filter over W and take an indexing of the W-propositions by ordinals. When p is indexed by α we call it p_α. Thanks to the fact that ordinals are well-ordered, there will be a least ordinal γ where every W-proposition has an index less than γ. Now we define a series of filters as follows.

$$T_0 = T$$

$$T_{\alpha+1} = \begin{cases} T_\alpha + p_\alpha, \text{ if that is a filter, and} \\ T_\alpha \text{ otherwise.} \end{cases}$$

$$T_\alpha = \bigcup_{\beta<\alpha} T_\beta \text{ for limit } \alpha$$

We want to show that each T_α in this series is a filter, and that if $\beta < \alpha$, then $T_\beta \subseteq T_\alpha$. This is done by induction on the ordinals. In the base case $T_\alpha = T_0$, which is a filter by assumption. For the induction step, we assume that if $\beta < \alpha$, then T_β is a filter.

When α is a successor ordinal, it is a filter by construction, and if $T_\alpha \neq T_{\alpha+1}$, then $T_\alpha \subseteq T_{\alpha+1}$ by lemma 4.i.

For a limit ordinal α, it's clear that if $\beta < \alpha$, then $T_\beta \subseteq T_\alpha$. But we need to show that it is a filter. To do that we have to show that it meets all three conditions of being a filter. For the third condition, if $[\![\bot]\!] \in T_\alpha$, then $[\![\bot]\!]$ would have to be a member of some T_β for $\beta < \alpha$, which contradicts the induction hypothesis that all such T_β are filters.

For implication, suppose that $q \in T_\alpha$ and q implies r. To be in T_α, q must be in some T_β for $\beta < \alpha$. But since T_β is a filter, it is closed under implication, so $r \in T_\beta$ and thus also in T_α.

For conjunction, suppose that q and r are in T_α. Then there are β_q and β_r, both less than α, where $q \in T_{\beta_q}$ and $r \in T_{\beta_r}$. If β is whichever is the larger of the two, then $q, r \in T_\beta$. But T_β is a filter, and so closed under conjunction, so $q \cap r \in T_\beta$. Since T_α includes T_β, $q \cap r$ is in T_α also.

So every T_α in our construction is a filter. Now we take T_γ, the union of all the T_α's in our construction. It is also a filter, by the above argument. We just need to show that it is an ultrafilter. To do this, note that if p is any W-proposition it is p_α for some α. If $T_\alpha + p_\alpha$ was a filter, then $p \in T_\alpha$ and so in T_γ. If instead $T_\alpha + p_\alpha$ was not a filter, then by lemma 4.iii, $-p$ was already in T_α, so $-p \in T_\gamma$. ∴

3 Intensional Spaces

Up until now we've been focused purely on the propositional side of things. We've introduced filters and ultrafilters as a special case of propositional theories — sets of propositions — and shown how to understand the central theorem on ultrafilters as a propositional version of Lindenbaumm's Lemma. But now that we've done this, what can we *do* with these tools?

Several things, in fact. But we'll focus on just one here, a *model-theoretic* one. The main theorem connecting ultrafilters to model theory is Łos's theorem which says, roughly, if you have an ultrafilter T over a space of worlds W, and if every world in W corresponds to some model (of a fixed language), then there is a model that makes all the propositions in T true.

We will precisify all of this in a little bit. But before we do, it will be helpful to connect our theory of W-propositions more directly to model theory.

3.1 The Model-Theoretic Background

We should start by fixing ideas about model theory. Suppose we have a first-order language $\mathscr{L}$, with a fixed stock of names and predicates, infinitely many variables, and primitive logical symbols $\sim$, $\wedge$, $\exists$, and $=$. Other logical symbols are understood as defined in the usual way.

A **model** $\mathcal{M}$ is an ordered pair $\langle D, I\rangle$ of a non-empty domain D and an interpretation function I. For any constant α in $\mathscr{L}$, $I(\alpha) \in D$, and for any n-adic predicate Π in $\mathscr{L}$, $I(\Pi) \subseteq D^n$. (That is, $I(\Pi)$ is a set of n-tuples of things drawn from D.)

For each model $\mathcal{M}$ of $\mathscr{L}$, **a variable assignment over $\boldsymbol{\mathcal{M}}$** is a function from variables of $\mathscr{L}$ to elements of D. A **term** is any constant or variable, and the **denotation** of a term α on a model $\mathcal{M}$ relative to a variable assignment a, written $\alpha^{\mathcal{M},a}$, is $I(\alpha)$ if a constant and $a(\alpha)$ if a variable. If a is a variable assignment over $\mathcal{M}$ and $o \in D$, $a[x \triangleright o]$ is the assignment just like a except that it maps x to o.

To avoid clutter, we write a sequence' $\alpha_1, \ldots, \alpha_n$' as '$\vec{\alpha}$'. Now we can recursively define truth of an open formula on a model relative to a variable assignment, written $\mathcal{M}, a \vDash \phi$. The definition runs:

Definition.

Let $\mathcal{M}$ be a model and a a variable assignment. Then

(i) $\mathcal{M}, a \vDash \Pi\vec{\alpha}$ iff $\langle \overrightarrow{\alpha^{\mathcal{M},a}} \rangle \in I(\Pi)$.

(ii) $\mathcal{M}, a \vDash \alpha = \beta$ iff $\alpha^{\mathcal{M},a} = \beta^{\mathcal{M},a}$.

(iii) $\mathcal{M}, a \vDash \sim\phi$ iff $\mathcal{M}, a \nvDash \phi$.

(iv) $\mathcal{M}, a \vDash \phi \wedge \psi$ iff $\mathcal{M}, a \vDash \phi$ and $\mathcal{M}, a \vDash \psi$.

(v) $\mathcal{M}, a \vDash \exists x \phi$ iff for some $o \in D$, $\mathcal{M}, a[x \triangleright o] \vDash \phi$.

When ϕ is a formula open in x, we allow ourselves to write it as $\phi(x)$. Finally, we say that a formula is true on ϕ, or $\mathcal{M} \vDash \phi$, when it is true on $\mathcal{M}$ relative to every variable assignment over $\mathcal{M}$; and $\mathcal{M} \vDash \Gamma$ if $\mathcal{M} \vDash \phi$ for every $\phi \in \Gamma$. And $\Gamma \vDash \phi$ iff, if $\mathcal{M} \vDash \Gamma$, then $\mathcal{M} \vDash \phi$.

It is routine to show that if a closed sentence is true on $\mathcal{M}$ relative to any variable assignment then it is true on $\mathcal{M}$ relative to all of them.

3.2 Models and Intensions

Models are *extensional*: they assign referents to names, and extensions to predicates. But when dealing with possible worlds semantics, we are interested in *intensional* assignments. These are often taken as functions from possible worlds to extensions. For instance, the extension of 'dog' is the set of all the dogs; the intension of 'dog' is the function from each possible world w to the set of things that are dogs in w.

Our 'possible worlds' don't have things in them — the set W is really just an index set — but if every world in W is associated with a model, we can use things in the model instead. For instance, if w is a world and $\mathcal{M}^w$ is its associated model, then there is the set of things that are in the extension of 'dog' in $\mathcal{M}^w$. So, for a set of worlds W, if each one has an associated model, we can have something like an 'intension' defined on these worlds and models.

First, a definition to make this idea of 'an associated model' precise.

Definition.

A **intensional space** over a language $\mathscr{L}$ is an ordered pair $\mathfrak{W} = \langle W, i \rangle$; W is a set of worlds and i a function where, for every $w \in W$, $i(w)$ is a model of $\mathscr{L}$.

When $\mathcal{M} = i(w)$, we write it as '$\mathcal{M}^w = \langle D^w, I^w \rangle$', which will help streamline notation a bit. Notice that every model in an intensional space is a model of the *same language* as every other.

Now we can define the notion of an intension, relative to an intensional space.

Definition.

Let $\mathfrak{W}$ be an intensional space over $\mathscr{L}$, Π a predicate of $\mathscr{L}$, and α a name of $\mathscr{L}$. Then:

- The **$\mathfrak{W}$-intension of Π**, written $[\![\Pi]\!]_{\mathfrak{W}}$, is the function from W defined by $[\![\Pi]\!]_{\mathfrak{W}}(w) = I^w(\Pi)$.
- The **$\mathfrak{W}$-intension of α**, written $[\![\alpha]\!]_{\mathfrak{W}}$, is the function from W defined

by $[\![\Pi]\!]_{\mathfrak{W}}(w) = I^w(\alpha)$.

Notice something about the second definition. Kripke (1972) taught us that names are 'rigid designators': They designate the same thing in every possible world. Whatever the merits of Kripke's lesson for modal semantics generally, they don't apply here. Instead, we follow Carnap's (1956) development, making names pick out *individual concepts*: functions from worlds to individuals. Despite the superficial similarities, we're doing logic here, not modal semantics, and our choice is the only one that makes sense from our standpoint.

Notice also that we've only defined intensions for particular symbols. But the structure itself is general. There will be plenty of functions from worlds to objects that aren't associated, in an intensional space, with any particular name. We can generalize them as follows:

Definition.

If $\mathfrak{W}$ is an intensional space, then:

- An n-adic **$\mathfrak{W}$-relation** (or '$\mathfrak{W}$-property', when $n = 1$) is a function f from W where, for each $w \in W$, $f(w) \subseteq (D^w)^n$. (That is: $f(w)$ is a set of n-tuples drawn from $\mathcal{M}^w$'s domain.)
- A **$\mathfrak{W}$-individual concept** is a function g from W where, for any w, $g(w) \in D^w$.

Now we can see that the $\mathfrak{W}$-intension of any predicate is a $\mathfrak{W}$-property or relation, and the $\mathfrak{W}$-intension of a name is a $\mathfrak{W}$-individual concept, or '$\mathfrak{W}$-concept' for short.

3.3 *Intensions and Variable Assignments*

Back in section 1.1 we helped ourselves to the notion of 'the set of worlds at which ϕ is true', and wrote it as $[\![\phi]\!]_W$. But since different assignments of worlds to models will result in different sets of worlds, we should have instead relativized that notion to an intensional space. We can fix that now:

Definition.

If $\mathfrak{W}$ is an intensional space over $\mathscr{L}$ and ϕ is any sentence of $\mathscr{L}$, then $[\![\phi]\!]_{\mathfrak{W}} = \{w \in W : \mathcal{M}^w \vDash \phi\}$.

In other words: $[\![\phi]\!]_{\mathfrak{W}}$ is the set of worlds whose associated models make ϕ true.

When it comes to models, we define truth-in-a-model by first defining truth in a model relative to a variable assignment. This gives the notion some nice properties, and makes it easy to work with. In short, it lets us treat the semantic properties of complex expressions (relative to an assignment) as determined by the semantic properties of their parts (relative to an assignment). The idea is that, while we can't say whether a formula like 'Philosopher(x)' is true or false *simpliciter*, we can say that it is true or false relative to some assignment to the variable 'x'; and this is helpful in defining truth-in-a-model for quantified sentences.

We can do something similar with open formulas and propositions. Given an intensional space $\mathfrak{W}$, we know what W-proposition 'Philosopher(jason)' expresses: it's the set of worlds w where $\mathcal{M}^w$ makes that sentence true. But we can also ask, relative to some way of assigning an intension to the variable 'x', what W-proposition 'Philosopher(x)' expresses. And the answer isn't hard to come by. The assigned intension will be a $\mathfrak{W}$-individual concept that returns, at each $\mathcal{M}^w$, something in the extension of 'Philosopher' at $\mathcal{M}^w$.

More generally, let a **$\mathfrak{W}$-assignment** be a function A from variables to $\mathfrak{W}$-individual concepts. We want to say what it is for an open formula to $\mathfrak{W}$-express a proposition (that is, a set of worlds) relative to A. To do this, it's worth noting that we can think of A in any of three different ways.

Way 1: The way we already described it. A is a function from variables to $\mathfrak{W}$-individual concepts.

Way 2: As a function from variable-world pairs x, w to objects in D^w. More precisely, if A is a $\mathfrak{W}$-assignment in the sense of Way 1, we can define another function A' where $A'(x, w) = A(x)(w)$ — that is, the result of applying the function $A(x)$ to the world w. Then $A'(x, w)$ will be an object in D^w.

Way 3: As a function from worlds to variable assignments. If we have a Way-2 assignment A', we can define a function A'' by $A''(w) = a$ iff a is a variable assignment on $\mathcal{M}^w$ where $a(x) = A'(x, w)$. Conversely, if we have a set of variable assignments, one for each $\mathcal{M}^w$, we can staple them together into one big Way-3 $\mathfrak{W}$-assignment A''.

Technically, the functions A, A', and A'' are all different, but only for boring reasons. They all represent the same general idea, so we won't distinguish between them. If we write $A(x)$, we are talking about the intension that A assigns to x. If we write A^w, we are talking about the variable assignment that A'' assigns to w. And if we write something of the form $A(x, w)$, we are talking about the thing that A' assigns to x and w. They can all be thought of different presentations of the same underlying function of two arguments.

With this in mind, we can state more directly a notion of variable-relative expression.

Definition.

A proposition p ***W*-expresses $\boldsymbol{\phi}$ relative to $\boldsymbol{A}$** if and only if p is the set of worlds w where $\mathcal{M}^w, A^w \models \phi$. In symbols:

$$[\![\phi]\!]_{\mathfrak{W},A} = \{w \in W : \mathcal{M}^w, A^w \models \phi\}.$$

When we defined truth-on-a-model, we made use of the notion one variable assignment being just like another except for what was assigned to a certain variable. That is, if a is a variable assignment on $\mathcal{M}$ and o is in $\mathcal{M}$'s domain, then $a[o \triangleright x]$ is the variable assignment just like x except that it assigns o to x.

We can do something similar here. If A is a $\mathfrak{W}$-assignment and g a $\mathfrak{W}$-concept, we can consider the variable assignment $A[g \triangleright x]$, which is just like A except that it assigns the concept g to x. This is Way-1 thinking. But we can also convert it into Way-3 thinking. The idea here is that, at each world w, the variable assignment $A[g \triangleright x]^w$ is the assignment just like A^w except that it assigns $g(w)$ to x — that is, $A[g \triangleright x]^w = A^w[g(w) \triangleright x]$.

3.4 Some Facts About $\mathfrak{W}$-Propositions

With the foregoing, we can show that, for a fixed intensional space $\mathfrak{W}$, proposition 1 extends to

Proposition 5.

If A is any $\mathfrak{W}$-assignment,

- $[\![\top]\!]_A = W$
- $[\![\bot]\!]_A = \varnothing$
- $[\![\sim\phi]\!]_A = -[\![\phi]\!]_A$
- $[\![\phi \wedge \psi]\!]_A = [\![\phi]\!]_A \cap [\![\psi]\!]_A$
- $[\![\phi]\!]_A$ implies $[\![\psi]\!]_A$ iff $[\![\phi \to \psi]\!]_A = [\![\top]\!]_A$.

I won't prove all of these here, but I'll talk through one of them by way of illustration. To show that $[\![\sim\phi]\!]_A = -[\![\phi]\!]_A$, we need to show that for any world w in $\mathfrak{W}$, $w \in [\![\sim\phi]\!]_A$ if and only if it is not in $[\![\phi]\!]_A$. So we let w be an arbitrary world. If it is in $[\![\sim\phi]\!]_A$, then $\mathcal{M}^w, A^w \vDash \sim\phi$, which means $\mathcal{M}^W, A^w \nvDash \phi$. But this means that $w \notin [\![\phi]\!]_A$. When w is not in $[\![\sim\phi]\!]_A$, a mirror-image of that argument tells us that w is in $[\![\phi]\!]_A$. The other instances of our expanded version of proposition 1 are shown in analogous ways.

Something else we can also show:

Proposition 6.

If $\phi \vDash \psi$, then for any $\mathfrak{W}$-assignment A, then ϕ implies ψ. (That is, $[\![\phi]\!]_A \subseteq [\![\psi]\!]_A$.)

The idea here is similar. If w is a world in $[\![\phi]\!]_A$, then $\mathcal{M}^w, A^w \vDash \phi$. But if $\phi \vDash \psi$, this means that $\mathcal{M}^w, A^w \vDash \psi$, too, so $w \in [\![\psi]\!]_A$. Proposition 6 can be very useful in streamlining proofs to come later.

Notice, in passing, that proposition 6's converse isn't in general true. It may turn out that there are countermodels to the argument from ϕ to ψ, but none of the worlds in $\mathfrak{W}$ get paired with any of them. In that case, $[\![\phi]\!]_A$ might $\mathfrak{W}$-imply $[\![\psi]\!]_A$ even though $\phi \nvDash \psi$.

The above facts are all, in some sense, propositional corollaries of truth-functional facts. But facts about how model theory deals with quantification and variable-binding also have propositional corollaries. First, we have

Proposition 7.

If A is a $\mathfrak{W}$-assignment and g a $\mathfrak{W}$-concept, then

$$[\![\phi]\!]_{A[g \triangleright x]} \subseteq [\![\exists x \phi]\!]_A.$$

The proof relies on the equivalence of $A[g \triangleright x]^w$ and $A^w[g(w) \triangleright x]$ noted at the end of the last section, and runs as follows.

Proof.

Suppose $w \in [\![\phi]\!]_{A[g \triangleright x]}$. Then $\mathcal{M}^w, A[g \triangleright x]^w \vDash \phi$, so $\mathcal{M}^w, A^w[g(w) \triangleright x] \vDash \phi$, and hence $\mathcal{M}^w, A^w \vDash \exists x \phi$. Thus $w \in [\![\exists x \phi]\!]_A$. $\therefore$

Notice that, while this is close to proposition 6, it doesn't quite follow from it, because in proposition 6 both $\mathfrak{W}$-expressions were relativized to the same variable assignment, and here they are not.

For our last observation, start by noting that, when we were doing model theory, it helped to have the concept of a term's (name or variable's) *denotation* on a model, relative to a variable assignment. This is just whatever that term names, relative to that assignment. Likewise, we can have the concept of what a term (name or variable) *expresses* on an intensional space $\mathfrak{W}$, relative to some $\mathfrak{W}$-assignment.

In fact, we already have half of this: if α is a name, we have already defined $[\![\alpha]\!]_{\mathfrak{W}}$, which is the function where $[\![\alpha]\!]_{\mathfrak{W}}(w) = I^w(\alpha)$. We can extend this to allow for variables, too, where the $\mathfrak{W}$-concept expressed by a variable x relative to an assignment A is just $A(x)$. More precisely,

Definition.

Let $\mathfrak{W}$ be an intensional space over $\mathscr{L}$, α a term of $\mathscr{L}$, and A a $\mathfrak{W}$-assignment.

$$(\!|\alpha|\!)_{\mathfrak{W},A} = \begin{cases} [\![\alpha]\!]_{\mathfrak{W}} \text{ if } \alpha \text{ is a name and} \\ A(\alpha) \text{ if } \alpha \text{ is a variable.} \end{cases}$$

As with other such definitions we drop the '$\mathfrak{W}$' when no confusion arises. Now we can show the following substitution principle:

Lemma 8.

$$[\![\phi(\vec{x})]\!]_{A[\overrightarrow{(\!|\alpha|\!)_A} \triangleright \vec{x}]} = [\![\phi(\vec{\alpha})]\!].$$

Proof.

For this, note that if α is a variable and $w \in W$, $(\!|\alpha|\!)_A(w) = A^w(\alpha) = \alpha^{\mathcal{M}^w, A^w}$, and if α

is a constant, $(\!|\alpha|\!)_A(w) = I^{\mathcal{M}^w}(\alpha) = \alpha^{\mathcal{M}^w, A^w}$.

$$\begin{aligned} w \in [\![\phi(x)]\!]_{A[\overrightarrow{(\!|\alpha|\!)_A} \rhd \vec{x}]} &\text{ iff } \mathcal{M}^w, A[\overrightarrow{(\!|\alpha|\!)_A} \rhd \vec{x}]^w \vDash \phi(\vec{x}) \\ &\text{ iff } \mathcal{M}^w, A^w[\overrightarrow{(\!|\alpha|\!)_A(w)} \rhd \vec{x}] \vDash \phi(\vec{x}) \\ &\text{ iff } \mathcal{M}^w, A^w[\overrightarrow{\alpha^{\mathcal{M}^w, A^w}} \rhd \vec{x}] \vDash \phi(\vec{x}) \\ &\text{ iff } \mathcal{M}^w, A^w \vDash \phi(\vec{\alpha}) \\ &\text{ iff } w \in [\![\phi(\vec{\alpha})]\!]_A \end{aligned}$$

The third biconditional relies on the identity noted above, and the penultimate one on the basic model-theoretic fact that $\mathcal{M}, a[\alpha^{\mathcal{M},a} \rhd x] \vDash \phi(x)$ iff $\mathcal{M}, a \vDash \phi(\alpha)$. ∴

Lemma 8 has a number of immediate corollaries. For instance, if two formulas differ only in a shuffling of the variables, and two $\mathfrak{W}$-assignments differ only in ways that compensate for those variable-shufflings, then the first formula will express the same thing on the first assignment that the second formula expresses on the second.

All of these details show a number of very tight connections between a sentence's truth-on-a-model, and the various principles governing it, and its expression of a $\mathfrak{W}$-proposition. In many ways these connections shouldn't be surprising; $\mathfrak{W}$-expression is defined in terms of truth-on-a-model, so we should expect many of the principles governing it to somehow 'project up' into the propositions. On the other hand, it is useful to have these principles ready to go. We may need them for some heavier lifting in what's to come.

3.5 *What a Propositional Theory 'Says'*

Everything so far has just been about $\mathfrak{W}$-propositions and what sentences express them. But we can also ask how these notions relate to particular propositional *theories*. Suppose, for instance, that T is a set of $\mathfrak{W}$-propositions and includes [[Philosopher(jason)]]. If so, there is a very clear sense in which the theory 'says' that Jason is a philosopher.

The scare quotes are important. [[Philosopher(jason)]] is just some old set of things we *call* 'worlds', picked out by some fancy set-theoretic apparatus. If there are genuine propositions, and they are sets of genuine possible worlds, there is no reason to think [[Philosopher(jason)]] is one of them, and much less the one that counts as the proposition that Jason is a philosopher. But insofar as we pretend that the 'worlds' in W are genuine possible worlds, and pretend that $\mathcal{M}^w \vDash \phi$ if and only if ϕ is true in the world w, then we will want to pretend that [[Philosopher(jason)]] is the proposition that Jason is a philosopher. Since the pretense helps us understand the underlying model theory, we'll stick with it here.

This picture of 'saying' applies to open formulas and $\mathfrak{W}$-assignments, too. Suppose that $[\![\text{Philosopher}(x)]\!]_A$ is in T. What does this 'say'? Roughly, that the $\mathfrak{W}$-concept assigned to x by A is a philosopher. More compactly: $A(x)$ is a philosopher. When this happens, we will say that $A(x)$ ***T*-satisfies** 'philosopher'. More precisely, and expanding for multiple variables:

Definition.

If $\vec{g}$ are all $\mathfrak{W}$-individual concepts and T a set of $\mathfrak{W}$-propositions, then $\vec{g}$ $\vec{x} \triangleright$ ***T*-satisfies** ϕ iff for any $\mathfrak{W}$-assignment A,

$$[\![\phi]\!]_{\mathfrak{W},A[\vec{g}\triangleright\vec{x}]} \in T.$$

Here, the idea is that some $\mathfrak{W}$-concepts T-satisfy some formula if the theory T 'says' that those concepts are the way the formula requires. For instance, if by assigning x to g and y to h, 'x loves y' gets us a proposition that is in T, then T 'says' that g loves h, and these concepts $x, y \triangleright T$-satisfy 'loves'. Often, it will be obvious (or won't matter) which variables go with which objects, in which case we drop the specification of variables and just say that $\vec{g}$ T-satisfies ϕ.

4 Łos's Theorem

All of this work has a point. We are going to prove that, if a set of sentences corresponds, in the right way, to a filter, then it has a model. More precisely, if $\mathfrak{W}$ is an intensional space over $\mathscr{L}$ and T an ultrafilter over $\mathfrak{W}$, then there is a model that makes true ϕ if and only if $[\![\phi]\!]_{\mathfrak{W}}$ is in T.

Before describing how the proof goes, let's remind ourselves of the technique Henkin (1949) used to prove completeness, which has now become the standard textbook method. The proof shows that if some theory is syntactically consistent, it has a model. We start with a syntactically consistent theory, and then expand it to a consistent theory that is negation-complete (via Lindenbaumm's lemma) and includes every instance of a witness axiom. We then build a model *out of* this theory: Roughly, the constants become the model's domain (with the witness axiom ensuring there are enough constants to go around) and we read off the extension's predicates from the theory itself. For instance, if 'Fc' is in the theory, we put 'c' in the extension of 'F', and if '$\sim Fd$' is in the theory, we keep 'd' out.

The proof of Łos's theorem is, in many ways, the propositional counterpart of this Henkin-style completeness proof. We start with a filter T on an intensional space $\mathfrak{W}$, which is our propositional analogue of a consistent theory. Then we extend it to an ultrafilter T', which is our propositional analogue of a negation-complete extension of T. Then, instead of using constants, we build our model out of $\mathfrak{W}$-individual concepts. Roughly, the $\mathfrak{W}$-concepts make up the model's domain. (We need no witness axiom to ensure there is enough this time.) And we read predicates' extensions from the theory itself: if T' 'says' that g satisfies 'F', then we put g in the extension of 'F', and if T' says that h does not satisfy 'F', then we keep g out.

If the language has an identity predicate, complications arise. In the case of first-order completeness, the problem is this: c and d may be different constants, but the theory may include the sentence '$c = d$'. Given the way the meaning of '='

is hardwired into the model, if c and d are two separate entities in the domain, the model will not be able to make '$c = d$' true. In the case of a propositional theory, the problem is similar. $\mathfrak{W}$-concepts g and h may T'-satisfy '$x = y$', even if they are different functions. Given the way the meaning of '$=$' is hardwired into the model, if we build our model by putting g and h in separately, we will have a hard time making sure the right formulas come out true on the right variable assignments.

The solution in both cases is the same: Rely on equivalence classes. In the Henkin proof, we call two constants c and d equivalent if the maximal consistent set includes '$c = d$'. In the proof of Łos's theorem, we say that two $\mathfrak{W}$-individual concepts are equivalent if the theory T' says they are identical; that is, if they T'-satisfy '$x = y$'. Once we do this, then rather than making our domains out of constants or $\mathfrak{W}$-concepts, we make them out of *equivalence classes* of such.

Let's set aside our comparisons with Henkin's proof and make our own more precise. We start with another definition.

Definition.

Let $\mathfrak{W}$ be an intensional space and T an ultrafilter on $\mathfrak{W}$. Then, if g and h are $\mathfrak{W}$-individual concepts, $g \equiv_{T,\mathfrak{W}} h$ iff g and h T-satisfy '$x = y$'.

The latter, recall, is equivalent to '$[\![x = y]\!]_{\mathfrak{W},A[g \triangleright x][h \triangleright y]} \in T$'. In general, when we have an relation $\equiv_{T,\mathfrak{W}}$ we know what the relevant T and $\mathfrak{W}$ are, so we leave them off.

With the resources of the previous section, it isn't difficult to verify that $\equiv_{T,\mathfrak{W}}$ is reflexive, transitive, and symmetric, and therefore an equivalence relation. For instance, when it comes to symmetry, the fact that $x = y \vDash y = x$ combines with proposition 5 and lemma 8 to tell us that the $\mathfrak{W}$-proposition witnessing $g \equiv h$ implies one witnessing $h \equiv g$. Since T is a filter, the latter is in T as well, so $h \equiv g$. (In the case of transitivity, we also use conjunction-closure so we can get a $\mathfrak{W}$-proposition expressing '$x = y \wedge y = z$' in T, to then imply one expressing '$x = z$'.)

Once we've done that, we can define equivalence classes over $\mathfrak{W}$-concepts; $[g]^T_{\mathfrak{W}}$ is the $\equiv_{T,\mathfrak{W}}$-class that contains g. As with other cases, we'll often drop the -scripts. Now, finally, we can define our model.

Definition.

Let $\mathfrak{W}$ be an intensional space and T an ultrafilter on $\mathfrak{W}$. Then define a model $\mathcal{M}^T = \langle D^T, I^T \rangle$ where

(i) $D^T = \{[g] : g \text{ is a } \mathfrak{W}\text{-individual concept}\}$,
(ii) $I^T(\alpha) = [[\![\alpha]\!]_{\mathfrak{W}}]$, and
(iii) $I^T(\Pi) = \{\overrightarrow{[g]} : \vec{g}\ T\text{-satisfy } \Pi\vec{x}\}$.

This model is called the **ultraproduct** of T on $\mathfrak{W}$, and is often written

$$\prod_{w \in W} \mathcal{M}^w / T.$$

To simplify notation, though, we will stick with $\mathcal{M}^T$.

Now that we've defined the ultraproduct, we're almost in a position to state what we need to prove. We just need one last piece of machinery. Where A is any $\mathfrak{W}$-assignment, we let $[A]$ be the assignment on $\mathcal{M}^T$ where $[A](x) = [A(x)]$. In other words, if A assigns some concept to x, then $[A]$ assigns its equivalence class to x.

Now we can prove, by an induction on the complexity of formulas, that

Theorem 9. Łos's Theorem

$$\mathcal{M}^T, [A] \models \phi \text{ iff } [\![\phi]\!]_A \in T.$$

4.1 *The Proof*

Before we can get on with proving Łos's Theorem, we need to verify that our definition of 'ultraproduct' was legit. Parts (i) and (ii) of the definition are going to be fine, no matter what. But (iii) only makes sense if the following holds:

Proposition 10.

If $g_i \equiv h_i$ for each i, $\vec{g}$ T-satisfy $\Pi\vec{x}$ iff $\vec{h}$ T-satisfy $\Pi\vec{x}$.

If proposition 10 didn't hold, then clause (iii) could give us contradictory instructions. For we could have $\overrightarrow{[g]} = \overrightarrow{[h]}$, $\vec{g}$ T-satisfying $\Pi\vec{x}$, and $\vec{h}$ not T-satisfying $\Pi\vec{x}$. The first of these would tell us to put $\overrightarrow{[g]}$ in $I^T(\Pi)$ and the second would tell us to keep it out, and given the identity we can't do both.

Fortunately, though, this can't happen.

Proof.

Suppose $g_i \equiv h_i$ for each i. Then (by appeal to lemma 8 to swap around variables) there is a $\mathfrak{W}$-assignment A where $A(x_i) = g_i$ and $A(y_i) = h_i$ for each i, and $[\![x_i = y_i]\!]_A \in T$. Now suppose that $\vec{g}$ $\vec{x}/T$-satisfies $\Pi\vec{x}$. Then, given how we defined A, $[\![\Pi\vec{x}]\!]_A \in T$. But since T is closed under conjunction,

$$[\![\Pi\vec{x} \wedge x_1 = y_1 \wedge \ldots \wedge x_n = y_n]\!]_A \in T.$$

However, we have that

$$\Pi\vec{x} \wedge x_1 = y_1 \wedge \ldots \wedge x_n = y_n \models \Pi\vec{y}.$$

By proposition 5 and T's closure under implication,

$$[\![\Pi\vec{y}]\!]_A \in T.$$

Given our specification of A, $\vec{h}$ $\vec{y}/T$-satisfies $\Pi\vec{y}$, and so by lemma 8 they also $\vec{x}/T$-satisfy $\Pi\vec{x}$. (The other direction is precisely the same.) $\therefore$

The proof of proposition 10 relied on T being a filter, not on its being an ultrafilter. As a result, our definition of $\mathcal{M}^T$ would have defined a model even if T was merely a filter instead of an ultrafilter. In that case, $\mathcal{M}^T$ is called a **reduced product** rather than an ultraproduct. However, the proof of Łos's Theorem *does* require T to be an ultrafilter, so for our purposes ultraproducts are more useful than reduced products.

Now that we know our ultraproduct is legit, we can prove

Lemma 11.

If A is a $\mathfrak{W}$-assignment, then

(i) $\overrightarrow{\alpha^{\mathcal{M}^T,[A]}} \in I^T(\Pi)$ iff $[\![\Pi\vec{\alpha}]\!]_A \in T$, and

(ii) $\alpha^{\mathcal{M}^T,[A]} = \beta^{\mathcal{M}^T,[A]}$ iff $(\!|\alpha|\!)_A \equiv (\!|\beta|\!)_A$.

Proof.

Note that for any term α, $\alpha^{\mathcal{M}^T,[A]} = [(\!|\alpha|\!)_A]$.

For (i), right-to-left: Suppose $[\![\Pi\vec{\alpha}]\!]_A \in T$. Then $[\![\Pi\vec{z}]\!]_{A[\overrightarrow{(\!|\alpha|\!)_A \triangleright z}]} \in T$ by lemma 8 and so $\overrightarrow{(\!|\alpha|\!)_A}$ T-satisfies $\Pi\vec{z}$. So by construction, $\overrightarrow{[(\!|\alpha|\!)_A]} = \overrightarrow{\alpha^{\mathcal{M}^T,[A]}} \in I^T(\Pi)$.

For (i), left-to-right: Suppose that $\overrightarrow{\alpha^{\mathcal{M}^T,[A]}} \in I^T(\Pi)$. Then for some $\vec{g}$, $[g_i] = \alpha^{\mathcal{M}^T,[A]}$. By construction, $\vec{g}$ T-satisfy $\Pi\vec{z}$. But each $[g_i] = [(\!|\alpha_i|\!)_A]$, so $g_i \equiv (\!|\alpha_i|\!)_A$. By proposition 10, $\overrightarrow{(\!|\alpha_i|\!)_A}$ T-satisfy $\Pi\vec{z}$, so by lemma 8, $[\![\Pi\vec{\alpha}]\!]_A \in T$.

For (ii): $\alpha^{\mathcal{M}^T,[A]} = \beta^{\mathcal{M}^T,[A]}$ iff (by the above observation) $[(\!|\alpha|\!)_A] = [(\!|\beta|\!)_A]$ iff $(\!|\alpha|\!)_A \equiv (\!|\beta|\!)_B$. $\therefore$

Now we are in a position to prove theorem 9. The proof is by induction on the complexity of formulas.

Proof.

The base case is immediate by lemma 11.

For the induction step, we suppose that for any ψ less complex than ϕ, $\mathcal{M}^T, [A] \vDash \psi$ iff $[\![\psi]\!]_A \in T$.

For negation, $\mathcal{M}^T, [A] \vDash {\sim}\psi$ iff $\mathcal{M}^T, A \nvDash \psi$ iff (by the induction hypothesis) $[\![\psi]\!]_A \notin T$ iff (since T is an ultrafilter) $-[\![\psi]\!]_A \in T$ iff $[\![{\sim}\psi]\!]_A \in T$.

For conjunction, $\mathcal{M}^T, [A] \vDash \psi \wedge \chi$ iff $\mathcal{M}^T, [A] \vDash \psi$ and $\mathcal{M}^T, [A] \vDash \chi$ iff (by the induction hypothesis) $[\![\phi]\!]_A \in T$ and $[\![\psi]\!]_A \in T$ iff (by the fact that T is a filter* and so meets condition (ii*)) $[\![\phi]\!]_A \cap [\![\psi]\!]_A \in T$ iff $[\![\phi \wedge \psi]\!]_A \in T$.

For quantification, we go in two directions. First, suppose that $\mathcal{M}^T, [A] \vDash \exists x\psi$. Then for some $\mathfrak{W}$-concept g, $\mathcal{M}^T[A]([g] \triangleright x) \vDash \psi$, so $\mathcal{M}^T, [A(g \triangleright x)] \vDash \psi$. By the induction hypothesis, $[\![\psi]\!]_{A[g \triangleright x]} \in T$. Thus by proposition 7 and T's closure under implication, $[\![\exists x\psi]\!]_A \in T$.

For the other direction, suppose $[\![\exists x\psi]\!]_A \in T$. Since T is a filter this proposition is not empty. Let $w \in [\![\exists x\psi]\!]_A$; then $\mathcal{M}^w, A^w \vDash \exists x\psi$, so for some $e_w \in D^w$, $\mathcal{M}^w, A^w[e_w \triangleright x] \vDash \psi$. Let $g(w) = e_w$ for each $w \in [\![\exists x\psi]\!]_A$. (Such a g is guaranteed by the axiom of choice.) Then each $w \in [\![\psi]\!]_{A[g\triangleright x]}$. By the induction hypothesis, $\mathcal{M}^T, [A[g \triangleright x]] \vDash \psi$, so $\mathcal{M}^T, [A][[g] \triangleright x] \vDash \psi$, so $\mathcal{M}^T, [A] \vDash \exists x\psi$. $\therefore$

4.2 *Compactness: An Application*

So how do we *use* Łos's theorem? One thing we can do with it is to prove

Theorem 12. Compactness

If every finite subset of Γ has a model, then Γ has a model.

Of course, this can be proven in other ways; if we have a complete proof system for which the analogous claim about consistency holds, then we can appeal to that analogous claim plus the completeness theorem. But one advantage of doing things this way is that it proves compactness in a purely model-theoretic way: No detour through proof theory is required.

The idea behind the proof is very simple. We assume that we have a Γ where each of its finite subsets has a model. We use these subsets and their models to construct an intensional space $\mathfrak{W}$ and a filter, where for each ϕ in Γ, $[\![\phi]\!]_{\mathfrak{W}}$ is in the filter. Then we extend the filter to an ultrafilter and apply Łos's theorem. This gets us a model of every sentence in Γ.

For the space $\mathfrak{W}$, let our worlds be the finite subsets of Γ themselves. But we are assuming for the proof that each finite subset — that is, each *world* — has a model. So we know that for every $w \in W$, there is a model $\mathcal{M}^w$ that corresponds to it. In fact, we know something stronger: $\mathcal{M}^w \vDash w$.

(I should note here that this last move requires the axiom of choice. That really isn't a big concession, since we're going to need the axiom of choice later to expand our propositional theory to an ultrafilter. But without the axiom of choice we might not be able to select a *single* model $\mathcal{M}^w$ to go with each w.)

At a first pass, we could just let our propositional theory to include $[\![\phi]\!]_{\mathfrak{W}}$ for each ϕ in Γ. But if we do this, it will be difficult to show that our theory is closed under conjunction. We know, though, not just that each sentence in Γ has a model, but also that each finite subset of Γ has a model, too. So, if ϕ and ψ are in Γ, then $\{\phi, \psi\}$ has a model, too, which will also be a model of $\phi \wedge \psi$. If we include not just $[\![\phi]\!]_{\mathfrak{W}}$ and $[\![\psi]\!]_{\mathfrak{W}}$ in our theory, but also $[\![\phi \wedge \psi]\!]_{\mathfrak{W}}$, we'll have a non-empty proposition that can help us with conjunction-completeness.

Of course, then we might have to repeat the process again for $\phi \wedge \psi$ and some other sentence χ in Γ, and then again for that conjunction, over and over. But rather than getting caught in a regress, we can just include, for any finite subset Δ of Γ, the conjunction of all the sentences in Δ. Since any two of these, conjoined, will be the conjunction of some *other* finite subset of Γ, by including all of these, we put ourselves in a good position to demonstrate the conjunction-closure of our theory.

More precisely: If Δ is any finite subset of Γ. let $\bigwedge \Delta$ be the conjunction of all the sentences in Δ. What we want is for our theory to include all propositions of the form $[\![\bigwedge \Delta]\!]_{\mathfrak{W}}$.

This helps with conjunction closure. But what about implication closure? Even if p is a set of worlds that includes $[\![\bigwedge \Delta]\!]_{\mathfrak{W}}$ for some finite Δ, it might not be of the form $[\![\bigwedge \Sigma]\!]_{\mathfrak{W}}$ for any finite subset Σ of Γ. To get around this, we simply include every proposition implied by one of these 'starter' propositions.

One way to do this is to first define the theory

$$T^- = \{[\![\bigwedge \Delta]\!]_{\mathfrak{W}} : \Delta \subseteq \Gamma \text{ is finite}\}$$

and then define

$$T = \{p : \text{ for some } q \in T^-, q \text{ implies } p\}.$$

Equivalently, we could define the whole thing all in one go, with:

$$T = \{p \subseteq W : \text{for some finite } \Delta \subseteq \Gamma, [\![\bigwedge \Delta]\!]_{\mathfrak{W}} \text{ implies } p\}.$$

Either way we get the same thing.

Now we just have to prove that T is a filter. It's clearly closed under implication, just by its definition. To show that it's closed under conjunction, suppose that p and q are both in T. Then for some Δ and Σ, $[\![\bigwedge \Delta]\!]$ and $[\![\bigwedge \Sigma]\!]$ are both in T, where the former implies p and the latter implies q. But if Δ and Σ are both finite, then $\Delta \cap \Sigma$ is also finite, so $[\![\bigwedge(\Delta \cap \Sigma)]\!] = [\![\bigwedge \Delta]\!] \cap [\![\bigwedge \Sigma]\!]$ is in T. But $[\![\bigwedge \Delta]\!] \cap [\![\bigwedge \Sigma]\!]$ will imply $p \cap q$; and since T is closed under implication, this means $p \cap q \in T$.[3]

We also have to show that T doesn't contain $[\![\bot]\!]$. Note that the only proposition that implies $[\![\bot]\!]$ is $[\![\bot]\!]$ itself. (Implication is just subsethood, and the empty set's only subset is the empty set itself.) So if the empty set is in T, it must be because it is $[\![\bigwedge \Delta]\!]$ for some finite subset Δ of Γ. But each of these finite subset has a model $\mathcal{M}^w$ for some w in the intensional space. So none of these are empty (they each have, at a minimum, w itself). So T doesn't contain $[\![\bot]\!]$.

This means T is a filter. Having established that, the rest of the work has already been done. T can be expanded to an ultrafilter T', and it will have an ultraproduct $\mathcal{M}^T$. But for any sentence ϕ, $\mathcal{M}^T \vDash \phi$ iff $[\![\phi]\!] \in T$. And if ϕ is any sentence in Γ, then $[\![\phi]\!]$ is in T.[4] So $\mathcal{M}^T \vDash \Gamma$. Thus, if every finite subset of Γ has a model, so does Γ.[5]

References

Boolos, G., J. P. Burgess, and R. C. Jeffrey (2002). *Computability and Logic* (4th ed.). Cambridge University Press.

[3]The relevant principle is: if $p \subseteq r$ and $q \subseteq s$, then $p \cap q \subseteq r \cap s$.

[4]$\{\phi\}$ is a finite subset of Γ and $\bigwedge\{\phi\} = \phi$.

[5]Thanks to Caleb Camrud and Juan Comesaña for helpful comments and conversation on an earlier draft of this paper.

Carnap, R. (1956). *Meaning and Necessity* (2nd ed.). Chicago: University of Chicago Press.

Chang, C. C. and H. J. Keisler (1990). *Model Theory* (3rd ed.), Volume 73 of *Studies in Logic and the Foundation of Mathematics*. Amsterdam: Elsevier.

Enderton, H. B. (2001). *A Mathematical Introduction to Logic* (2nd ed.). San Diego: Harcourt Academic Press.

Gamut, L. T. F. (1991a). *Logic, Language, and Meaning*, Volume 1, Introduction to Logic. Chicago: University of Chicago Press.

Gamut, L. T. F. (1991b). *Logic, Language, and Meaning*, Volume 2, Intensional Logic and Logical Grammar. Chicago: University of Chicago Press.

Henkin, L. (1949). The completeness of the first-order functional calculus. *The Journal of Symbolic Logic 14*(3), 159–166.

Kripke, S. (1972). *Naming and Necessity*. Harvard University Press.

Mendelson, E. (1997). *Introduction to Mathematical Logic* (4th ed.). London: Chapman & Hall.

Sider, T. (2010). *Logic for Philosophy*. Oxford: Oxford Univerisity Press.