Journal of Universal Computer Science, vol. 11, no. 12 (2005), 1945-1962
submitted: 9/6/05, accepted: 5/10/05, appeared: 28/12/05 © J.UCS

Axiomatic Classes of Intuitionistic Models[1]

Robert Goldblatt
(Victoria University of Wellington, New Zealand
rob.goldblatt@vuw.ac.nz)

Abstract: A class of Kripke models for intuitionistic propositional logic is 'axiomatic' if it is the class of all models of some set of formulas (axioms). This paper discusses various structural characterisations of axiomatic classes in terms of closure under certain constructions, including images of bisimulations, disjoint unions, ultrapowers and 'prime extensions'. The prime extension of a model is a new model whose points are the prime filters of the lattice of upwardly-closed subsets of the original model. We also construct and analyse a 'definable' extension whose points are prime filters of definable sets.

A structural explanation is given of why a class that is closed under images of bisimulations and invariant under prime/definable extensions must be invariant under arbitrary ultrapowers. This uses iterated ultrapowers and saturation.
Key Words: intuitionistic logic, Kripke model, bisimulation, disjoint union, prime filter, ultraproduct, iterated ultrapower, saturated model
Category: F.4.1

1 Introduction

This is a contribution to the model theory of intuitionistic logic, the logic that underlies a good deal of the mathematical research of Douglas Bridges. The paper is written in honour of his 60th birthday.

Our interest is in various structural characterisations of classes of Kripke models for intuitionistic propositional logic (IPC) that are *axiomatic*, which means being the class of all models of some set of formulas (axioms). It was shown in [Rodenburg, 1986, 13.8] that a class of IPC-models is axiomatic if, and only if, it is

> *closed under images of total bisimulation relations, inner submodels, disjoint unions, ultrapowers and ultraroots*

(these notions will be explained later). An analogous theorem for models of Boolean modal propositional logic was given in [Venema, 1999]: a class of modal Kripke models is the class of all models of a set of modal formulas if, and only if, it is

> *closed under images of bisimulation relations, disjoint unions and ultrafilter extensions, while its complement is closed under ultrafilter extensions.*

[1] C. S. Calude, H. Ishihara (eds.). *Constructivity, Computability, and Logic. A Collection of Papers in Honour of the 60th Birthday of Douglas Bridges.*

Here the ultrafilter extension of a model $\mathcal{M}$ is a new model whose points are all the ultrafilters on the underlying set of $\mathcal{M}$.

We will show that a characterisation of this second kind is available for IPC-axiomatic classes if ultrafilter extensions are replaced by *prime* extensions whose points are the prime filters in the Heyting algebra of upwardly-closed subsets of an IPC-model. An equivalent characterisation results if we replace the prime extensions by a notion of *definable* extension, restricting the construction to the Heyting algebra of definable subsets of the model. The main aim of the paper is to explore the structural relationships between prime/definable extensions and ultrapowers, showing how they are connected by bisimulations, and how the various types of characterisation come to be equivalent.

An IPC-model is also a model for a certain first-order language $\mathcal{L}$, and IPC-formulas translate into $\mathcal{L}$-formulas with a single free variable. In this way the model theory of IPC can be identified with that of a fragment of the Boolean logic of $\mathcal{L}$. We take advantage of the fact that $\mathcal{L}$ is countable, applying a standard fact about the *saturation* of ultrapowers for countable languages, namely that an ultrapower modulo a countably incomplete ultrafilter is $\aleph_1$-saturated. This is used in an iterated ultrapower construction to give a structural explanation of why a class that is closed under images of bisimulations and invariant under prime/definable extensions must be invariant under ultrapowers.

Here is a summary of the paper. In the next Section we describe the language and semantics of IPC. Sections 3 and 4 review the basic theory of truth preserving model-constructions, including bisimulations, bounded morphisms, inner submodels and disjoint unions. Section 5 explains why the relation of logical equivalence is a bisimulation between sufficiently saturated models. Section 6 and 7 define the prime and definable extensions of a model and gives their basic properties and relationships. Section 8 is about ultraproducts and ultrapowers and gives the proof that a class of models is invariant under ultrapowers if it is closed under bisimulation images and invariant under definable extensions. The final Section 9 gives the main result setting out a number of equivalent structural characterisations of axiomatic classes.

2 Languages and Models

Formulas of IPC are constructed from an infinite set $\{p_n : n \in \omega\}$ of propositional variables and the constant $\bot$ by the connectives $\wedge$, $\vee$ and $\to$. The negation of formula φ can be defined to be the formula $\varphi \to \bot$, and $\varphi \leftrightarrow \psi$ is an abbreviation for $(\varphi \to \psi) \wedge (\psi \to \varphi)$. We denote the set of all IPC-formulas by Φ.

A *quasiorder* on a non-empty set X is a reflexive and transitive relation $\leq$. A subset Y of X is *up-closed* if $y \in Y$ whenever $x \in Y$ and $x \leq y$. If $[x) = \{y \in X : x \leq y\}$, then $[x)$ is the smallest up-closed set containing x. In general, Y is up-closed iff $[x) \subseteq Y$ for all $x \in Y$. The set of all up-closed subsets of $(X, \leq)$ will be denoted $U(\leq)$. An *IPC-model* is a structure

$$\mathcal{M} = (X, \leq, P_0, \dots, P_n, \dots),$$

where $\leq$ is a quasiorder on X and each P_n is a member of $U(\leq)$. The satisfaction relation $\mathcal{M}, x \models \varphi$, expressing "formula φ is true/satisfied at x in $\mathcal{M}$", is defined by induction on the formation of the formula $\varphi \in \Phi$ as follows:

$\mathcal{M}, x \models p_n$ iff $x \in P_n$;
$\mathcal{M}, x \not\models \bot$;
$\mathcal{M}, x \models \varphi \wedge \psi$ iff $\mathcal{M}, x \models \varphi$ and $\mathcal{M}, x \models \psi$;
$\mathcal{M}, x \models \varphi \vee \psi$ iff $\mathcal{M}, x \models \varphi$ or $\mathcal{M}, x \models \psi$;
$\mathcal{M}, x \models \varphi \to \psi$ iff for all $y \geq x$, if $\mathcal{M}, y \models \varphi$ then $\mathcal{M}, y \models \psi$.

The collection $U(\leq)$ of up-closed sets forms a Heyting algebra under the partial order $\subseteq$, with lattice meet and join being the set-theoretic operations $\cap$ and $\cup$, and with least element $\emptyset$, greatest element X, and relative pseudo-complement operation $\Rightarrow$ defined by

$$Y \Rightarrow Z = \{x \in X : [x) \cap Y \subseteq Z\}$$

(see [Rasiowa and Sikorski, 1963] or [Balbes and Dwinger, 1974] for the general theory of Heyting algebras).

The "truth set" $\mathcal{M}(\varphi) := \{x \in X : \mathcal{M}, x \models \varphi\}$ of any formula turns out to be up-closed, and indeed the satisfaction conditions are equivalent to the equations

$\mathcal{M}(p_n) = P_n$;
$\mathcal{M}(\bot) = \emptyset$;
$\mathcal{M}(\varphi \wedge \psi) = \mathcal{M}(\varphi) \cap \mathcal{M}(\psi)$;
$\mathcal{M}(\varphi \vee \psi) = \mathcal{M}(\varphi) \cup \mathcal{M}(\psi)$;
$\mathcal{M}(\varphi \to \psi) = \mathcal{M}(\varphi) \Rightarrow \mathcal{M}(\psi)$.

It follows that $U(\mathcal{M}) = \{\mathcal{M}(\varphi) : \varphi \in \Phi\}$ is a sub-Heyting algebra of $U(\leq)$.

Formula φ is *true in the model* $\mathcal{M}$, written $\mathcal{M} \models \varphi$, if $\mathcal{M}, x \models \varphi$ for all $x \in X$, i.e. if $\mathcal{M}(\varphi) = X$. In this case we also say that $\mathcal{M}$ *is a model of* φ. For a set $\Sigma \subseteq \Phi$ we put $\mathcal{M} \models \Sigma$ if $\mathcal{M} \models \varphi$ for all $\varphi \in \Sigma$, and write $\mathrm{Mod}\,\Sigma$ for the class $\{\mathcal{M} : \mathcal{M} \models \Sigma\}$ of all *models of* Σ. A class $\mathbb{C}$ of IPC-models is called *axiomatic* if there exists some set Σ of formulas such that $\mathbb{C} = \mathrm{Mod}\,\Sigma$. The formulas that are true in all IPC-models are precisely those that are theorems of Heyting's intuitionistic propositional calculus. This model theory is due to [Kripke, 1965].

An IPC-model can be viewed as a structure for the first-order language $\mathcal{L}$ having a binary relation symbol $\leq$ interpreted as the quasi-order and unary relation symbols π_n interpreted as the sets P_n. As such, each IPC-model satisfies the $\mathcal{L}$-sentence σ_{qo} expressing that $\leq$ is a quasiorder. Each $\varphi \in \Phi$ can be translated into an $\mathcal{L}$-formula $\varphi^t(v)$ with a single free variable v, as follows:

$p_n^t = \pi_n(v)$; $\bot^t = \bot$;

$(\varphi \wedge \psi)^t(v) = \varphi^t(v) \wedge \psi^t(v)$; $\quad (\varphi \vee \psi)^t(v) = \varphi^t(v) \vee \psi^t(v)$;
$(\varphi \to \psi)^t(v) = \forall w(v \leq w \to (\varphi^t(w/v) \to \psi^t(w/v))$, where $w \neq v$ and v is free for w in $\varphi^t(v)$.

Then in general,

$$\mathcal{M}, x \models \varphi \quad \text{iff} \quad \mathcal{M} \models \varphi^t[x],$$

where the notation '$\mathcal{M} \models \varphi^t[x]$' means that φ^t is satisfied in the $\mathcal{L}$-structure $\mathcal{M}$ in the usual Tarskian sense for first-order logic when the variable v is assigned the value x. In this way IPC can be viewed as a special fragment of first-order logic. In particular,

$$\mathcal{M} \models \varphi \quad \text{iff} \quad \mathcal{M} \models \forall v \varphi^t,$$

so for any $\Sigma \subseteq \Phi$, an arbitrary $\mathcal{L}$-structure $\mathcal{M}$ belongs to $\mathrm{Mod}\,\Sigma$ iff

$$\mathcal{M} \models \{\sigma_{\mathrm{qo}}\} \cup \{\forall v \varphi^t : \varphi \in \Sigma\}.$$

Thus any axiomatic class is also an *elementary class*, i.e. the class of all $\mathcal{L}$-models of a set of $\mathcal{L}$-sentences. Our aim is to clarify just which elementary classes are of the form $\mathrm{Mod}\,\Sigma$.

3 Bisimulations

A *bisimulation* from IPC-model $\mathcal{M}$ to IPC-model $\mathcal{M}' = (X' \leq', P'_n)_{n \in \omega}$ is a binary relation $R \subseteq X \times X'$ such that for all $x \in X$ and $x' \in X'$, if xRx' then:

B1: $x \in P_n \quad \text{iff} \quad x' \in P'_n$.

B2: $x' \leq' y' \quad \text{implies} \quad \exists y(x \leq y \text{ and } yRy')$.

B3: $x \leq y \quad \text{implies} \quad \exists y'(x' \leq y' \text{ and } yRy')$.

When this holds, it follows that for all $\varphi \in \Phi$,

$$xRx' \quad \text{implies} \quad [\mathcal{M}, x \models \varphi \text{ iff } \mathcal{M}', x' \models \varphi]. \tag{3.1}$$

This is readily shown by induction on the formation of φ, with B1 taking care of the case that $\varphi = p_n$, and the 'back-and-forth' conditions B2 and B3 used for the inductive case that φ has the form $\varphi_1 \to \varphi_2$.

A bisimulation is *surjective* if its image $\{x' \in X' : \exists x(xRx')\}$ is X' itself. $\mathcal{M}'$ is a *bisimulation image* of $\mathcal{M}$ if there exists a surjective bisimulation from $\mathcal{M}$ to $\mathcal{M}'$. In that case, it follows from the above that $\mathcal{M} \models \varphi$ implies $\mathcal{M}' \models \varphi$. Thus an axiomatic class $\mathrm{Mod}\,\Sigma$ is *closed under bismulation images*: if it contains $\mathcal{M}$ then it contains all bisimulation images of $\mathcal{M}$.

Dually, a bisimulation is *total* if its domain $\{x \in X : \exists x'(xRx')\}$ is X. Then $\mathcal{M}' \models \varphi$ implies $\mathcal{M} \models \varphi$, so an axiomatic class is closed under domains of total bisimulations: if it contains $\mathcal{M}'$ then it contains any $\mathcal{M}$ having a total bisimulation to $\mathcal{M}'$. Alternatively, this can be seen from the fact that the definition of a bisimulation is symmetric, in the sense that if R is a bisimulation from $\mathcal{M}$ to $\mathcal{M}'$, then its inverse R^{-1} is a bisimulation from $\mathcal{M}'$ to $\mathcal{M}$. Moreover, R is total iff R^{-1} is surjective (and vice versa). Thus closure of any class $\mathbb{C}$ of models under bisimulation images implies closure under domains of total bisimulations. Hence it implies *invariance* under total surjective bisimulations, in the sense that if there exists a total surjective bisimulation from $\mathcal{M}$ to $\mathcal{M}'$, then $\mathcal{M} \in \mathbb{C}$ iff $\mathcal{M}' \in \mathbb{C}$.

A *bounded morphism* $f : \mathcal{M} \to \mathcal{M}'$ can be defined as a function $f : X \to X'$ whose graph $\{(x, f(x)) : x \in X\}$ is a bisimulation, and hence a *total* bisimulation. This is equivalent to the more common definition that $x \in P_n$ iff $f(x) \in P'_n$, and

$$f(x) \leq' y' \quad \text{iff} \quad \exists y(x \leq y \text{ and } f(y) = y'). \tag{3.2}$$

If f is surjective, then it is called a bounded *epi*morphism. Thus an axiomatic class is closed under bounded epimorphic images, and under domains of arbitrary bounded morphisms. A bijective bounded morphism is precisely an isomorphism of models in the usual sense.

$\mathcal{M}$ is called an *inner submodel* of $\mathcal{M}'$ if $X \subseteq X'$ and the inclusion function $X \hookrightarrow X'$ is a bisimulation from $\mathcal{M}$ to $\mathcal{M}'$. Then the inverse of the graph of the inclusion is a surjective bisimulation from $\mathcal{M}'$ to $\mathcal{M}$, showing that $\mathcal{M}' \models \varphi$ implies $\mathcal{M} \models \varphi$. Hence axiomatic classes are closed under inner submodels. An alternative definition of inner submodel is that $X \subseteq X'$; $P_n = P'_n \cap X$; $\leq$ is the restriction of $\leq'$ to X; and X is up-closed in $(X', \leq')$. Thus any $X \in U(\leq')$ becomes an inner submodel of $\mathcal{M}'$ by restricting $\leq'$ and the P'_n's to X. In particular, if R is a bisimulation from $\mathcal{M}$ to $\mathcal{M}'$, then the domain of R is an inner submodel of $\mathcal{M}$, while the image of R is an inner submodel of $\mathcal{M}'$.

For each point x of a model $\mathcal{M}$ we denote by $\mathcal{M}_x$ the inner submodel of $\mathcal{M}$ *generated by* x, which by definition is the submodel based on the up-closed set $[x)$. Since the inclusion $\mathcal{M}_x \hookrightarrow \mathcal{M}$ is a bisimulation we get $\mathcal{M}_x, y \models \varphi$ iff $\mathcal{M}, y \models \varphi$ for all $y \in [x)$. It follows that

$$\mathcal{M} \models \varphi \quad \text{iff} \quad \text{for all } x \text{ in } \mathcal{M},\ \ \mathcal{M}_x \models \varphi. \tag{3.3}$$

In modal logic, bounded epimorphisms are often called 'p-morphisms' (this terminology comes from [Segerberg, 1970, Segerberg, 1971], while total surjective bisimulations were first introduced in [van Benthem, 1983] as 'p-relations' . There are many relationships between these concepts. For instance, for any class $\mathbb{C}$ of models the following properties are equivalent:

- $\mathbb{C}$ is closed under bisimulation images.
- $\mathbb{C}$ is closed under total bisimulation images and inner submodels.

- $\mathbb{C}$ is invariant under bounded epimorphic images and closed under inner submodels.
- $\mathbb{C}$ is closed under domains of bounded morphisms and under bounded epimorphic images.

In particular, every axiomatic class $\mathrm{Mod}\,\Sigma$ has these properties.

4 Disjoint and Bounded Unions

Let $\{\mathcal{M}^i : i \in I\}$ be a set of IPC-models, with $\mathcal{M}^i = (X^i, \leq^i, P^i_n)_{n\in\omega}$. The *disjoint union* $\coprod_I \mathcal{M}^i$ is simply the union of a collection of pairwise disjoint copies of the $\mathcal{M}^i$'s. Formally we take this to be the model based on the set $\bigcup_I (X^i \times \{i\})$ whose quasiorder and P_n-relations are the disjoint unions of the corresponding relations in the $\mathcal{M}^i$'s. For each $i \in I$, the map $x \mapsto (x, i)$ is an injective bounded morphism $\mathcal{M}^i \rightarrowtail \coprod_I \mathcal{M}^i$ making $\mathcal{M}^i$ isomorphic to an inner submodel of $\coprod_I \mathcal{M}^i$ (viz. $\mathcal{M}^i \times \{i\}$). Since this map is a bisimulation it shows that $\mathcal{M}^i, x \models \varphi$ iff $\coprod_I \mathcal{M}^i, (x, i) \models \varphi$. Since every member of $\coprod_I \mathcal{M}^i$ is of the form (x, i), this implies that

$$\textstyle\coprod_I \mathcal{M}^i \models \varphi \quad \text{iff} \quad \text{for all } i \in I,\ \ \mathcal{M}^i \models \varphi.$$

Consequently, every axiomatic class is *closed under disjoint unions*: if $\{\mathcal{M}^i : i \in I\} \subseteq \mathrm{Mod}\,\Sigma$, then $\coprod_I \mathcal{M}^i \in \mathrm{Mod}\,\Sigma$.

A model $\mathcal{M}$ is the *bounded union* of a collection $\{\mathcal{M}^i : i \in I\}$ if the $\mathcal{M}^i$'s are all inner submodels of $\mathcal{M}$ and their union is $\mathcal{M}$ itself. Then the map $(x, i) \mapsto x$ defines a bounded epimorphism $\coprod_I \mathcal{M}^i \twoheadrightarrow \mathcal{M}$ from the disjoint union of the $\mathcal{M}^i$'s onto $\mathcal{M}$. This shows that every axiomatic class is closed under bounded unions, and also gives an alternative explanation for (3.3). More generally it implies that if a class is closed under bounded epimorphic images and disjoint unions, then it is closed under bounded unions.

Notice that any IPC-model $\mathcal{M}$ is the bounded union of the collection $\{\mathcal{M}_x : x \text{ in } \mathcal{M}\}$ of its point-generated inner submodels. Combining this with the last observation provides the following result that will be needed later:

Lemma 4.1 *If a class $\mathbb{C}$ of IPC-models is closed under bisimulation images and disjoint unions, then for any model $\mathcal{M}$,*

$$\mathcal{M} \in \mathbb{C} \quad \textit{iff} \quad \textit{for all } x \textit{ in } \mathcal{M},\ \ \mathcal{M}_x \in \mathbb{C}.$$

5 Bisimilarity From Saturation

The union of all bisimulations from $\mathcal{M}$ to $\mathcal{M}'$ is itself a bisimulation, known as the *bisimilarity* relation. This notion was developed in the theory of process algebra as a formalisation of the relation of 'observational equivalence' between states of transition systems. Hennessy and Milner [1985] proposed the idea of devising a logical system to characterise bisimilarity as the relation of 'logical equivalence' of states. Here we will say that point x of model $\mathcal{M}$ is *logically equivalent* to point x' of $\mathcal{M}'$ if for all $\varphi \in \Phi$, $\mathcal{M}, x \models \varphi$ iff $\mathcal{M}', x' \models \varphi$. If this holds we write $\mathcal{M}, x \equiv \mathcal{M}', x'$, or just $x \equiv x'$ if the models are understood.

For IPC-models, as for modal logic, the Hennessy-Milner proposal can be fulfilled in models that are saturated to some degree. In fact this needs only the weak assumption of '2-saturation', which refers to the addition of *one* constant (i.e. fewer than 2). To define this, let $\mathcal{L}_c$ be the expansion of the language $\mathcal{L}$ by the addition of a single individual constant c. An $\mathcal{L}_c$-structure has the form $(\mathcal{M}, x_c)$ with x_c being a member of the $\mathcal{L}$-structure $\mathcal{M}$ interpreting c. A set Γ of $\mathcal{L}_c$-formulas that have at most one free variable v is *satisfiable* in this structure if there is some y in $\mathcal{M}$ such that $(\mathcal{M}, x_c) \models \sigma[y]$ for all $\sigma \in \Gamma$. We may write $(\mathcal{M}, x_c) \models \Gamma[y]$ when this happens. Γ is *finitely satisfiable* in the structure if each of its finite subsets is satisfiable. An $\mathcal{L}$-structure $\mathcal{M}$ is 2-*saturated* if for each member x_c of $\mathcal{M}$, any set of $\mathcal{L}_c$-formulas that is finitely satisfiable in $(\mathcal{M}, x_c)$ must itself be satisfiable in that structure.

For any cardinal $\aleph$, $\aleph$-saturation is defined like this but using an expansion of $\mathcal{L}$ by fewer than $\aleph$ constants. In Section 8 we will observe that ultrapowers can be used to construct models that are $\aleph_1$-saturated, and hence 2-saturated. The following is a typical use of 2-saturation in Kripke models, a technique first introduced for modal logic in [Fine, 1975].

Theorem 5.1 *If $\mathcal{M}$ and $\mathcal{M}'$ are 2-saturated IPC-models, then the logical equivalence relation $\equiv$ is a bisimulation from $\mathcal{M}$ to $\mathcal{M}'$.*

Proof. Suppose $\mathcal{M}, x \equiv \mathcal{M}', x'$. Then for all $n \in \omega$, $\mathcal{M}, x \models p_n$ iff $\mathcal{M}', x' \models p_n$, which shows that bisimulation-condition B1 holds.

For the 'back' condition B2, suppose that $x' \leq' y'$ in $\mathcal{M}'$. We have to show there is some y in $\mathcal{M}$ with $x \leq y$ and $x \equiv y$. Let

$$\Gamma = \{\varphi^t(v) : \varphi \in \Phi \text{ and } \mathcal{M}', y' \models \varphi\},$$
$$\Delta = \{\neg\varphi^t(v) : \varphi \in \Phi \text{ and } \mathcal{M}', y' \not\models \varphi\}$$

(here $\neg$ is the Boolean negation symbol of $\mathcal{L}$). We will show that the set of formulas $\{c \leq v\} \cup \Gamma \cup \Delta$ is finitely satisfiable in the $\mathcal{L}_c$-structure $(\mathcal{M}, x)$.

Let $\mathcal{M}', y' \models \varphi_i$ for all $i \leq n$ and $\mathcal{M}', y' \not\models \psi_j$ for all $j \leq m$. As $x' \leq y'$, the IPC-semantics of Φ then gives that the formula

$$\varphi_1 \wedge \cdots \wedge \varphi_n \to \psi_1 \vee \cdots \vee \psi_m$$

is not true at x' in $\mathcal{M}'$. Since $x \equiv x'$, this formula is not true at x in $\mathcal{M}$, so there is some z in $\mathcal{M}$ with $x \leq z$, $\mathcal{M}, z \models \varphi_i$ for all $i \leq n$, and $\mathcal{M}, z \not\models \psi_j$ for all $j \leq m$. Hence the set

$$\{c \leq v\} \cup \{\varphi_1^t(v), \ldots, \varphi_n^t(v), \neg\psi_1^t(v), \ldots, \neg\psi_m^t(v)\}$$

is satisfiable in the $\mathcal{L}_c$-model $(\mathcal{M}, x)$ by interpreting v as z.

This confirms that $\{c \leq v\} \cup \Gamma \cup \Delta$ is finitely satisfiable in $(\mathcal{M}, x)$, so by 2-saturation of $\mathcal{M}$ it is satisfiable in $(\mathcal{M}, x)$ by some y. Then $x \leq y$ and for all $\varphi \in \Phi$, if $\mathcal{M}', y' \models \varphi$ then $\mathcal{M} \models \varphi^t[y]$, while if $\mathcal{M}', y' \not\models \varphi$ then $\mathcal{M} \not\models \varphi^t[y]$, so $x \equiv y$ as desired.

The proof of B3 is symmetric to this one for B2, using the 2-saturation of $\mathcal{M}'$. $\square$

Notice that if R is any bisimulation from $\mathcal{M}$ to $\mathcal{M}'$, then by (3.1), xRx' implies $x \equiv x'$. So logical equivalence is indeed the union (largest) of all bisimulation relations between 2-saturated models.

6 Prime Extensions

The collection $U(\leq)$ of up-closed subsets of a quasiordered set $(X, \leq)$ is a distributive lattice. New models can be built from the prime filters of this lattice. A non-empty subset F of $U(\leq)$ is a *prime filter* iff it has $\emptyset \notin F$; $Y \cap Z \in F$ iff $Y \in F$ and $Z \in F$; and $Y \cup Z \in F$ iff $Y \in F$ or $Z \in F$, for all up-closed Y, Z. For example, $F_x = \{Y \in U(\leq) : x \in Y\}$ is a prime filter for each $x \in X$.

For $H, K \subseteq U(\leq)$, we say that H is *separated from* K if for any finite subsets H' of H and K' of K we have $\bigcap H' \not\subseteq \bigcup K'$. In this context the classical Birkhoff-Stone result on the existence of prime filters takes the form of

Lemma 6.1 *If H is separated from K, then H is included in a prime filter of $U(\leq)$ that is disjoint from K.* $\square$

We define the *prime extension* of an IPC-model $\mathcal{M} = (X, \leq, P_n)_{n \in \omega}$ to be the structure

$$\mathcal{M}^* = (X^*, \subseteq, P_0^*, \ldots, P_n^*, \ldots),$$

where X^* is the set of all prime filters of $U(\leq)$, and $P_n^* = \{F \in X^* : P_n \in F\}$.

Lemma 6.2 *For any formula $\varphi \in \Phi$:*

(1) *For all $F \in X^*$, $\mathcal{M}^*, F \models \varphi$ iff $\mathcal{M}(\varphi) \in F$.*

(2) *$\mathcal{M} \models \varphi$ iff $\mathcal{M}^* \models \varphi$.*

Proof.

(1) By induction of the formation of φ. The case of $\varphi = \bot$ holds because $\mathcal{M}(\bot) = \emptyset \notin F$; and the case of $\varphi = p_n$ follows from the definition of $P_n^* = \mathcal{M}^*(p_n)$ because $P_n = \mathcal{M}(p_n)$. The inductive cases for the connectives $\wedge$ and $\vee$ are straightforward from the above-listed properties of a prime filter.

Now suppose $\varphi = (\varphi_1 \to \varphi_2)$ and assume the result for φ_1 and φ_2. Let $\mathcal{M}(\varphi) \in F$. Then for all $G \in X^*$, if $F \subseteq G$ and $\mathcal{M}^*, G \models \varphi_1$, then $\mathcal{M}(\varphi_1) \in G$ by induction hypothesis on φ_1, and $\mathcal{M}(\varphi_1 \to \varphi_2) \in G$. But $\mathcal{M}(\varphi_1) \cap \mathcal{M}(\varphi_1 \to \varphi_2) \subseteq \mathcal{M}(\varphi_2)$ by the semantics of implication, so as G is a filter $\mathcal{M}(\varphi_2) \in G$, hence $\mathcal{M}^*, G \models \varphi_2$ by hypothesis on φ_2. This establishes that $\mathcal{M}^*, F \models \varphi_1 \to \varphi_2$.

Conversely, suppose $\mathcal{M}^*, F \models \varphi_1 \to \varphi_2$. Then if $F \cup \{\mathcal{M}(\varphi_1)\}$ was separated from $\{\mathcal{M}(\varphi_2)\}$, by Lemma 6.1 there would be some $G \in X^*$ with $F \subseteq G$, $\mathcal{M}(\varphi_1) \in G$, and $\mathcal{M}(\varphi_2) \notin G$; hence $\mathcal{M}^*, G \models \varphi_1$ and $\mathcal{M}^*, G \not\models \varphi_2$ by hypothesis. But this situation contradicts $\mathcal{M}^*, F \models \varphi_1 \to \varphi_2$. Hence $F \cup \{\mathcal{M}(\varphi_1)\}$ is not separated from $\{\mathcal{M}(\varphi_2)\}$, so as F is closed under finite intersections there must be some $Y \in F$ with $Y \cap \mathcal{M}(\varphi_1) \subseteq \mathcal{M}(\varphi_2)$. This implies $Y \subseteq \mathcal{M}(\varphi_1 \to \varphi_2)$, hence $\mathcal{M}(\varphi_1 \to \varphi_2) \in F$ as F is a filter.

Thus the result holds in all cases.

(2) If $\mathcal{M} \models \varphi$, then $\mathcal{M}(\varphi)$ is X, which belongs to every prime filter, so $\mathcal{M}^*, F \models \varphi$ for all $F \in X^*$ by part (1).

Conversely, if $\mathcal{M}^* \models \varphi$, then for each $x \in X$, $\mathcal{M}^*, F_x \models \varphi$, hence $\mathcal{M}(\varphi) \in F_x$ by (1), which means that $\mathcal{M}, x \models \varphi$. □

Part (2) of this Lemma implies that axiomatic classes are *invariant under prime extensions*: $\mathcal{M} \in \mathrm{Mod}\,\Sigma$ iff $\mathcal{M}^* \in \mathrm{Mod}\,\Sigma$.

7 Definable Extensions

The collection $U(\mathcal{M}) = \{\mathcal{M}(\varphi) : \varphi \in \Phi\}$ of 'definable' up-closed subsets of a model $\mathcal{M}$ is always countable, so may be much smaller than $U(\leq)$. But it is a distributive lattice in its own right – indeed a sub-Heyting-algebra of $U(\leq)$ – and so has its own prime filters. We define

$$\mathcal{M}^\delta = (X^\delta, \subseteq, P_0^\delta, \dots, P_n^\delta, \dots),$$

where X^δ is the set of all prime filters of $U(\mathcal{M})$, and $P_n^\delta = \{F \in X^\delta : P_n \in F\}$. $\mathcal{M}^\delta$ will be called the *definable extension* of $\mathcal{M}$.

A version of Lemma 6.2 can be proved for $\mathcal{M}^\delta$, implying that *axiomatic classes are invariant under definable extensions*. But we can also deduce this from invariance under prime extensions, by analysing the relationship between

$\mathcal{M}^\delta$ and $\mathcal{M}^*$. Note first that X^δ is not a subset of X^*, since a prime filter of $U(\mathcal{M})$ will be a filter of $U(\leq)$ but may not be prime in $U(\leq)$. The exact relationship between the two constructions is given by the map $f_\mathcal{M} : X^* \to X^\delta$ specified by $f_\mathcal{M}(F) = F \cap U(\mathcal{M})$ for all $F \in X^*$.

Lemma 7.1 *f_M is a bounded epimorphism $\mathcal{M}^* \twoheadrightarrow \mathcal{M}^\delta$.*

Proof. This is an instance of a well-established result in the duality theory of Heyting algebras: $f_\mathcal{M}$ is the dual map to the inclusion homomomorphism $U(\mathcal{M}) \hookrightarrow U(\leq)$. Details can be found for instance in [Goldblatt, 1989, Section 2]. Here is a sketch of the main points.

First, to show $f_\mathcal{M}$ is surjective, for each $H \in X^\delta$, H is separated from $U(\mathcal{M}) - H$, so by Lemma 6.1 there is a prime filter $F \in X^*$ extending H and disjoint from $U(\mathcal{M}) - H$, hence $F \cap U(\mathcal{M}) = H$.

Clearly if $F \subseteq G$ in $U(\leq)$, then $f_\mathcal{M}(F) \subseteq f_\mathcal{M}(G)$ in $U(\mathcal{M})$, so the right-to-left implication of (3.2) holds. For the converse, if $F \in X^*$ and $f_\mathcal{M}(F) \subseteq H$ in $U(\mathcal{M})$, then $F \cup H$ is separated from $U(\mathcal{M}) - H$, so there exists $G \in X^*$ that extends $F \cup H$ and is disjoint from $U(\mathcal{M}) - H$, hence has $F \subseteq G$ and $F \cap U(\mathcal{M}) = H$. □

Corollary 7.2 *Let $\mathbb{C}$ be a class of IPC-models that is closed under images of total bisimulations. Then for all $\mathcal{M}$, $\mathcal{M}^* \in \mathbb{C}$ iff $\mathcal{M}^\delta \in \mathbb{C}$. Hence $\mathbb{C}$ is closed/invariant under prime extensions iff it is closed/invariant under definable extensions.*

Proof. The graph of $f_\mathcal{M}$ and its inverse give total surjective bisimulations in each direction between $\mathcal{M}^*$ and $\mathcal{M}^\delta$. □

Corollary 7.3 *For any formula $\varphi \in \Phi$:*

(1) *For all $F \in X^\delta$, $\mathcal{M}^\delta, F \models \varphi$ iff $\mathcal{M}(\varphi) \in F$.*

(2) *$\mathcal{M} \models \varphi$ iff $\mathcal{M}^\delta \models \varphi$.*

Proof.

(1) If $F = f_\mathcal{M}(G)$, then $\mathcal{M}^\delta, F \models \varphi$ iff $\mathcal{M}^*, G \models \varphi$ by (3.1). But also $\mathcal{M}(\varphi) \in F$ iff $\mathcal{M}(\varphi) \in G$, so the result follows from Lemma 6.2(1).

(2) From Lemma 6.2(2), as $\mathcal{M}^* \in \mathrm{Mod}\,\varphi$ iff $\mathcal{M}^\delta \in \mathrm{Mod}\,\varphi$ by Corollary 7.2. □

$\mathcal{M}^\delta$ need not be a genuine 'extension' of $\mathcal{M}$: it may have lower cardinality. The natural map $x \mapsto \{\mathcal{M}(\varphi) : \mathcal{M}, x \models \varphi\}$ of X into X^δ identifies any two points that are logically equivalent. Hence this map will be injective iff $x \equiv y$ implies $x = y$ in $\mathcal{M}$. The natural map $x \mapsto F_x$ of X into X^* is injective iff the quasiorder $\leq$ is anti-symmetric.

We now study the relationship between definable extensions and 2-saturated models.

Theorem 7.4 *For any IPC-model $\mathcal{M}$, $\mathcal{M}^\delta$ is a bounded epimorphic image of any 2-saturated model $\mathcal{N}$ such that for all $\varphi \in \Phi$, $\mathcal{N} \models \varphi$ iff $\mathcal{M} \models \varphi$.*

Proof. Let $\mathcal{N} = (Y, \leq^{\mathcal{N}}, \dots)$. Define a map $\eta : Y \to X^\delta$ by putting, for any $x \in Y$,

$$\eta(x) = \{\mathcal{M}(\varphi) : \mathcal{N}, x \models \varphi\}.$$

This is well-defined, because if $\mathcal{M}(\varphi) = \mathcal{M}(\psi)$ then $\mathcal{M} \models \varphi \leftrightarrow \psi$, hence $\mathcal{N} \models \varphi \leftrightarrow \psi$ by hypothesis on $\mathcal{N}$, and so $\mathcal{N}, x \models \varphi$ iff $\mathcal{N}, x \models \psi$. It is readily checked that $\eta(x)$ is a prime filter of $U(\mathcal{M})$, so belongs to X^δ.

If $x \leq^{\mathcal{N}} y$, then $\mathcal{N}, x \models \varphi$ implies $\mathcal{N}, y \models \varphi$, hence $\eta(x) \subseteq \eta(y)$; so the right-to-left implication of (3.2) holds. For the converse, let $\eta(x) \subseteq y'$ in $\mathcal{M}^\delta$. We have to show there is some y in $\mathcal{N}$ with $x \leq^{\mathcal{N}} y$ and $\eta(y) = y'$. The proof is similar to that of Theorem 5.1. Let

$$\begin{aligned} \Gamma &= \{\varphi^t(v) : \varphi \in \Phi \text{ and } \mathcal{M}^\delta, y' \models \varphi\}, \\ \Delta &= \{\neg\varphi^t(v) : \varphi \in \Phi \text{ and } \mathcal{M}^\delta, y' \not\models \varphi\}. \end{aligned}$$

It suffices to show that the set $\{c \leq v\} \cup \Gamma \cup \Delta$ is finitely satisfiable in the $\mathcal{L}_c$-structure $(\mathcal{N}, x)$. For then by 2-saturation of $\mathcal{N}$ it is satisfiable in $(\mathcal{N}, x)$ by some y. Then $x \leq^{\mathcal{N}} y$ and $(\mathcal{N}, y) \equiv (\mathcal{M}^\delta, y')$. Using Corollary 7.3(1), this implies

$$\mathcal{M}(\varphi) \in y' \quad \text{iff} \quad \mathcal{M}^\delta, y' \models \varphi \quad \text{iff} \quad \mathcal{N}, y \models \varphi \quad \text{iff} \quad \mathcal{M}(\varphi) \in \eta(y),$$

so $y' = \eta(y)$ as desired.

For the proof of finite satisfiability, let $\mathcal{M}^\delta, y' \models \varphi_i$ for all $i \leq n$ and $\mathcal{M}^\delta, y' \not\models \psi_j$ for all $j \leq m$. Let φ be the formula

$$\varphi_1 \wedge \dots \wedge \varphi_n \to \psi_1 \vee \dots \vee \psi_m. \tag{7.1}$$

As $\eta(x) \subseteq y'$, φ is not true at $\eta(x)$ in $\mathcal{M}^\delta$, hence $\mathcal{M}(\varphi) \notin \eta(x)$ by Corollary 7.3(1), so $\mathcal{N}, x \not\models \varphi$. Hence there is some z in Y such that the set

$$\{c \leq v\} \cup \{\varphi_1^t(v), \dots, \varphi_n^t(v), \neg\psi_1^t(v), \dots, \neg\psi_m^t(v)\}$$

is satisfiable in the $\mathcal{L}_c$-model $(\mathcal{N}, x)$ by interpreting v as z. This completes the proof that η is a bounded morphism.

Finally, to show η is surjective, we take any $y' \in X^\delta$ and show that the set $\Gamma \cup \Delta$ as above is finitely satisfiable in $(\mathcal{N}, x)$. Hence it is satisfiable by some y which then has $(\mathcal{N}, y) \equiv (\mathcal{M}^\delta, y')$ and so $y' = \eta(y)$ as before.

So, suppose again that $\mathcal{M}^\delta, y' \models \varphi_i$ for $i \leq n$ and $\mathcal{M}^\delta, y' \not\models \psi_j$ for $j \leq m$. Then if φ is the formula (7.1), we have $\mathcal{M}^\delta, y' \not\models \varphi$. This time we infer $\mathcal{M}^\delta \not\models \varphi$, hence $\mathcal{M} \not\models \varphi$ by Corollary 7.3(2), hence $\mathcal{N} \not\models \varphi$ by hypothesis on $\mathcal{N}$. Thus there is some z in $\mathcal{N}$ satisfying $\{\varphi_1^t(v), \dots, \varphi_n^t(v), \neg\psi_1^t(v), \dots, \neg\psi_m^t(v)\}$ as required. $\square$

It is also possible to construct bounded epimorphisms from 2-saturated models onto the *prime* extension $\mathcal{M}^*$, but only by working with models for a typically uncountable language extending $\mathcal{L}$ by adding monadic predicates defining each up-closed subset of $\mathcal{M}$. Existence theorems for saturated models for such languages are demanding: to construct them as ultrapowers requires the theory of 'good' ultrafilters [Chang and Keisler, 1973, Section 6.1].

8 Ultraproducts and Ultrapowers

Let $\{\mathcal{M}^i : i \in I\}$ be a set of IPC-models, and D an *ultrafilter* over the index set I. Recall that this means that D is a *filter*, i.e. in general $Y \cap Z \in D$ iff $Y \in D$ and $Z \in D$, and that exactly one of Y and $I - Y$ belongs to D for each $Y \subseteq I$.

We review the definition of the *ultraproduct*

$$\prod_D \mathcal{M}^i = (\prod_D X^i, \leq^D, P_0^D, \dots, P_n^D, \dots).$$

An equivalence relation $f =_D g$ between functions $f, g \in \prod_I X^i$ is defined to mean that $\{i \in I : f(i) = g(i)\} \in D$. Then $\prod_D X^i$ is the set of equivalence classes of $\prod_I X^i$ under $=_D$. Writing f^D for the equivalence class of f, we have

$$\begin{aligned} f^D \leq^D g^D \quad &\text{iff} \quad \{i \in I : f(i) \leq^i g(i)\} \in D, \\ f^D \in P_n^D \quad &\text{iff} \quad \{i \in I : f(i) \in P_n^i\} \in D. \end{aligned}$$

If all the $\mathcal{M}^i$'s are the same model $\mathcal{M}$, then the ultraproduct is denoted $\prod_D \mathcal{M} = (\prod_D X, \dots)$ and called an *ultrapower* of $\mathcal{M}$. There is a natural map $x \mapsto x^D$ from X into $\prod_D X$ defined by $x^D = f_x^D$, where f_x is the constant function having $f_x(i) = x$ for all $i \in I$.

The fundamental theorem of Łoś states that for any $\mathcal{L}$-formula $\sigma(v_1, \dots, v_n)$, with free variables amongst those listed, and any $f_1, \dots, f_n \in \prod_I X^i$,

$$\prod_D \mathcal{M}^i \models \sigma[f_1^D, \dots, f_n^D] \quad \text{iff} \quad \{i \in I : \mathcal{M}^i \models \sigma[f_1(i), \dots, f_n(i)]\} \in D.$$

Hence if σ is a *sentence*,

$$\prod_D \mathcal{M}^i \models \sigma \quad \text{iff} \quad \{i \in I : \mathcal{M}^i \models \sigma\} \in D.$$

Taking σ to be the sentence σ_{qo} expressing that $\leq$ is a quasiorder, or the sentence $\forall v \forall w(v \leq w \wedge \pi(v) \to \pi(w))$ expressing that P_n is up-closed, then shows that $\prod_D \mathcal{M}$ is an IPC-model. Taking the cases that σ is $\varphi^t(v)$ or $\forall v \varphi^t$ for some $\varphi \in \Phi$, we get in terms of IPC-semantics that

$$\prod_D \mathcal{M}^i, f^D \models \varphi \quad \text{iff} \quad \{i \in I : \mathcal{M}^i, f(i) \models \varphi\} \in D,$$

and

$$\prod_D \mathcal{M}^i \models \varphi \quad \text{iff} \quad \{i \in I : \mathcal{M}^i \models \varphi\} \in D.$$

For ultrapowers these imply that for all x in $\mathcal{M}$,

$$\prod_D \mathcal{M}, x^D \models \varphi \quad \text{iff} \quad \mathcal{M}, x \models \varphi; \tag{8.1}$$

and that

$$\prod_D \mathcal{M} \models \varphi \quad \text{iff} \quad \mathcal{M} \models \varphi. \tag{8.2}$$

Łoś's Theorem entails that an elementary class of $\mathcal{L}$-structures is closed under ultraproducts, and its complement is closed under ultrapowers. In particular, this holds for the axiomatic classes Mod Σ. For a class like Mod φ that is axiomatized by a single sentence, the complement is closed under ultraproducts.

We will need to use the following ultraproduct version of the Compactness Theorem.

Lemma 8.1 *Let $\mathbb{C}$ be a class of $\mathcal{L}$-structures that is closed under ultraproducts. If Γ is a set of $\mathcal{L}$ formulas having at most one free variable, and each finite subset of Γ is satisfiable in a model from $\mathbb{C}$, then Γ is satisfiable in a model from $\mathbb{C}$.*

Proof. This is standard [Chang and Keisler, 1973, 4.1.11]. Let I be the set of all finite subsets of Γ. For each $i \in I$ there is some $\mathcal{M}^i \in \mathbb{C}$ and some x_i in $\mathcal{M}^i$ with $\mathcal{M} \models i[x_i]$. There is an ultrafilter D over I such that for each $\sigma \in \Gamma$, D contains the set $J_\sigma = \{i : \sigma \in i\}$. Put $f(i) = x_i$. Then $J_\sigma \subseteq \{i \in I : \mathcal{M}^i \models \sigma[f(i)]\}$, so by Łoś's Theorem $\prod_D \mathcal{M}^i \models \sigma[f^D]$. Thus Γ is satisfied by f^D in $\prod_D \mathcal{M}^i \in \mathbb{C}$. □

There is a significant relationship between ultraproducts and ultrapowers of Kripke models that was first identified by the author in the modal context. Here is takes the following form:

Theorem 8.2 *For any set $\{\mathcal{M}^i : i \in I\}$ of IPC-models and any ultrafilter D over I, there is an injective bounded morphism*

$$\prod_D \mathcal{M}^i \rightarrowtail \prod_D (\coprod_I \mathcal{M}^i)$$

making the ultraproduct $\prod_D \mathcal{M}^i$ isomorphic to an inner submodel of the D-ultrapower of the disjoint union $\coprod_I \mathcal{M}^i$ of the $\mathcal{M}^i$'s.

Proof. For $f \in \prod_D X$, define $\hat{f}(i) = (f(i), i)$ to get $\hat{f} \in \prod_I (\coprod_I X^i)$. Then the asserted bounded morphism is $f^D \mapsto \hat{f}^D$ – see [Goldblatt, 1989, 3.8.3]. □

Corollary 8.3 *If a class $\mathbb{C}$ of IPC-models is closed under bisimulation images, disjoint unions and ultrapowers, then it is closed under ultraproducts.*

Proof. Closure under bisimulation images implies closure under inner submodels and isomorphism. □

One advantage of working with Φ and $\mathcal{L}$ is that for countable languages 2-saturated ultrapowers are readily available. To explain this, recall that an ultrafilter D over a set I is *countably incomplete* if it there is a countable set $E \subseteq D$ with $\bigcap E = \emptyset$. For example, if I is itself countable then any nonprincipal D over I is countably incomplete, as shown by taking $E = \{I - \{i\} : i \in I\}$. The following result is proven in [Chang and Keisler, 1973, Theorem 6.1.1], and holds for models for any *countable* first-order language.

Theorem 8.4 *If D is a countably incomplete ultrafilter over a set I, then for any set $\{\mathcal{M}^i : i \in I\}$ of models, the ultraproduct $\prod_D \mathcal{M}^i$ is $\aleph_1$-saturated.* □

We use this result to show how closure under definable extensions can lead to closure under ultrapowers, by iterating the ultrapower construction.

Theorem 8.5 *If a class $\mathbb{C}$ of IPC-models is closed under total bisimulation images and invariant under definable extensions, then it is invariant under ultrapowers.*

Proof. Let $\prod_D \mathcal{M}$ be any ultrapower of some model $\mathcal{M}$. Take any countably incomplete ultrafilter E (e.g. any nonprincipal ultrafilter on ω). By Theorem 8.4, the ultrapower $\prod_E(\prod_D \mathcal{M})$ is $\aleph_1$-saturated, and using (8.2) twice we get $\prod_E(\prod_D \mathcal{M}) \models \varphi$ iff $\mathcal{M} \models \varphi$. Hence by Theorem 7.4 there is a bounded epimorphism $f : \prod_E(\prod_D \mathcal{M}) \twoheadrightarrow \mathcal{M}^\delta$. The inverse of f is a total bisimulation from $\mathcal{M}^\delta$ onto $\prod_E(\prod_D \mathcal{M})$.

Thus if $\mathcal{M} \in \mathbb{C}$, then $\mathcal{M}^\delta \in \mathbb{C}$ and hence $\prod_E(\prod_D \mathcal{M}) \in \mathbb{C}$ by the given closure conditions. Applying Theorem 7.4 now to $\prod_D \mathcal{M}$, since $\prod_E(\prod_D \mathcal{M}) \models \varphi$ iff $\prod_D \mathcal{M} \models \varphi$, there is a bounded epimorphism $\prod_E(\prod_D \mathcal{M}) \twoheadrightarrow (\prod_D \mathcal{M})^\delta$, so $(\prod_D \mathcal{M})^\delta \in \mathbb{C}$, and finally $\prod_D \mathcal{M} \in \mathbb{C}$ by invariance under definable extensions.

This proves that $\mathbb{C}$ is closed under ultrapowers. But now if $\prod_D \mathcal{M} \in \mathbb{C}$ then $\prod_E(\prod_D \mathcal{M}) \in \mathbb{C}$ by this closure just proven, hence $\mathcal{M}^\delta \in \mathbb{C}$ by closure under bounded epimorphic images, which finally gives $\mathcal{M} \in \mathbb{C}$ by invariance under definable extensions. □

In this proof, if D is principal then $\prod_D \mathcal{M} \cong \mathcal{M}$, while if D is countably incomplete then we can apply Theorem 7.4 directly to get a bounded epimorphism from $\prod_D \mathcal{M}$ onto $\mathcal{M}^\delta$, hence $\prod_D \mathcal{M} \in \mathbb{C}$ iff $\mathcal{M}^\delta \in \mathbb{C}$ iff $\mathcal{M} \in \mathbb{C}$. So, intriguingly, the use of the iterated ultrapower is required only to cover the case that D is a nonprincipal but countably *complete* ultrafilter, something whose existence is equivalent to that of a measurable cardinal and cannot be proved in ZFC.

9 Characterizing Axiomatizability

We are now ready to put together our main result:

Theorem 9.1 *For any class $\mathbb{C}$ of IPC-models, the following are equivalent.*

(1) $\mathbb{C}$ *is axiomatic, i.e.* $\mathbb{C} = Mod\,\Sigma$ *for some* $\Sigma \subseteq \Phi$.

(2) $\mathbb{C}$ *is closed under bisimulation images and disjoint unions, and invariant under prime extensions.*

(3) $\mathbb{C}$ *is closed under bisimulation images and disjoint unions, and invariant under definable extensions.*

(4) $\mathbb{C}$ *is closed under bisimulation images and disjoint unions, and invariant under ultrapowers.*

Proof. (1) implies (2): this has been explained in Sections 3, 5 and 6.

(2) implies (3): Corollary 7.2.

(3) implies (4): Theorem 8.5.

(4) implies (1): this is essentially the argument of [Rodenburg, 1986, 13.8]. Suppose (4) holds, and let $\Sigma = \{\varphi \in \Phi : \mathbb{C} \models \varphi\}$ be the set of all IPC-formulas that are true in every member of $\mathbb{C}$. Then $\mathbb{C} \subseteq \mathrm{Mod}\,\Sigma$ by definition, and we prove the converse inclusion.

Let $\mathcal{M} \in \mathrm{Mod}\,\Sigma$. To show $\mathcal{M} \in \mathbb{C}$ it suffices, by Lemma 4.1, to show that each point-generated inner submodel of $\mathcal{M}$ belongs to $\mathbb{C}$. But each such submodel belongs to $\mathrm{Mod}\,\Sigma$ by (3.3), so we may as well assume that $\mathcal{M}$ is generated by one of its points x. Then we prove that x in $\mathcal{M}$ is logically equivalent to a point of some model in $\mathbb{C}$. A variant of this argument has already been used twice: we set

$$\Gamma = \{\varphi^t(v) : \varphi \in \Phi \text{ and } \mathcal{M}, x \models \varphi\}$$
$$\Delta = \{\neg\varphi^t(v) : \varphi \in \Phi \text{ and } \mathcal{M}, x \not\models \varphi\},$$

and show that $\Gamma \cup \Delta$ is finitely satisfiable in $\mathbb{C}$. If $\mathcal{M}, x \models \varphi_i$ for all $i \leq n$ and $\mathcal{M}, x \not\models \psi_j$ for all $j \leq m$, then if φ is the formula

$$\varphi_1 \wedge \cdots \wedge \varphi_n \to \psi_1 \vee \cdots \vee \psi_m,$$

we have $\mathcal{M}, x \not\models \varphi$, so $\varphi \notin \Sigma$ as $\mathcal{M} \models \Sigma$. By definition of Σ, φ must then be false at some point of some member of $\mathbb{C}$: that point realises the set

$$\{\varphi_1^t(v), \ldots, \varphi_n^t(v), \neg\psi_1^t(v), \ldots, \neg\psi_m^t(v)\}.$$

This shows that $\Gamma \cup \Delta$ is finitely satisfiable in $\mathbb{C}$. But $\mathbb{C}$ is closed under ultraproducts by Corollary 8.3, so by Lemma 8.1 there is some model $\mathcal{N} \in \mathbb{C}$ and some point y of $\mathcal{N}$ such that $\mathcal{N} \models (\Gamma \cup \Delta)[y]$, hence $(\mathcal{N}, y) \equiv (\mathcal{M}, x)$.

Now let D be a countably incomplete ultrafilter. The ultrapowers $\prod_D \mathcal{N}$ and $\prod_D \mathcal{M}$ are both 2-saturated (Theorem 8.4), and $\prod_D \mathcal{N} \in \mathbb{C}$ by (4). Moreover, by (8.1) we have $(\prod_D \mathcal{N}, y^D) \equiv (\mathcal{N}, y)$ and $(\mathcal{M}, x) \equiv (\prod_D \mathcal{M}, x^D)$, hence $(\prod_D \mathcal{N}, y^D) \equiv (\prod_D \mathcal{M}, x^D)$.

But by Theorem 5.1, the logical equivalence relation $\equiv$ is a bisimulation from $\prod_D \mathcal{N}$ to $\prod_D \mathcal{M}$, so if we can show it is *surjective*, then we will get $\prod_D \mathcal{M} \in \mathbb{C}$ by closure under bisimulation images, and then $\mathcal{M} \in \mathbb{C}$ by invariance under ultrapowers, completing the proof that $\mathbb{C} = \mathrm{Mod}\,\Sigma$ and establishing the Theorem.

Now the $\mathcal{L}$-formula $\forall w(v \leq w)$ is satisfied by x in $\mathcal{M}$, since x generates $\mathcal{M}$, and so by Łoś's Theorem this formula is satisfied by x^D in $\prod_D \mathcal{M}$. Hence for any point z' of $\prod_D \mathcal{M}$ we have $x^D \leq^D z'$, so as $y^D \equiv x^D$ the bisimulation condition B2 gives some z in $\prod_D \mathcal{N}$ such that $(y^D \leq z$ and) $z \equiv z'$. This proves $\equiv$ is surjective as required. $\square$

Of course we can obtain further characterizations of axiomatic classes by replacing "closed under bisimulation images" in any of (2)–(4) by any of the equivalent alternatives listed at the end of Section 3.

Finally, to characterize classes of the form $\mathrm{Mod}\,\varphi$ for a single formula φ, just replace "invariant under ultrapowers" in Theorem 9.1 by "closed under ultrapowers, and the complement $\overline{\mathbb{C}} = \{\mathcal{M} : \mathcal{M} \notin \mathbb{C}\}$ is closed under ultraproducts". The proof of this is standard: if the stronger condition holds for $\mathbb{C}$, then $\mathbb{C}$ has form $\mathbb{C} = \mathrm{Mod}\,\Sigma$, and there must be some finite $\Sigma_i \subseteq \Sigma$ such that $\mathbb{C} = \mathrm{Mod}\,\Sigma_i = \mathrm{Mod}\,(\bigwedge \Sigma_i)$. For if not, then for each such Σ_i there would be a model $\mathcal{M}^i \models \Sigma_i$ with $\mathcal{M}^i \not\models \Sigma$, hence $\mathcal{M}^i \in \overline{\mathbb{C}}$. But then by the construction in the proof of Lemma 8.1, we could construct an ultraproduct of these $\mathcal{M}^i$'s having $\prod_D \mathcal{M}^i \models \Sigma$, contradicting the closure of $\overline{\mathbb{C}}$ under ultraproducts.

10 Related and Further Work

Our Theorem 9.1 shows that a certain *logically* specified notion, viz. an axiomatic model class, has a *structural* characterisation in terms of closure under algebraic constructions. The first characterisation of this kind was the famous "variety theorem" of [Birkhoff, 1935], which showed that the *equationally* definable classes of abstract algebras are just those that are closed under homomorphic images, subalgebras and direct products. There have been many other such theorems developed subsequently, a notable example being the celebrated Keisler-Shelah characterisation of elementary (i.e. first-order definable) classes of structures as those that are closed under isomorphism and ultraproducts and have their complements closed under ultrapowers.

Results of this kind have recently been developed for certain classes of *coalgebras*, adapting ideas from modal logic to coalgebraic theory through the observation [Rutten, 1995] that Kripke models for propositional modal logic are coalgebras for a particular functor $T : \mathbf{Set} \to \mathbf{Set}$ on the category of sets. In [Goldblatt, 2001, Goldblatt, 2003a] a study is made of so-called *polynomial* coalgebras, in which T is any functor built from the identity functor and/or constant functors using the polynomial operations of products, coproducts and exponentials with constant exponent. A notion of *ultrafilter enlargement* of a polynomial coalgebra is developed, and it is shown that a class of polynomial coalgebras is the class of all models of a set of Boolean combinations of equations of a certain type precisely when it is closed under bisimulation images, disjoint unions and ultrafilter enlargements. In [Goldblatt, 2003b], ultrafilter enlargements are replaced in this result by a certain modified ultrapower construction. Section 8 of that paper gives a discussion of the analogy between such results and Birkhoff's theorem, as well as surveying the relevant literature in the theory of coalgebras.

Polynomial functors provide a rather specific class of coalgebras, to be thought of as *deterministic* transition systems. To model non-determinism requires use of the powerset functor, as indeed does the representation of a Kripke model for modal logic as a coalgebra. So it would be of interest to extend these characterisation theorem to coalgebras of functors whose formation involves powersets, or indeed to any kind of endofunctor on **Set**. A notion of ultrafilter extension for such general functors has been very recently developed in [Kupke *et al.*, 2005], raising the question of how to develop a logical specification of classes of coalgebras closed under the construction. It would also be of interest to adapt this line of enquiry to coalgebraic abstractions of IPC-models. Here it may be relevant to consider the observations of [Palmigiano, 2004] about duality between Heyting algebras and coalgebras for a certain Vietoris functor on partially-ordered Stone spaces, as well as the coalgebraic perspective on Heyting duality of [Davey and Galati, 2003].

References

[Balbes and Dwinger, 1974] R. Balbes and P. Dwinger. *Distributive Lattices.* University of Missouri Press, Columbia MO, 1974.

[Birkhoff, 1935] Garrett Birkhoff. On the structure of abstract algebras. *Proceedings of the Cambridge Philosophical Society*, 31:433–454, 1935.

[Chang and Keisler, 1973] C. C. Chang and H. J. Keisler. *Model Theory.* North-Holland, Amsterdam, 1973.

[Davey and Galati, 2003] Brian A. Davey and John C. Galati. A coalgebraic view of Heyting duality. *Studia Logica*, 75(3):259–270, 2003.

[Fine, 1975] Kit Fine. Some connections between elementary and modal logic. In Stig Kanger, editor, *Proceedings of the Third Scandinavian Logic Symposium*, pages 15–31. North-Holland, 1975.

[Goldblatt, 1989] Robert Goldblatt. Varieties of complex algebras. *Annals of Pure and Applied Logic*, 44:173–242, 1989.

[Goldblatt, 2001] Robert Goldblatt. What is the coalgebraic analogue of Birkhoff's variety theorem? *Theoretical Computer Science*, 266:853–886, 2001.

[Goldblatt, 2003a] Robert Goldblatt. Enlargements of polynomial coalgebras. In Rod Downey et al., editors, *Proceedings of the 7th and 8th Asian Logic Conferences*, pages 152–192. World Scientific, 2003.

[Goldblatt, 2003b] Robert Goldblatt. Observational ultraproducts of polynomial coalgebras. *Annals of Pure and Applied Logic*, 123:235–290, 2003.

[Hennessy and Milner, 1985] Matthew Hennessy and Robin Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the Association for Computing Machinery*, 32:137–161, 1985.

[Kripke, 1965] Saul A. Kripke. Semantical analysis of intuitionistic logic I. In J. N. Crossley and M. A. E. Dummett, editors, *Formal Systems and Recursive Functions*, pages 92–130. North-Holland, 1965.

[Kupke *et al.*, 2005] Clemens Kupke, Alexander Kurz, and Dirk Pattinson. Ultrafilter extensions for coalgebras. In J. Fiadeiro, N. Harmann, M. Roggenbach, and J. Rutten, editors, *Algebra and Coalgebra in Computer Science: First International Conference, CALCO 2005*, volume 3629 of *Lecture Notes in Computer Science*, pages 263–277. Springer, 2005.

[Palmigiano, 2004] Alessandra Palmigiano. A coalgebraic view on positive modal logic. *Theoretical Computer Science*, 327:175–195, 2004.

[Rasiowa and Sikorski, 1963] H. Rasiowa and R. Sikorski. *The Mathematics of Metamathematics*. PWN–Polish Scientific Publishers, Warsaw, 1963.

[Rodenburg, 1986] Pieter Hendrik Rodenburg. *Intuitionistic Correspondence Theory*. PhD thesis, University of Amsterdam, 1986.

[Rutten, 1995] J. J. M. M. Rutten. A calculus of transition systems (towards universal coalgebra). In Alban Ponse, Maarten de Rijke, and Yde Venema, editors, *Modal Logic and Process Algebra*, CSLI Lecture Notes No. 53, pages 231–256. CSLI Publications, Stanford, California, 1995.

[Segerberg, 1970] Krister Segerberg. Modal logics with linear alternative relations. *Theoria*, 36:301–322, 1970.

[Segerberg, 1971] Krister Segerberg. *An Essay in Classical Modal Logic*, volume 13 of *Filosofiska Studier*. Uppsala Universitet, 1971.

[van Benthem, 1983] J. F. A. K. van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Naples, 1983.

[Venema, 1999] Yde Venema. Model definability, purely modal. In Jelle Gerbrandy, Maarten Marx, Maarten de Rijke, and Yde Venema, editors, *JFAK. Essays Dedicated to Johan van Benthem on the Occasion of his 50th Birthday*, Vossiuspers. Amsterdam University Press, 1999. ISBN 90 5629 104 1, `www.illc.uva.nl/j50`.