Understanding Ultrafilters as (Almost) Propositions

Jason Turner

December 6, 2019

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

> 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.

Here is the idea. We can think of models as analogous to possible worlds. Then sets of models are like sets of possible worlds — that is, propositions.[2] The set of models making ϕ true is like the proposition expressed by ϕ. Filters and ultrafilters are then kinds of sets of propositions. Filters are sets of propositions that are closed under implication and conjunction. Ultrafilters are filters that are also maximal — for every proposition, either it or its negation is in the ultrafilter.

Here's why this matters. 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.

[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.

[2] Or at least suitable surrogates for propositions. We'll leave that particular debate aside..

1 Fixing Ideas

Before we can get to work, we need to do some preliminary tidying. For one thing, models *aren't* possible worlds, at least not the way philosophers usually understand them. There are models that make '$\exists x(\mathsf{Red}(x) \wedge {\sim}\mathsf{Colored}(x))$' true, though I take it that no possible world contains a red-but-uncolored thing. This needn't detain us here; as long as we remember that models are analogous-to-but-distinct-from possible worlds, we shan't be led astray.

Something else should detain us, though. In general, there won't be any such thing as 'the set of all models that make ϕ true'. The reason is simple: there are too many models to form a set. In fact, a sentence that has any models at all has too many models to form a set.[3]

But we aren't really interested in sets of *all* models of ϕ anyway. I was fudging a bit when I said that. We're interested in sets of models of ϕ *drawn from a background set*. A filter is always a filter relative to some background or 'index' set I. Filters can still be thought of as sets of propositions, but the relation between these 'propositions' and sentences is more attenuated. If P is in a filter and we're tempted to call it 'the proposition that ϕ', that won't be thanks fo P containing all models of ϕ, but instead to it containing all models of ϕ that are *in* I. Since there will generally be lots of models of ϕ not in I, if we give in to temptation our terminology will mislead. So we need to get some better terminology.

We'll do that soon, in section 1.2. First, however, let's write down our model theory. There's nothing particularly new here, but any model theory has to make some strategic and some terminological choices. It will be good to fix ideas and avoid any ambiguity that might otherwise arise.

1.1 *Model Theory*

We will restrict our attention to first-order languages $\mathcal{L}$ with constants and predicates of fixed adicity, but no other non-logical symbols. We will also suppose that the primitive logical symbols are $\sim$, $\wedge$, $\exists$, and $=$, with other symbols defined in the usual way.

A **model** $\mathcal{M}$ is an ordered pair $\langle D_{\mathcal{M}}, I_{\mathcal{M}} \rangle$ of a non-empty domain $D_{\mathcal{M}}$ and an interpretation function $I_{\mathcal{M}}$. (Don't confuse $I_{\mathcal{M}}$, which is an interpretation function for the model $\mathcal{M}$, with I, which we will use for a set of models.) For any constant α in $\mathcal{L}$, $I_{\mathcal{M}}(\alpha) \in D$, and for any n-adic predicate Π in $\mathcal{L}$, $I_{\mathcal{M}}(\Pi) \subseteq (D_{\mathcal{M}})^n$.

For each model $\mathcal{M}$ of $\mathcal{L}$, **a variable assignment over $\mathcal{M}$** is a function from

[3]Here's an easy way to see this. Suppose ϕ has some model; call it $\mathcal{M}$. Let a be some element of $\mathcal{M}$. Then, for any ordinal α, let $\mathcal{M}_\alpha$ be the result of systematically replacing α for a. (If α is already in $\mathcal{M}$, the two trade places; if not, then α enters the model and a leaves.) For the most part each $\mathcal{M}_\alpha$ is distinct from $\mathcal{M}$; but there are as many $\mathcal{M}_\alpha$'s as their are ordinals, so there will be too many of them to form a set.

variables of $\mathcal{L}$ to elements of $D_{\mathcal{M}}$. 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 $I_{\mathcal{M}}(\alpha)$ if a constant and $a(\alpha)$ if a variable. If a is a variable assignment over $\mathcal{M}$ and $o \in D_{\mathcal{M}}$, $a[x \triangleright o]$ is the assignment just like a except that it maps x to o. (If $a(x) = o$, then $a[x \triangleright o] = a$.)

We now recursively define truth of an open formula on a model relative to a variable assignment, written $\mathcal{M}, a \vDash \phi$. The definition runs:

(i) $\mathcal{M}, a \vDash \Pi\alpha_1 \ldots \alpha_n$ iff $\langle \alpha_1^{\mathcal{M},a}, \ldots, \alpha_n^{\mathcal{M},a} \rangle \in I_{\mathcal{M}}(\Pi)$.

(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 $o \in D_{\mathcal{M}}$, $\mathcal{M}, a[x \triangleright o] \vDash \phi$.

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 $\phi \in \Gamma$. 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.

1.2 *I*-Propositions and *I*-Expression

A **space** of models I is a set of models all of the same language. An ***I*-proposition** is any subset of I. In possible-worlds semantics, the set of all worlds is the 'necessary proposition' and the empty set is the 'impossible proposition'. Accordingly, the empty set is the **impossible proposition** and the set I is the ***I*-necessary proposition**. (Different spaces of models have different 'necessary propositions', so we need to relativize that in our definition. Each space of models has the very same impossible proposition, though; there is only one empty set.)

If P is the subset of I that contains exactly the members of I that make a sentence ϕ true, we say that ϕ ***I*-expresses** P, and we write P as $[\![\phi]\!]_I$.[4] It should be clear that, if $\bot$ is any logical falsehood and $\top$ any logical truth, then $[\![\bot]\!]_I = \varnothing$ and $[\![\top]\!]_I = I$. But notice that the converse doesn't hold. It may be that $[\![\phi]\!]_I$ is empty because all the ϕ-models happen to lie outside I; and it may be that $[\![\psi]\!]_I$ contains all of I because the non-ψ-models happen to lie outside I.

In possible-worlds semantics, the conjunction of two propositions is their intersection, and the disjunction of two propositions is their union. That goes for I-propositions as well. It is vindicated in part by the following observation:

[4]Note that, at this point, we are presuming that the ϕ, ψ, etc. in '$[\![\phi]\!]_I$' are *sentences* — that is, closed formulas — and not formulas generally. We'll put off considering open formulas until section 3.1.2.

- $[\![\phi \wedge \psi]\!]_I = [\![\phi]\!]_I \cap [\![\psi]\!]_I$
- $[\![\phi \vee \psi]\!]_I = [\![\phi]\!]_I \cup [\![\psi]\!]_I$

But it goes beyond this. Some I-propositions won't be I-expressed by any sentence. Those I-propositions still have conjunctions and disjunctions, and their conjunctions and disjunctions are still given by intersections and unions.

What about negation? If we have sets X and Y where X is a subset of Y, the **complement** of X with respect to Y, written $Y - X$, is the set of all members of Y that are not in X. When the set Y is held fixed, we often just write $-X$. When it comes to I-propositions, the negation of an I-proposition P is its complement with respect to I, $I - P$, or just $-P$. (When we are talking about I-propositions, we only ever take complements with respect to I.) As before, this is vindicated by

- $[\![\sim\phi]\!]_I = -[\![\phi]\!]_I$.

And as before, this is only part of the story, since some propositions won't be I-expressed by any sentence but will still have negations.

Finally, in possible-worlds semantics, implication corresponds to subsethood: If P implies Q, then every P-world is a Q-world, so the set of P-worlds is a subset of the set of Q-worlds. Since our I-propositions are relativized to the space I, we don't quite get subsethood corresponding to (model-theoretic) entailment. Rather, we get subsethood corresponding to a sort of implication-with-respect-to-I:

- $[\![\phi]\!]_I \subseteq [\![\psi]\!]_I$ iff $[\![\phi \rightarrow \psi]\!]_I = I$.

As a consequence of this, if $\phi \vDash \psi$, then $[\![\phi]\!]_I \subseteq [\![\psi]\!]_I$. But the converse won't always hold. It may be that $\phi \nvDash \psi$ but all the countermodels to the inference lie outside of I.

2 Filters and Ultrafilters

Officially, filters and ultrafilters are purely set-theoretic constructs. A set doesn't have to be filled with I-propositions in order to be a filter or an ultrafilter. On the other hand, we're mainly concerned here with filters and ultrafilters *of I-propositions*. This is partially pragmatic — filters of I-propositions work well for the uses we want to put them to — and partially pedagogical. We philosophers tend to have better-developed intuitions about conjunction, implication, and negation than we do about intersection, subset-hood, and complementation.

As a compromise, I will do the following. In definitions, I will use set-theoretic formalism, but use the more familiar notions of conjunction, implication, and so on when writing things out longhand. It is good to remember that the definitions apply generally to any set-theoretical constructions whatsoever, but I will often talk as though they do not.

Additionally, although a few interesting results will be described, I will generally content myself with intuitive, I-proposition-theoretic explanations as to why various results hold. Full proofs will be reserved for an appendix.

2.1 *Filters*

A filter is any set of I-propositions that is closed under implication, closed under conjunction, and 'non-trivial', in the sense that it does not contain the impossible proposition. More precisely:

> Let I be any set. Then a **filter over *I*** is a non-empty set $\mathscr{F}$ of subsets of I which is:
>
> **Closed Under Implication:** If $S \in \mathscr{F}$ and $S \subseteq T \subseteq I$, then $T \in \mathscr{F}$.
>
> **Closed Under Conjunction:** If $S, T \in F$, then $S \cap T \in \mathscr{F}$.
>
> **Non-trivial:** $\emptyset \notin \mathscr{F}$.

Suppose we have a set S of I-propositions. Why care if it is a filter? One reason is that if S is a filter, then it is more-or-less a *consistent* set of I-propositions. For if it weren't consistent, then some I-propositions in S would imply an impossibility. And if some I-propositions in S implied the impossible proposition $\emptyset$, then by conjoining those I-propositions and looking at their implications we'd end up with the impossible proposition. Since filters are closed under conjunction and entailment and *don't* contain the impossible proposition, filers are consistent sets of I-propositions.

Calling filters 'more-or-less' consistent is advisable, and for at least two reasons. The first is pedantic. In formal contexts, 'consistent' is generally reserved for a proof-theoretic notion: a set of formulas Γ is consistent if and only if $\Gamma \nvdash \bot$. In the above discussion we are using a rough-and-ready, intuitive notion of 'consistent' that isn't really well-defined. Still, we more-or-less know what we mean, and the gloss does well enough to get the intuitive point across.

The second reason is more important. Strictly speaking, filters are closed under *finite* conjunction. (The definition tells us that they are closed under *pairwise* conjunction. Repeated applications will get us any finite conunction.) But the conjunction of an *infinite* set of I-propositions need not be in a filter. So, for all that has been said, a filter may in fact contain an inconsistent set of I-propositions, so long as that inconsistency is *essentially* infinite.[5]

You may think that, at least in cases of first-order logic, this will not be an issue thanks to the compactness theorem. That theorem says that if every finite subset of some set of formulas has a model, so does the entire thing. The I-proposition-theoretic analogue would say that if every finite subset of I-propositions is 'consistent', then so is the entire set. But whether that is true

[5] If we think omega-inconsistency is a kind of essentially infinite inconsistency, then it gives us precisely such a case, for there are intuitively omega-inconsistent filters. Cf. note 7.

or not, it isn't something we can appeal to now. We can't appeal to the I-proposition-theoretic analogue of compactness because we have not proved it. And we cannot appeal to the more ordinary model-theoretic compactness theorem because it is something we want to prove using ultrafilters later.

2.2 *Alternative Definitions*

Sometimes texbooks provide alternative definitions of filters. A reasonably common variant drops the non-triviality condition from the definition, and goes on to distinguish 'proper filters' from 'improper filters'.[6] This is mere teminological deviation. If we wanted, we could call non-empty sets of subsets of I that are closed under implication and conjunction 'proto-filters'. Then our 'proto-filters' would correspond to their 'filters' and our 'filters' would correspond to their 'proper filters'. (Their 'improper filters' would be our 'proto-filters that are not filters.')

A different definition for 'filter' (in *our* 'proper' sense) is also possible:

> Let I be any set. Then a **filter* over *I*** is a non-empty set $\mathscr{F}$ of subsets of I which is:
>
> **Conjunction Complete:** If $S, T \subseteq I$, then $S \cap T \in \mathscr{F}$ iff both $S \in \mathscr{F}$ and $T \in \mathscr{F}$.
>
> **Non-universal:** $\mathscr{F}$ does not contain all subsets of I.

It takes a bit of thinking to see that the two definitions are equivalent, but they in fact are. (Proof is in the appendix.)

2.3 *Ultrafilters*

Suppose that $\mathscr{F}$ is a filter over I. Then $\mathscr{F}$ might (or might not) also be:

> **Maximal:** No filter over I (properly) extends $\mathscr{F}$. That is, if $\mathscr{F}'$ is a filter over I and $\mathscr{F} \subseteq \mathscr{F}'$, then $\mathscr{F} = \mathscr{F}'$.
>
> **Prime:** If $P, Q \subseteq I$, then $P \cup Q \in \mathscr{F}$ iff either $P \in \mathscr{F}$ or $Q \in \mathscr{F}$.
>
> **Ultra:** For any $P \subseteq I$, either $P \in \mathscr{F}$ or $-P \in \mathscr{F}$.

If it's maximal, then it is in a sense as 'big' a filter over I that there can be. If it is prime, then it is disjunction-complete: it contains a disjunction precisely when it contains at least one of the disjuncts. And if it is ultra, it is negation-complete: for every I-proposition, it contains either it or its negation.

A filter that has any of these properties has all three; they are provably equivalent. Such filters are **ultrafilters**.

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

Proving that these are equivalent is left to the appendix, but here's the intuitive idea. First, if a filter $\mathscr{F}$ was ultra, it would already contain each I-proposition or its negation. The only way to expand it would be to add a new I-proposition P; but if P wasn't already in the filter $-P$ would have been, making the expanded thing no longer a filter. So if it's ultra, it's maximal.

Next, suppose $\mathscr{F}$ is maximal. Each I-proposition implies any disjunction of it with anything else, so if P is in $\mathscr{F}$, so is $P \cup Q$. Conversely, suppose that $P \cup Q$ is in $\mathscr{F}$. If neither of P or Q were, then at least one could be added, meaning that $\mathscr{F}$ wasn't maximal after all. Since it is maximal, either P or Q is in it. So if $\mathscr{F}$ is maximal, it is prime.

Finally, suppose $\mathscr{F}$ is prime. Note that $I = P \cup -P$ for any I-proposition P. Since I is in $\mathscr{F}$, either P or $-P$ must be true by disjunction-closure. So $\mathscr{F}$ is ultra.

So we're able to start with a filter's being ultra, move through the other two properties, and end up back with its being ultra again. That's enough to show the three are equivalent.

2.4 *The Central Theorem on Ultrafilters*

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

Theorem 1 (The Central Theorem on Ultrafilters) *Every filter can be extended to an ultrafilter.*

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

Its proof is similar to the proof of Lindenbaumm's Lemma, in which a (proof-theoretically) consistent theory (that is, set of sentences) is expanded to 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 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 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 I-propositions. Then we take our starting filter and go through each of these I-propositions in order. If the result of adding an I-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 I-proposition, and thanks to a few facts about how set theory works, for each I-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 an I-proposition to a filter and closing the result up under conjunction and implication'. We do this with **additions**: if $\mathscr{F}$ is a filter, then $\mathscr{F} + X$ is, intuitively, the result of adding X to $\mathscr{F}$ and closing the result up under conjunction and implication.

We can define this officially by:

Addition: $\mathscr{F} + X = \{Y \subseteq I : \text{for some } A \in \mathscr{F}, A \cap X \subseteq Y\}$.

As with some of our other definitions, it takes a bit of fluency with set theory to see how this does the job we want it to. Here's the basic idea. I is in any filter, and (since X is an I-proposition) $X = I \cap X$. So I is the filter that stands as witness to X's meeting the condition needed to be in $\mathscr{F} + X$. Likewise, if P is already in $\mathscr{F}$, then it's implied by the conjunction of it with X, so it stays in $\mathscr{F} + X$. So this tells us that $\mathscr{F} + X$ contains X plus everything in $\mathscr{F}$. Closure under conjunction and implication then comes from the way that those two notions are embedded in the definition of $\mathscr{F} + X$. (This is all more carefully verified in the appendix.)

This is then the notion that we use in proving the Central Theorem. When we get to each I-proposition P in our well-ordering, we trade in our previous filter $\mathscr{F}$ for the new one $\mathscr{F} + P$ if that is, in fact, a filter. Happily, it's possible to show that, if $\mathscr{F}$ is a filer, then $\mathscr{F} + P$ is a filter if and only if $-P$ is not already in $\mathscr{F}$. This means that, by the time we're done, for each I-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 I is infinite, the axiom of choice must be assumed so that the I-propositions can be well-ordered. (Since the set of all I-propositions is the powerset of I, even if I is only countably infinite, there will uncountably many I-propositions.) The Central Theorem — and indeed, any result that follows on from it — should be understood as requiring choice.

3 Łos's Theorem

All this work about filters has a point: Ultrafilters give us a powerful tool for proving that certain kinds of models exist. The theorem says, in essence, that if I is a space of models and $\mathscr{F}$ is an ultrafilter over them, then there is a model $\mathcal{M}^I_{\mathscr{F}}$, called the **ultraproduct of $\mathscr{F}$ over I**, where

$$[\![\phi]\!]_I \in \mathscr{F} \text{ iff } \mathcal{M}^I_{\mathscr{F}} \vDash \phi.$$

(Actually, Łos's theorem says slightly more even than this, as we will soon see.) Note that, thanks to its definition, the ultraproduct won't be in the space I. It's possible that there is some other model with the above property in I; but it's also possible that there isn't.[7] So Łos's theorem is genuinely informative. It tells us about models with properties that might not be had by any model in the space that generates it.

Łos's theorem is proved by simply *building* the ultrafilter out of the resources that I provides. It is not unlike the standard Henkin-style proofs of completeness. In those proofs, we start with a consistent theory, expand it to a maximal consistent theory (via Lindenbaumm's Lemma), and then build a model from it. In the model, the domain is built out of constants, and the rest of the model is 'read off' of the maximal consistent set.

Something similar happens here. We start with a consistent set of I-propositions (a filter) and expand it to a maximal consistent set (an ultrafilter, via the Central Theorem). Then we build a model from it, where we 'read off' the model from the ultrafilter. However, we first have to figure out what the domain will be made of. If our construction really does mirror the Henkin-style one, then we will want to find some kind of objects that stand to I-propositions as names stand to sentences. We introduce these next.

3.1 *I-Concepts and I-Satisfaction*

3.1.1 I-Concepts

In modal possible-worlds semantics, what stands to propositions as names stand to sentences? Your answer will depend on your other commitments. Since Kripke (1972), orthodoxy has held that names are *rigid designators*: they denote the same thing in every possible world. As a result, the proposition 'Jason is a philosopher' is true at (and therefore is the set of) all and only the worlds where *I* — me, the very thing! — am a philosopher. From this perspective, the proposition-theoretic correlates of names are individuals.

The Kripkean orthodoxy contrasts with an earlier tradition (stemming back to Carnap 1956) treating names instead as *individual concepts*: functions from worlds to individuals. On the Carnapian view, 'Jason' doesn't have to denote *me* at every possible world; at some possible worlds it might denote something else, a 'representor' of me, for instance. If j is the individual concept associated with 'Jason', then 'Jason is a philosopher' is true at (and therefore is the set of) all and only the worlds w where $j(w)$ is a philosopher. From this perspective, the proposition-theoretic correlates of names are individual concepts: functions from worlds to individuals in those worlds.

Whatever the right view of *modal* semantics, it is pretty clear that names are not *model-theoretically* rigid. Sure, there are models where 'Jason' denotes me; but there are also models where it denotes one of my colleagues, or a fried

[7] Here's one such case. Let I be the space of standard models of arithmetic that also interpret one additional constant, a. Let $\mathscr{F}$ be an ultrafilter that includes the I-propositions I-expressing the Peano axioms plus $[\![\mathsf{a} > \underline{\mathsf{n}}]\!]_I$ for each natural number n. Any model that makes true the Peano axioms and $\mathsf{a} > \underline{\mathsf{n}}$ for each n will be a non-standard model, which means that any model making all these sentences true will be non-standard and so not in I.

egg, or an inaccessible cardinal. So when we look for something to stand to I-propositions as names stand to sentences, we will want to look towards the Carnapian picture for guidance.

And the guidance is pretty easy to come by. An individual constant is a function from worlds to things in those worlds; so we can let our model-theoretic analogue be a function from models to things in those models. For a space of models I, an ***I*-concept** is a function from models $\mathcal{M}$ in I to elements of $\mathcal{M}$'s domain. (Formally, functions g that meet the constraint $g(\mathcal{M}) \in D_{\mathcal{M}}$.) These will be our correlates of names, and we can use them to build domains for our models.

Notice that I-concepts are correlates of names in another way, too. For any name c in the language of I, there is a function that takes each model $\mathcal{M}$ in I to c's denotation on $\mathcal{M}$. This is the I-concept 'associated' with c. We can say that c ***I*-denotes** it, and write it as $[\![c]\!]_I$.

3.1.2 I-Satisfaction and Open Formulas

Consider the open formula

(†) x is a philosopher.

What proposition does it express? None, presumably. The variable 'x' is just hanging out there, free; without specifying just what it stands for, 'x is a philosopher' doesn't succeed in saying anything, and so can't even have a truth-value, much less express a proposition.

We can, however, think of variables as temporary names. Then we can ask what proposition (†) expresses *relative to* an assignment to the variables. If we think of names as rigid designators, then we assign an individual i to x and say that, relative to that assignment, (†) expresses the proposition that i is a philosopher — that is, the set of possible worlds in which i is a philosopher. If instead we think of names as individual concepts, then we assign an individual concept g to 'x' and say that, relative to that assignment, (†) expresses the set of possible worlds w where $g(w)$ is a philosopher.

Something similar goes for I-propositions. It doesn't really make sense to ask what I-proposition

(‡) $\mathsf{Philosopher}(\mathsf{x})$

I-expresses.[8] But now that we know what kind of things stand in for names,

[8] Actually, given our definitions the question technically makes perfect sense, and its answer is the set of all models in I which make the open sentence (‡) true — which is the set of models which make it true relative to every variable assignment over them. But this answer is an artefact of some arbitrary choices we made in our model theory, and furthermore isn't particularly useful. It makes $[\![\mathsf{Philosopher}(\mathsf{x})]\!]_I$ equivalent to $[\![\forall\mathsf{x}(\mathsf{Philosopher}(\mathsf{x}))]\!]_I$, and so doesn't recognize the 'openness' of the formula in question. We still need a notion of 'I-expression relative to a variable assignment' that would serve as the I-proposition-theoretic analogue to what we want to say about (†).

we know what it would be to assign, at the level of I-propositions, something to that variable. We want to assign it an I-concept.

Let an ***I*-assignment** be a function from variables to I-concepts. If A is an I-assignment, then $A(x)$ is an I-concept. Intuitively, it is a concept that assigns different things to the variable x on different models in I.

Once we have this, we can say what I-proposition a formula expresses relative to an I-assignment. Defining this explicitly takes one more piece of machinery, though.

Suppose you took an I-assignment A and then picked a model $\mathcal{M}$ from I. If you feed each variable into A and then apply the resulting I-concept to $\mathcal{M}$, you'll get an object in $\mathcal{M}$ each time. So the result of that process will be a variable assignment on $\mathcal{M}$. This means that every I-assignment induces, for each model, a variable assignment on that model; we call it the **$\mathcal{M}$-reduct of A**, formally defined by $A_{\mathcal{M}}(x) = A(x)(\mathcal{M})$.

We can use these reducts to now define the I-expression of an I-proposition relative to an I-assignment:

***I*-Expression:** $[\![\phi]\!]_{I,A} = \{\mathcal{M} \in I : \mathcal{M}, A_{\mathcal{M}} \vDash \phi\}$.

This then cashes out the intuitive idea of 'what I-proposition an open formula expresses when variables are assigned to I-concepts.'

We can use this apparatus to make sense of a further idea. If A is an I-assignment, then we use $A[x \triangleright g]$ for the I-assignment just like it except that x is assigned to g.[9] Intuitively, then, $[\![Fx]\!]_{I,A[x \triangleright g]}$ is the proposition that 'says' that the I-concept g is F.

Filters (and ultrafilters) are the propositional analogues of theories. If $[\![Fx]\!]_{I,A[x \triangleright g]}$ is in a filter $\mathscr{F}$, it's as though the 'theory' $\mathscr{F}$ 'says' that g is F. In this case, we say that g **$\mathscr{F}$-satisfies** the formula Fx. We can expand this terminology to formulas open in more variables, saying, for instance, that g and h $\mathscr{F}$-satisfy Rxy iff for every I-assignment A, $[\![Rxy]\!]_{I,A[x \triangleright g][y \triangleright h]} \in \mathscr{F}$. This expansion runs the risk of ambiguity, since for more complex formulas it will leave it open which I-concept goes with which variable. Still, we will use it when no confusion threatens.

3.2 *Building Models*

Henkin's method for proving completeness builds a model out of names from our theory and then 'reads off' the model's interpretation function from the theory's atomic sentences. For instance, if Fa is in the theory, we put the item built from a in the extension of F, and if it isn't, we don't.

But when we build the domain out of names, we don't *identify* the domain with the set of names. That would cause problems. For our theory might include '$c = d$'; but c and d are different names. On a model with the names c

[9]Note that $A[x \triangleright g](y)(\mathcal{M}) = A_{\mathcal{M}}[x \triangleright g(\mathcal{M})](y)$ for any $\mathcal{M} \in I$ and variable y.

and d in the domain, where c and d each denote themselves, $c = d$ will count as false, whereas we might want to make a model which counts it as true. So instead we make our model out of *equivalence classes* of names, where two names c and d are 'equivalent' when $c = d$ is in our theory.

The same thing goes for proving Łos's theorem. We are going to build a model from I-concepts in our space, and 'read off' the interpretation function from the atomic I-propositions in the filter. For instance, if an I-proposition g $\mathscr{F}$-satisfies Fx, we put the item built from g in the extension of F, and if it doesn't, we don't.

But for the same reason, we don't *identify* the domain with the set of I-interpretations. We might want our model to make $c = d$ true, even though $[\![c]\!]_I \neq [\![d]\!]_I$, for instance. If we just put these two items in the model and had c and d denote each of them respectively, we'd make that sentence false. So instead we make our model out of *equivalence classes* of I-concepts, where two I-concepts are 'equivalent' iff they $\mathscr{F}$-satisfy $x = y$.

More precisely, for any filter $\mathscr{F}$ over a space of models I, we say that I-concepts g and h are **$\mathscr{F}$-equivalent**, written $g \approx_{\mathscr{F}} h$, iff they $\mathscr{F}$-satisfy $x = y$. This turns out to be an equivalence relation, so we let $[g]_{\mathscr{F}}$ be the **$\mathscr{F}$-equivalence class** containing g, and build our model out of $\mathscr{F}$-equivalence-classes of I-concepts. As a notational matter, when the choice of $\mathscr{F}$ is obvious, we simply write $g \approx h$ and $[g]$, and call them equivalence classes.

Now that we have a domain, we finish building our model by reading off its extensions and denotations from the ultrafilter $\mathscr{F}$. For any constant c, we have it denote $[[\![c]\!]_I]$, and for any n-placed predicate Π, we put $\langle g_1, \ldots, g_n \rangle$ in its extension if and only if $g_1, \ldots, g_n$ $\mathscr{F}$-satisfy $\Pi x_1 \ldots x_n$.[10] When $\mathscr{F}$ is an ultrafilter, this model is our ultraproduct of $\mathscr{F}$ over I, $\mathcal{M}^I_{\mathscr{F}}$.

3.3 *Making Sentences True*

So now we have a model of a language $\mathcal{L}$, built from a space of models I and an ultrafilter $\mathscr{F}$. We want to show that a sentence of $\mathcal{L}$ is true on this model if and only if it I-expresses an I-proposition in $\mathscr{F}$. Before we can do that, there is one minor wrinkle still to iron out.

The wrinkle is this. What we want to prove is a universal claim about all sentences. The usual way of doing that is by induction on the length of formulas. But these inductions make us consider open formulas as well as closed ones, to deal with complex quantificational sentences. We usually deal with this by proving the stronger universal claim about all sentences, relativized to variable assignments. So we are going to need to think about variable assignments for our new models and how they are supposed to relate to the I-expression of I-propositions.

[10]This could go awry if, for instance, g $\mathscr{F}$-satisfied a predicate F, h did *not* $\mathscr{F}$-satisfy F, and $[g] = [h]$. Happily, we can rely on the definition of $\approx$ and the fact that $\mathscr{F}$ is a filter to show that this can't happen; the details are in the appendix.

We already have I-assignments, which pair variables with I-concepts. These are close to what we want, but aren't quite there yet. We need assign variables to *equivalence classes* of I-concepts, and the I-assignments only pair them with the I-concepts themselves. But we can easily build what we need out of I-assignments. If A is an I-assignment, let $[A]$ be the function where $[A](x) = [A(x)]$ for any variable x. Then if A is any I-assignment, $[A]$ is a variable assignment over the ultraproduct $\mathcal{M}^I_{\mathscr{F}}$. (Notice also that every variable assignment over $\mathcal{M}^I_{\mathscr{F}}$ is $[A]$ for *some* A, so this notation won't lead us to miss out on any variable assignments.

Now we can use a relatively straightforward induction to prove that ϕ's being true on the ultraproduct, relative to an assignment $[A]$, coincides precisely with $[\![\phi]\!]_{I,A}$'s being in the ultrafilter. The base case for the induction falls out pretty quickly from our definition of the ultraproduct; this is no surprise, since we built the ultraproduct precisely to give us this base case. The only other wrinkle is that, in the induction step, when we come to the case of quantifications, we have to use the axiom of choice to select a representative I-concept out of each equivalence class. But this shouldn't be too worrying; we already had to appeal to choice to extend our filter to an ultrafilter anyway, so a further appeal is no further commitment.

All this work, then, ends up establishing:

Theorem 2 (Łos's Theorem) *If $\mathscr{F}$ is an ultrafilter over a space I of models of a language $\mathcal{L}$, then there exists a model $\mathcal{M}^I_{\mathscr{F}}$ where, for any formula ϕ of $\mathcal{L}$,*

$$\mathcal{M}^I_{\mathscr{F}}, [A] \vDash \phi \textit{ iff } [\![\phi]\!]_{I,A} \in \mathscr{F}.$$

4 Compactness: An Application

If $\mathcal{L}$ is a first-order language, let a **theory** be any non-empty set of sentences of $\mathcal{L}$, and a **subtheory** be any non-empty subset of a theory. Most advanced logic textbooks at some point prove

Theorem 3 (Compactness) *If every finite subtheory of a theory Γ has a model, then Γ also has a model.*

This is an easy corollary of the completeness theorem, and is usually proved that way: having a model coincides (via completeness) with being proof-theoretically consistent, and since proof-theoretic consistency is essentially finite, the only way a theory could be inconsistent would be by having an inconsistent finite subtheory.

But that's a proof of a model-theoretic theorem that takes a detour through proof theory. We might wonder: Is there any way of proving compactness *directly*, just using model theoretic resources? The answer is yes, and the way uses ultrafilters.

Here's the idea. Suppose that every finite subtheory of Γ is consistent. That means that every finite subset of Γ has a model. We let I be the space of all these models. Then we create a filter over the finite subtheories in a way that gives every theory a 'representative' in the filter. We use the Central Theorem to generate an ultrafilter made out of subtheories, and then trade it in for an ultrafilter over the models of those subtheories. The representative for each sentence ϕ thus get traded in for I-propositions that I-express ϕ. Finally, Łos's theorem tells us that the ultraproduct makes all the I-propositions in the ultrafilter true; but every sentence in Γ I-expresses some I-proposition in the ultrafilter. So the ultraproduct is a model of Γ.

Let S be the set of all finite subtheories of Γ. (Hereafter, when I say 'subtheory' I will mean 'finite subtheory' unless I say otherwise.) Our ultrafilter will be made up of *sets* of these subtheories. For each subtheory Δ, we'll let its representative be the set of all subtheories that include it. That is, if Δ is a subtheory, then its representative Δ^* is the set $\{\Sigma \in S : \Delta \subseteq \Sigma\}$.

We want our filter to include all the representatives. But if it includes *only* the representatives, it might not be closed under implication, which would make it not a filter. Suppose, for instance, there are infinitely many non-overlapping subtheories. Including each of their representatives won't give us any set of theories containing all of them; but since such a set is a superset of each of them, we'll need it to be included to make the result a filter. We fix this by defining our filter to include all the representatives plus all their supersets. More precisely, we define

$$\mathscr{F}^- = \{T \subseteq S : \text{ for some } \Delta \in S, \Delta^* \subseteq T\}.$$

This gets to be closed under implication by design; and it doesn't take much to show that it is non-trivial and closed under conjunction. So it's a filter. Thus, by the Central Theorem, it can be expanded to an ultrafilter $\mathscr{F}^S$.

Now we simply trade each theory in S for the model in I that makes it true, which trades the ultrafilter $\mathscr{F}^S$ over S for another one, $\mathscr{F}$, over I. By Łos's theorem, the ultraproduct of $\mathscr{F}$ over I makes ϕ true if and only if $[\![\phi]\!]_I \in \mathscr{F}$. So to finish up, we just have to show that, if $\phi \in \Gamma$, then $[\![\phi]\!]_I \in \mathscr{F}$.

But if $\phi \in \Gamma$, then $\{\phi\}$ is a finite subtheory of Γ, and so has a representative $\{\phi\}^*$ in $\mathscr{F}^S$. Every theory in $\{\phi\}^*$ includes ϕ. When we trade S in for I, we also trade in $\{\phi\}^*$ for a set T of models, each one of which will make ϕ true. But this means that $T \subseteq [\![\phi]\!]_I$. Since $T \in \mathscr{F}$, which is closed under implication, $[\![\phi]\!]_I \in \mathscr{F}$ as well. Compactness follows.

Proofs

A.1 *Preliminary Results*

Proposition 4 *A set $\mathscr{F}$ is a filter iff it is a filter*.*

Proof *Left-to-right.* Suppose $\mathscr{F}$ is a filter over I. We need to show it is conjunction-complete and non-universal. The latter is simple, because as a filter $\mathscr{F}$ is non-trivial, so $\emptyset \notin \mathscr{F}$. For the former, suppose $X_1, X_2 \subseteq I$. If $X_1, X_2 \in \mathscr{F}$, then $X_1 \cap X_2 \in \mathscr{F}$ by conjunction closure. Conversely, if $X_1 \cap X_2 \in \mathscr{F}$, since $X_1 \cap X_2 \subseteq X_1$ (and similarly for X_2), by implication closure each are in $\mathscr{F}$.. So $\mathscr{F}$ is conjunction-complete.

Right-to-left. Suppose $\mathscr{F}$ is a filter* over I. We need to show it is non-trivial and closed under implication and conjunction. Conjunction closure is easy, since it is one direction of conjunction-completeness. For implication closure, note that, if $S \in \mathscr{F}$ and $S \subseteq T \subseteq I$, then $S = S \cap T$. So by conjunction completeness, $T \in \mathscr{F}$. Thus $\mathscr{F}$ is closed under implication. Since it is closed under implication, if $\emptyset \in \mathscr{F}$, then $\mathscr{F}$ would contain all subsets of I, contrary to non-universality. QED.

Lemma 5 *Suppose $\mathscr{F}$ is a filter over I and $X \subseteq I$. Then*

(i) $\mathscr{F} \subseteq \mathscr{F} + X$ *and* $X \in \mathscr{F} + X$.

(ii) $\mathscr{F} + X$ *is closed under implication and conjunction.*

(iii) $\mathscr{F} + X$ *is a filter iff* $-X \notin \mathscr{F}$.

Proof Recall that $\mathscr{F} + X = \{Y \subseteq I : \text{for some } A \in \mathscr{F}, A \cap X \subseteq Y\}$.

(i). Follows since $I \cap X = X$ and $X \cap Y \subseteq Y$.

(ii). For implication, suppose that $Y \in \mathscr{F} + X$ and $Y \subseteq Z$. Since $Y \in \mathscr{F} + X$, for some $A \in \mathscr{F}$, $A \cap X \subseteq Y$. But then $A \cap X \subseteq Z$, so $Z \in \mathscr{F} + X$.

For conjunction, suppose $Y_1, Y_2 \in \mathscr{F} + X$. Then for some $A_1, A_2 \in \mathscr{F}$, $A_1 \cap X \subseteq Y_1$ and $A_2 \cap X \subseteq Y_2$. In this case, $(A_1 \cap A_2) \cap X = (A_1 \cap X) \cap (A_2 \cap X) \subseteq Y_1 \cap Y_2$. But since $A_1, A_2 \in \mathscr{F}$ which is a filter, $A_1 \cap A_2 \in \mathscr{F}$, so $Y_1 \cap Y_2 \in \mathscr{F} + X$.

(ii). By part (ii), $\mathscr{F} + X$ is a filter iff $\emptyset \notin \mathscr{F} + X$. We need to show that $\emptyset \in \mathscr{F} + X$ iff $-X \in \mathscr{F}$.

Suppose first that $-X \in \mathscr{F}$. Then $-X \in \mathscr{F} + X$ and $X \in \mathscr{F} + X$ by part (i), so since $\mathscr{F} + X$ is closed under conjunction from part (ii), $\emptyset = X \cap -X \in \mathscr{F}$.

Suppose next that $\emptyset \in \mathscr{F} + X$. Then there is an $A \in \mathscr{F}$ such that $A \cap X = \emptyset$. Thus A and X have no members in common, so $A \subseteq -X$. Since $\mathscr{F}$ is closed under implication, $-X \in \mathscr{F}$. QED.

Lemma 6 *A filter is maximal iff it is prime iff it is ultra.*

Proof Let $\mathscr{F}$ be a filter.

Suppose $\mathscr{F}$ is maximal. It is clear that if either $P \in \mathscr{F}$ or $Q \in \mathscr{F}$, then $P \cup Q \in \mathscr{F}$ by closure under implication. For the other direction, suppose $P \cup Q \in \mathscr{F}$, and suppose for reductio that neither P nor Q is in $\mathscr{F}$. Now consider $\mathscr{F} + P$. If this is a filter, then $\mathscr{F}$ is not maximal; since it is not a filter, $-P \in \mathscr{F}$ by Lemma 5.iii. Analogous reasoning shows $-P \in \mathscr{F}$. But then $-P \cap -Q \in \mathscr{F}$, and so $(-P \cap -Q) \cap (P \cup Q) = \emptyset \in \mathscr{F}$, contrary to its being a filter. Reductio complete; either $P \in \mathscr{F}$ or $Q \in \mathscr{F}$. So $\mathscr{F}$ is prime.

Suppose $\mathscr{F}$ is prime. Let $P \subseteq I$. Then $P \cup -P = I \in \mathscr{F}$; since it is prime, either $P \in \mathscr{F}$ or $-P \in \mathscr{F}$. So it is ultra.

Suppose $\mathscr{F}$ is not maximal. Then there is a filter $\mathscr{F}'$ which extends $\mathscr{F}$. Let $P \in \mathscr{F}'$ where $P \notin \mathscr{F}$. If $-P \in \mathscr{F}$, then $-P \in \mathscr{F}'$, so $P \cap -P = \emptyset \in \mathscr{F}'$, contrary to its being a filter. So $-P \notin \mathscr{F}'$, which means $\mathscr{F}$ is not ultra. Contrposing, if $\mathscr{F}$ is ultra, it is maximal.

A.2 The Central Theorem on Ultrafilters

Now to prove that every filter can be extended to an ultrafilter. The proof uses transfinite induction and relies on a consequence of the well-ordering theorem, which is equivalent to the Axiom of Choice.

Proof Let $\mathscr{F}$ be a filter over I. By the well-ordering theorem, each subset of I can be indexed by an ordinal. Call this ordinal γ. We will write these subsets, thus indexed, as P_α. Note that for each P_α, $\alpha < \gamma$.

For $\alpha \leq \gamma$, define a (transfinite) chain of sets as follows:

$$\mathscr{F}_0 = \mathscr{F}$$
$$\mathscr{F}_{\alpha+1} = \begin{cases} \mathscr{F}_\alpha + P_\alpha \text{ if that is a filter, and} \\ \mathscr{F}_\alpha \text{ otherwise.} \end{cases}$$
$$\mathscr{F}_\alpha = \bigcup_{\beta<\alpha} \mathscr{F}_\beta \text{ for limit } \alpha$$

We will show, by transfinite induction, that the $\mathscr{F}_\alpha$'s are a chain ordered by subsethood and that each $\mathscr{F}_\alpha$ is a filter, from which it follows that the final one, $\mathscr{F}_\gamma$ is a filter.

The base case holds by assumption. For successor ordinals, $\mathscr{F}_\alpha \subseteq \mathscr{F}_{\alpha+1}$ by theorem 5.i, and $\mathscr{F}_\alpha$ is a filter by construction. So we need to consider limit ordinal α. Assume that for each $\beta < \alpha$, (i) $\mathscr{F}_\beta$ is a filter and (ii) if $\beta' < \beta$, $\mathscr{F}_{\beta'} \subseteq \mathscr{F}_\beta$. Let $\mathscr{F}' = \bigcup_{\beta<\alpha} \mathscr{F}_\beta$. It is clear that if $\beta < \alpha$, $\mathscr{F}_\beta \subseteq \mathscr{F}'$. So we just need to show that $\mathscr{F}'$ is a filter. To do that we have to show that is closed under implication and intersection and is non-trivial. For non-triviality, it suffices to note that, if $\emptyset \in \mathscr{F}'$, then for some $\beta < \alpha$, $\emptyset \in \mathscr{F}_\beta$, contradicting $\mathscr{F}_\beta$'s being a filter.

For implication closure, suppose $X \in \mathscr{F}'$ and $X \subseteq Y \subseteq I$. Then for some $\beta < \alpha$, $X \in \mathscr{F}_\beta$; since $\mathscr{F}_\beta$ is a filter it is closed under implication, so $Y \in \mathscr{F}_\beta \subseteq \mathscr{F}'$.

For conjunction closure, suppose $X_1, X_2 \in \mathscr{F}'$. Then for some β_1, β_2, $X_1 \in \mathscr{F}_{\beta_1}$ and $X_2 \in \mathscr{F}_{\beta_2}$. Since β_1 and β_2 are ordinals, either $\beta_1 \leq \beta_2$ or vice versa, which means either $\mathscr{F}_{\beta_1} \subseteq \mathscr{F}_{\beta_2}$ or vice versa. Suppose it's the first (without loss of generality). Then $X_1 \in \mathscr{F}_{\beta_2}$; since it is a filter with X_2 in it as well, $X_1 \cap X_2 \in \mathscr{F}_{\beta_2} \subseteq \mathscr{F}'$.

Thus every $\mathscr{F}_\alpha$ for $\alpha \leq \gamma$ is a filter. Now we will show that $\mathscr{F}_\gamma$ is an ultrafilter. Let $P \subseteq I$; then $P = P_\alpha$ for some $\alpha < \gamma$. Consider $\mathscr{F}_\alpha + P_\alpha$. If that is a filter, then it is $\mathscr{F}_{\alpha+1} \subseteq \mathscr{F}_\gamma$, in which case $P \in \mathscr{F}_\gamma$. If it is not a filter, then by theorem 5.iii, $-P_\alpha \in \mathscr{F}_\alpha \subseteq \mathscr{F}_\gamma$. So either $P \in \mathscr{F}_\gamma$ or $-P \in \mathscr{F}_\gamma$; since P was arbitrary, $\mathscr{F}_\gamma$ is an ultrafilter. QED.

A.3 Łos's theorem

In this section, to clean up a little bit on clutter, we'll eliminate the '$\mathscr{F}$ subscript on the notation for equivalence relations and equivalence classes. (There's still a fair bit of notational clutter. Sorry about that!)

A.3.1 Ensuring Well-Definedness for 'Ultrafilter"

We want to show that 'ultraproduct' is well-defined. The worry is that we might have (for instance) $[g] = [h]$, but $[\![\Pi x]\!]_{I,A[x \triangleright g]} \in \mathscr{F}$ while $[\![\Pi x]\!]_{I,A[x \triangleright h]} \notin \mathscr{F}$, which would make our attempted definition incoherent. We want to show that this can't happen. We'll use another bit of shorthand for the result and its proof, to ease on clutter: when we have a list '$x_1 \dots x_n$', '$g_1, \dots, g_n$', etc., we will simply write it as $\boldsymbol{x}$, $\boldsymbol{g}$, etc. Likewise, we will abbreviate $[x_1 \triangleright g_1] \dots [x_n \triangleright g_n]$ as $[\boldsymbol{x} \triangleright \boldsymbol{g}]$.

With this in hand, we have:

Proposition 7 *If each $[g_i] = [h_i]$ for $1 \leq i \leq n$, then*

$$[\![\Pi \boldsymbol{x}]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}]} \in \mathscr{F} \textit{ iff } [\![\Pi \boldsymbol{x}]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{h}]} \in \mathscr{F}.$$

Proof Begin by noting that $[\![\Pi \boldsymbol{x}]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}]} = [\![\Pi \boldsymbol{y}]\!]_{I,A[\boldsymbol{y} \triangleright \boldsymbol{g}]}$. Note also that, if some variables $\boldsymbol{z}$ don't show up in ϕ, then $[\![\phi]\!]_{I,B} = [\![\phi]\!]_{I,B[\boldsymbol{z} \triangleright \boldsymbol{j}]}$ for any I-concepts $\boldsymbol{j}$ and I-assignment B. So what we will actually show is

$$[\![\Pi \boldsymbol{x}]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}][\boldsymbol{y} \triangleright \boldsymbol{h}]} \in \mathscr{F} \text{ iff } [\![\Pi \boldsymbol{y}]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}][\boldsymbol{y} \triangleright \boldsymbol{h}]} \in \mathscr{F}.$$

First, suppose that for $1 \leq i \leq n$, $[g_i] = [h_i]$. Thus each $g_1 \approx h_1$ and so $[\![x = y]\!]_{I,A[x \triangleright g_i][y \triangleright h_i]} \in \mathscr{F}$. By swapping variables, $[\![x_i = y_i]\!]_{I,A[x_i \triangleright g_i][y_i \triangleright h_i]} = [\![x_i = y_i]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}][\boldsymbol{y} \triangleright \boldsymbol{h}]} \in \mathscr{F}$. Now suppose that $[\![\Pi \boldsymbol{x}]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}][\boldsymbol{y} \triangleright \boldsymbol{h}]} \in \mathscr{F}$. By closure under conjunction for $\mathscr{F}$ we reason

$$[\![\Pi \boldsymbol{x}]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}][\boldsymbol{y} \triangleright \boldsymbol{h}]} \cap [\![x_1 = y_1]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}][\boldsymbol{y} \triangleright \boldsymbol{h}]} \cap \dots \cap [\![x_n = y_n]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}][\boldsymbol{y} \triangleright \boldsymbol{h}]} \in \mathscr{F}, \qquad \text{so}$$
$$[\![\Pi \boldsymbol{x} \wedge x_1 = y_1 \wedge \dots \wedge x_n = y_n]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}][\boldsymbol{y} \triangleright \boldsymbol{h}]} \in \mathscr{F}.$$

But for every variable assignment a and model $\mathcal{M}$, if $\mathcal{M}, a \vDash \Pi \boldsymbol{x} \wedge x_1 = y_1 \wedge \dots \wedge x_n = y_n$, then $\mathcal{M}, a \vDash \Pi \boldsymbol{y}$. So

$$[\![\Pi \boldsymbol{x} \wedge x_1 = y_1 \wedge \dots \wedge x_n = y_n]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}][\boldsymbol{y} \triangleright \boldsymbol{h}]} \subseteq [\![\Pi \boldsymbol{y}]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}][\boldsymbol{y} \triangleright \boldsymbol{h}]}.$$

By closure under implication, $[\![\Pi \boldsymbol{y}]\!]_{I,A[\boldsymbol{x} \triangleright \boldsymbol{g}][\boldsymbol{y} \triangleright \boldsymbol{h}]} \in \mathscr{F}$. Precisely the same argument, in the other direction, finishes the biconditional. QED.

A.3.2 The Theorem Itself

We need to show, recall,

$$\mathcal{M}^I_{\mathscr{F}}, [A] \vDash \phi \text{ iff } [\![\phi]\!]_{I,A} \in \mathscr{F}.$$

We'll start with a lemma about denotations.

Lemma 8 *For any term α of $\mathcal{L}$, $\alpha^{\mathcal{M}^I_{\mathscr{F}},[A]} = [[\![\alpha]\!]_{I,A}]$.*

Proof Suppose α is a variable; then $[\![\alpha]\!]_{I,A} = A(\alpha)$. Then $\alpha^{\mathcal{M}^I_{\mathscr{F}},[A]_{\mathscr{F}}} = [A]_{\mathscr{F}}(\alpha) = [A(\alpha)] = [[\![\alpha]\!]_{I,A}]$. If instead α is a term, then $[[\![\alpha]\!]_{I,A}]$ = (by construction) $\mathcal{M}^I_{\mathscr{F}}(\alpha) = \alpha^{\mathcal{M}^I_{\mathscr{F}},[A]_{\mathscr{F}}}$. QED.

Now we prove Łos's Theorem itself. Let $\mathscr{F}$ be an ultrafilter over a space of models I of a language $\mathcal{L}$, and $\mathcal{M}^I_{\mathscr{F}}$ its ultraproduct. The proof is by induction on the length of ϕ.

Proof In the base case, ϕ is either $\Pi\alpha_1 \ldots \alpha_n$ or $\alpha = \beta$.

For the first,

$$\begin{aligned} [\![\Pi\alpha_1 \ldots \alpha_n]\!]_{I,A} \in \mathscr{F} &\text{ iff } \langle [[\![\alpha_1]\!]_{I,A}], \ldots, [[\![\alpha_1]\!]_{I,A}] \rangle \in I_{\mathcal{M}^I_{\mathscr{F}}}(\Pi) && \text{(df. ultraproduct)} \\ &\text{ iff } \langle \alpha_1^{\mathcal{M}^I_{\mathscr{F}},[A]}, \ldots, \alpha_n^{\mathcal{M}^I_{\mathscr{F}},[A]} \rangle \in I_{\mathcal{M}^I_{\mathscr{F}}}(\Pi) && \text{(lemma 8)} \\ &\text{ iff } \mathcal{M}^I_{\mathscr{F}}, [A] \vDash \Pi\alpha_1 \ldots \alpha_n. \end{aligned}$$

The argument is essentially the same when ϕ is $\alpha = \beta$.

For the induction step, ϕ is either of the form $\sim\psi$, $\psi \wedge \chi$, or $\exists x \chi$.

When ϕ is $\sim\psi$, then $[\![\sim\psi]\!]_{I,A} \in \mathscr{F}$ iff $-[\![\psi]\!]_{I,A} \in \mathscr{F}$ iff (since $\mathscr{F}$ is an ultrafilter) $[\![\psi]\!]_{I,A} \notin \mathscr{F}$ iff (by the induction hypothesis) if $\mathcal{M}^I_{\mathscr{F}}, [A] \nvDash \psi$ iff $\mathcal{M}^I_{\mathscr{F}}, [A] \vDash \sim\psi$.

When ϕ is $\psi \wedge \psi$, we rely on the fact that filters are also filters*. Then $[\![\psi \wedge \chi]\!]_{I,A} \in \mathscr{F}$ iff $[\![\psi]\!]_{I,A} \cap [\![\psi]\!]_{I,A} \in \mathscr{F}$ iff (since $\mathscr{F}$ is a filter*) $[\![\psi]\!]_{I,A} \in \mathscr{F}$ and $[\![\chi]\!]_{I,A} \in \mathscr{F}$ iff (by the induction hypothesis) $\mathcal{M}^I_{\mathscr{F}}, [A] \vDash \psi$ and $\mathcal{M} \vDash \chi$ iff $\mathcal{M}^I_{\mathscr{F}}, [A] \vDash \phi \wedge \chi$.

When ϕ is $\exists x \psi$, the general form of reasoning is:

$$\begin{aligned} \mathcal{M}^I_{\mathscr{F}}, [A] \vDash \exists x \psi &\text{ iff for some } I\text{-concept } g,\ \mathcal{M}^I_{\mathscr{F}}, [A][x \triangleright [g]] \vDash \psi \\ &\text{ iff for some } I\text{-concept } g,\ [\![\psi]\!]_{I,A[x \triangleright g]} \in \mathscr{F} && \text{(induction hyp.)} \\ &\text{ iff } [\![\exists x \psi]\!]_{I,A} \in \mathscr{F}. \end{aligned}$$

To finish the proof we have to verify the equivalence of the last two lines.

First, suppose that a model $\mathcal{M}$ is in $[\![\psi]\!]_{I,A[x \triangleright [g]]}$. Let a be the assignment where $a(x) = A_{\mathcal{M}}(x)$. Then this means that $\mathcal{M}, a[x \triangleright [A_{\mathcal{M}}(g)] \vDash \psi$, in which case $\mathcal{M}, a \vDash \exists x \psi$, so $\mathcal{M} \in [\![\exists x \psi]\!]_{I,A}$. So $[\![\psi]\!]_{I,A[x \triangleright g]} \subseteq [\![\exists x \psi]\!]_{I,A}$. Since $\mathscr{F}$ is closed under implication, this means that if $[\![\psi]\!]_{I,A[x \triangleright [g]]} \in \mathscr{F}$, then $[\![\exists x \psi]\!]_{I,A} \in \mathscr{F}$.

For the other direction, suppose that $[\![\exists x \psi]\!]_{I,A} \in \mathscr{F}$. Then for each $\mathcal{M} \in [\![\exists x \psi]\!]_{I,A}$, $\mathcal{M}, a_{\mathcal{M}} \vDash \exists x \psi$, which means that in each domain $D_{\mathcal{M}}$ there is an object $o_{\mathcal{M}}$ where $\mathcal{M}, a_{\mathcal{M}}[x \triangleright o_{\mathcal{M}}] \vDash \psi$. Let g be a function that takes each $\mathcal{M}$ to such an $o_{\mathcal{M}}$. (This exists by the axiom of choice.) So for every model $\mathcal{M}$ in $[\![\exists x \psi]\!]_{I,A}$, $\mathcal{M}, A[x \triangleright g]_{\mathcal{M}} \vDash \psi$. Thus $[\![\exists x \psi]\!]_{I,A} \subseteq [\![\psi]\!]_{I,A[x \triangleright g]}$. Since $\mathscr{F}$ is closed under implication, $[\![\psi]\!]_{I,A[x \triangleright g]} \in \mathscr{F}$. Thus for some I-concept g, $[\![\psi]\!]_{I,A[x \triangleright g]} \in \mathscr{F}$. QED.

References

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

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.

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.