Sahlqvist correspondence for modal mu-calculus

Johan van Benthem, Nick Bezhanishvili, Ian Hodkinson

May 20, 2011

Abstract

We define analogues of modal Sahlqvist formulas for the modal mu-calculus, and prove a correspondence theorem for them.

1 Introduction

The modal mu-calculus provides a perspicuous way of isolating essential laws of induction and recursion generalizing computational logics such as PDL, CTL, and CTL*. This paper adds one more strand to its exploration, going back to a traditional modal concern: frame correspondence theory. It was observed in [5] how the usual method for obtaining frame correspondents for Sahlqvist-type axioms can be applied to non-first-order axioms like Löb's Axiom whose antecedents have a special 'PIA syntax' supporting a minimal valuation that is definable in the classical fixed-point language FO+LFP. It is then natural to look for a balance on both sides, in terms of generalized Sahlqvist forms in the language of the modal mu-calculus that support this style of analysis. Such a generalization is found in this paper, by employing additional notions and techniques from [6]. We will use only semantic standard models here, but the latter paper also considers generalized models for the mu-calculus with restrictions on the predicates that are available in the process of fixed-point approximation.

We will not look into completeness versions of Sahlqvist's Theorem in this paper, except for a few remarks on the existence of proof systems that match semantic frame correspondence arguments. However, this research is part of a larger project on analyzing special-purpose logics based on the modal mu-calculus, and finding general techniques for their completeness proofs, which are still lacking today. An important bridge in obtaining completeness from correspondence results for Sahlqvist axioms has been the celebrated Esakia Lemma [11] tying modal semantics to topological spaces. This is just one of the many strategic points in research on modal logic and beyond where Leo Esakia has shown the way to so many of us. We are happy to dedicate this article to the memory of this great teacher, colleague, and friend.

2 Preliminaries

Before we start, we briefly go through the background material and notation needed for the paper. Our terseness is due to lack of space, not of respect.

2.1 Modal mu-calculus

We fix disjoint sets $\mathcal{P}$ of propositional atoms and $\mathcal{V}$ of fixed point variables. We write $p, q, s, \ldots$ for propositional atoms, and $X, Y, Z, \ldots$ for fixed point variables.

Any element of $\mathcal{P} \cup \mathcal{V}$ is a modal mu-formula, as are $\top, \bot$. If φ, ψ are modal mu-formulas then so are $\varphi \wedge \psi$, $\varphi \vee \psi$, $\Diamond\varphi$, $\Box\varphi$, and if $X \in \mathcal{V}$ and every free occurrence of X in φ is positive (in the scope of an even number of negations), then $\mu X \varphi$ and $\nu X \varphi$ are modal mu-formulas. We use the usual abbreviations $\to, \leftrightarrow$. An occurrence of X in φ is said to be *bound* if it is in the scope of a μX or νX, and *free,* otherwise. For convenience, occurrences of propositional atoms will also be called 'free' occurrences. A *sentence* is a modal mu-formula with no free fixed point variables.

We write $\varphi(p_1, \ldots, p_n, X_1, \ldots, X_m)$ to indicate that the atoms and free variables in φ are among $p_1, \ldots, p_n$ and $X_1, \ldots, X_m$, respectively. It will be implicit that $p_1, \ldots, p_n, X_1, \ldots, X_m$ are pairwise distinct. For modal mu-formulas φ and ψ, and $\xi \in \mathcal{P} \cup \mathcal{V}$, $\varphi(\psi/\xi)$ denotes what we get by replacing all free occurrences of ξ in φ by ψ.

A *frame* is a pair $\mathcal{F} = (W, R)$, where W is a non-empty set and $R \subseteq W \times W$. An *assignment into* $\mathcal{F}$ is a map $h : \mathcal{P} \cup \mathcal{V} \to \wp(W)$. For $\xi \in \mathcal{P} \cup \mathcal{V}$ and $U \subseteq W$, we write h^U_ξ for the assignment that agrees with h on all symbols other than ξ and whose value on ξ is U. We define $[\![\varphi]\!]_h \subseteq W$ by induction on φ; the frame $\mathcal{F}$ is implicit in the notation. For $\varphi \in \mathcal{P} \cup \mathcal{V}$ we put $[\![\varphi]\!]_h = h(\varphi)$. $[\![\top]\!]_h = W$, and $[\![\bot]\!]_h = \emptyset$. We put $[\![\neg\varphi]\!]_h = W \setminus [\![\varphi]\!]_h$, $[\![\varphi \wedge \psi]\!]_h = [\![\varphi]\!]_h \cap [\![\psi]\!]_h$, $[\![\varphi \vee \psi]\!]_h = [\![\varphi]\!]_h \cup [\![\psi]\!]_h$, $[\![\Diamond\varphi]\!]_h = \{a \in W : \exists b(R(a,b) \wedge b \in [\![\varphi]\!]_h)\}$, and $[\![\Box\varphi]\!]_h = \{a \in W : \forall b(R(a,b) \to b \in [\![\varphi]\!]_h)\}$. Finally, for a mu-formula φ and $X \in \mathcal{V}$ with only positive free occurrences in φ, we note that the map $f : \wp(W) \to \wp(W)$ given by $f(U) = [\![\varphi]\!]_{h^U_X}$ is monotonic (this can be proved by induction on φ), and define

$$
\begin{aligned}
[\![\mu X \varphi]\!]_h &= \bigcap\{U \subseteq W : [\![\varphi]\!]_{h^U_X} \subseteq U\}, \\
[\![\nu X \varphi]\!]_h &= \bigcup\{U \subseteq W : [\![\varphi]\!]_{h^U_X} \supseteq U\}.
\end{aligned}
$$

By the Knaster–Tarski theorem [21], these are (respectively) the least and greatest fixed points of f. As alternative notation, for a mu-formula φ we write $(\mathcal{F}, h), a \models \varphi$ iff $a \in [\![\varphi]\!]_h$.

Let φ be any modal mu-formula. It can be checked by induction that if $S \subseteq \mathcal{P} \cup \mathcal{V}$ and no $\xi \in S$ occurs free in φ, then $[\![\varphi]\!]_g = [\![\varphi]\!]_h$ for all assignments g, h into the same frame that agree except perhaps on symbols in S. We say that φ is *positive (negative)* if every atom and free fixed point variable in φ occurs under an even (odd) number of negations. Suppose that π is positive and ν negative. It can be checked by induction that π is *monotonic* and ν

is *antimonotonic:* that is, if h, h' are assignments into the same frame and $h(\xi) \subseteq h'(\xi)$ for all $\xi \in \mathcal{P} \cup \mathcal{V}$, then $[\![\pi]\!]_h \subseteq [\![\pi]\!]_{h'}$ and $[\![\nu]\!]_{h'} \subseteq [\![\nu]\!]_h$.

We say that φ is *valid in a frame* $\mathcal{F} = (W, R)$ if $[\![\varphi]\!]_h = W$ for every assignment h into $\mathcal{F}$, and *valid* if it is valid in every frame. We let '$\equiv$' denote logical equivalence: $\varphi \equiv \psi$ iff $\varphi \leftrightarrow \psi$ is valid.

The *dual* operators to $\wedge, \vee, \Box, \Diamond, \mu, \nu$ are $\vee, \wedge, \Diamond, \Box, \nu, \mu$, respectively. As well as the usual $\neg(\varphi \wedge \psi) \equiv \neg\varphi \vee \neg\psi$, $\neg\Box\varphi \equiv \Diamond\neg\varphi$, etc, it can be checked that $\neg\mu X\varphi(X) \equiv \nu X\neg\varphi(\neg X/X)$ and $\neg\nu X\varphi(X) \equiv \mu X\neg\varphi(\neg X/X)$.

2.2 First-order logic plus fixed points (FO+LFP)

We will be very brief here, since first-order logic plus fixed point operators is a well known and well understood system. We refer the reader to [10] for much more information on it. We will use 'FO+LFP' to stand for first-order logic augmented by least and also greatest fixed point operators. We work in the signature with a binary relation symbol R and unary relation symbols P, X for each $p \in \mathcal{P}$ and $X \in \mathcal{V}$. The atomic formulas of FO+LFP are $x = y$, $R(x, y)$, $\top$, $\bot$, $P(x)$, and $X(x)$, for any variables x, y, and $p \in \mathcal{P}$, $X \in \mathcal{V}$. If φ, ψ are formulas then so are $\neg\varphi$, $\varphi \wedge \psi$, $\varphi \vee \psi$, $\forall x\varphi$, and $\exists x\varphi$. If φ is a formula, x a variable, and S a unary relation symbol (arising either from $\mathcal{P}$ or $\mathcal{V}$) all of whose free occurrences in φ are positive, then $[LFP(S, x)\varphi]$ and $[GFP(S, x)\varphi]$ are formulas with the same free first-order variables as φ, but in which S is now bound. The semantics is as usual; in particular, if all free occurrences of S in $\varphi(x, y_1, \ldots, y_n, S)$ are positive, then $M \models [LFP(S, x)\varphi](a, b_1, \ldots, b_n)$ iff a is in the least fixed point of the (monotone) map $f : \wp M \to \wp M$ given by $f(U) = \{c \in M : M \models \varphi(c, b_1, \ldots, b_n, U)\}$. Semantics of $[GFP(S, x)\varphi]$ are defined similarly, using greatest fixed points. Occasionally we will take fixed points of higher-arity relations.

As in the mu-calculus, $\equiv$ will denote the relation of logical equivalence. Any formula positive in P is monotonic in P as well.

2.3 Standard translations

For a first-order variable x, every modal mu-formula $\varphi(p_1, \ldots, p_n, X_1, \ldots, X_m)$ has a *standard translation* $ST_x(\varphi)$: a formula $\varphi'(x, P_1, \ldots, P_n, X_1, \ldots, X_m)$ of FO+LFP defined as follows:

1. $ST_x(p) = P(x)$, $ST_x(X) = X(x)$, $ST_x(\top) = \top$, and $ST_x(\bot) = \bot$,
2. $ST_x(\neg\varphi) = \neg ST_x\varphi$, etc.,
3. $ST_x(\Diamond\varphi) = \exists y(R(x, y) \wedge ST_y(\varphi))$, for some variable $y \neq x$,
4. $ST_x(\Box\varphi) = \forall y(R(x, y) \to ST_y(\varphi))$, for some variable $y \neq x$,
5. $ST_x(\mu X\varphi) = [LFP(X, x)ST_x\varphi]$,
6. $ST_x(\nu X\varphi) = [GFP(X, x)ST_x\varphi]$.

For any frame $\mathcal{F} = (W, R)$, any assignment h into $\mathcal{F}$, any $a \in W$, and any modal mu-formula $\varphi(p_1, \ldots, p_n, X_1, \ldots, X_m)$ with $ST_x\varphi = \varphi'(x, P_1, \ldots, P_n, X_1, \ldots, X_m)$, we have $(\mathcal{F}, h), a \models \varphi$ iff $\mathcal{F} \models \varphi'(a, h(p_1), \ldots, h(p_n), h(X_1), \ldots, h(X_m))$. We remark that if φ is positive in p_i then $ST_x(\varphi)$ is positive in P_i.

3 Sahlqvist's theorem and the mu-calculus

Here we will describe the existing work that led us to the position recorded in this paper.

3.1 Classical Sahlqvist correspondence

Sahlqvist formulas originated in [19]. In spite of (or perhaps because of) their importance in modal logic today, there seems to be no universally agreed modern definition of them. We will adopt the following simple definition.

DEFINITION 3.1 [Sahlqvist formula]

1. Any positive formula is a Sahlqvist formula.
2. Any formula of the form $\neg\Box^n s$ (a negated 'boxed atom') is a Sahlqvist formula, where $n \geq 0$, $\Box^0\varphi = \varphi$, $\Box^{n+1}\varphi = \Box(\Box^n\varphi)$, and s is a propositional atom.
3. If φ, ψ are Sahlqvist formulas then so are $\varphi \vee \psi$ and $\Box\varphi$.

Many commonly arising modal axioms are equivalent to Sahlqvist formulas. To illustrate, the formula $\Box p \to p$ is equivalent to $\neg\Box p \vee p$, which is constructed from the negated boxed atom $\neg\Box p$ (clause 2) and the positive formula p (clause 1) using $\vee$ (clause 3). It is common to include $\varphi \wedge \psi$ in clause 3 above — for example, the definition of Sahlqvist formulas in [7, definition 3.51] boils down to this. We do not allow $\wedge$ in clause 3 for two reasons. First, any formula obtained by adding $\wedge$ to clause 3 is in any case equivalent to a conjunction of Sahlqvist formulas as defined above, because any occurrence of $\wedge$ can be moved up through the $\vee$s and $\Box$s using distributivity. Second, the argument coming up in a moment is simpler without $\wedge$ in clause 3. But when we come to Sahlqvist mu-formulas, we will want to include $\wedge$.

Sahlqvist formulas have two key properties:

Correspondence. For any Sahlqvist formula φ, there is a first-order sentence χ_φ, called the *frame correspondent* of φ, that is true in an arbitrary Kripke frame iff φ is valid in that frame. Moreover, χ_φ can be computed from φ by a simple algorithm.

Completeness. For any Sahlqvist formula φ, the basic modal logic K augmented with φ as an extra axiom is sound and complete for the class of frames defined by χ_φ.

The properties are of course related, and further algebraic properties have been established (e.g., [14]). The celebrated 'Esakia lemma' [11] is used in a key step in the proof of completeness (e.g., [20]). In this paper we are concerned only with correspondence, and we confine our discussion to that topic. There are several proofs of Sahlqvist correspondence in the literature: e.g., [19, 2, 20, 7]. But the idea can be simply explained, as follows. It will be familiar to many readers, but we (briefly) go through the steps because we intend to generalise them later.

Let φ be a Sahlqvist formula and $\mathcal{F} = (W, R)$ a Kripke frame.

Step 1. Assume that φ is *not* valid in $\mathcal{F}$. This says that there is a model $\mathcal{M} = (\mathcal{F}, h)$, for some assignment h of atoms into $\mathcal{F}$, and some world $a \in W$, such that $\mathcal{M}, a \models \neg\varphi$. Now $\neg\varphi$ is plainly equivalent to a formula of the form

$$\sigma(\nu_1, \ldots, \nu_m, \beta_1, \ldots, \beta_n), \tag{1}$$

where $\sigma(p_1, \ldots, p_m, q_1, \ldots, q_n)$ is a formula made from distinct atoms $p_1, \ldots,$ p_m $q_1, \ldots, q_n$ using only $\wedge$ and $\Diamond$ (the duals of the operations in clause 3 of definition 3.1); each of $q_1, \ldots, q_n$ occurs exactly once in σ; $\nu_1, \ldots, \nu_m$ are negative formulas; $\beta_1, \ldots, \beta_n$ are boxed atoms; and (1) is shorthand for the result

$$\sigma(\nu_1/p_1, \ldots, \nu_m/p_m, \beta_1/q_1, \ldots, \beta_n/q_n)$$

of uniformly replacing each atom p_i in σ by ν_i and each q_j by β_j. So φ is not valid in $\mathcal{F}$ iff there are a, h with

$$(\mathcal{F}, h), a \models \sigma(\nu_1, \ldots, \nu_m, \beta_1, \ldots, \beta_n). \tag{2}$$

Step 2. Now we observe the following critical fact. Let x be any first-order variable.

LEMMA 3.2 *The standard translation $ST_x(\sigma(p_1, \ldots, p_m, q_1, \ldots, q_n))$ of σ is equivalent to a formula $\psi(x, P_1, \ldots, P_m, Q_1, \ldots, Q_n)$ of the form*

$$\exists y_1 \ldots y_n \Big(\psi(x, P_1, \ldots, P_m, \bar{y}) \wedge \bigwedge_{1 \le j \le n} Q_j(y_j)\Big), \tag{3}$$

for some first-order formula $\psi(x, P_1, \ldots, P_m, \bar{y})$ positive in each of $P_1, \ldots, P_m$, where $\bar{y} = (y_1, \ldots, y_n)$ is a tuple of distinct variables different from x.

The proof is a simple induction on the structure of σ, and it can be done precisely because (as a result of clause 3 of definition 3.1) σ only involves $\wedge$ and $\Diamond$, and each q_j occurs exactly once in σ. If we allowed $\wedge$ in clause 3, ψ would be more complicated: a disjunction of formulas of the form (3).

With (3) at hand, we see that (2) *literally says* that for some a, h,

(∗) there are $b_1, \ldots, b_n \in W$, standing in a certain relation to a and to each other specified by ψ (formally, by $\mathcal{F} \models \psi(a, [\![\nu_1]\!]_h, \ldots, [\![\nu_m]\!]_h, b_1, \ldots, b_n)$) and such that $(\mathcal{F}, h), b_j \models \beta_j$ for each $1 \le j \le n$.

Step 3. The next critical step is to observe that without loss of generality we can replace h by a *'minimal assignment'* h°, satisfying $h^\circ(s) \subseteq h(s)$ for every atom s occurring in φ. In fact, h° is the assignment where each $h^\circ(s)$ is as small as possible subject to the condition that $(\mathcal{F}, h^\circ), b_j \models \beta_j$ for $1 \leq j \leq n$. The definition of h° is uniform in $b_1, \ldots, b_n$.

To find h°, for each atom s we collect up all the boxed atoms β_j involving s. To illustrate, suppose that there are just two of them: $\beta_3 = \Box^2 s$, and $\beta_7 = \Box^0 s = s$. (So β_1, β_5, etc., are boxed atoms involving other atoms than s.) Then $(*)$ states that $(\mathcal{F}, h), b_3 \models \Box^2 s$ and $(\mathcal{F}, h), b_7 \models s$. This will be preserved if we replace h by an assignment g with $g(s) = \{w \in W : \mathcal{F} \models \exists z(R(b_3, z) \wedge R(z, w))\} \cup \{b_7\}$. This is the 'minimal' assignment satisfying $(\mathcal{F}, g), b_3 \models \Box^2 s$ and $(\mathcal{F}, g), b_7 \models s$. Any assignment g' making β_3 true at b_3 and β_7 at b_7 must plainly satisfy $g(s) \subseteq g'(s)$, and in particular, we have $g(s) \subseteq h(s)$. Let h° be the 'minimal assignment' that assigns the minimal value $g(s)$ to each atom s as just explained. If s does not occur in any β_j then $h^\circ(s) = \emptyset$.

Now $h^\circ(s) \subseteq h(s)$ for all atoms s. Consequently, by antimonotonicity of negative formulas, $[\![\nu_i]\!]_h \subseteq [\![\nu_i]\!]_{h^\circ}$ for $i = 1, \ldots, m$. Since $P_1, \ldots, P_m$ occur only positively in ψ, the truth of $\psi(a, [\![\nu_1]\!]_h, \ldots, [\![\nu_m]\!]_h, b_1, \ldots, b_n)$ in $(*)$ is unaffected by our replacing h by h°.

So if $(*)$ holds for some assignment h, then it holds for h°. Since if $(*)$ holds for h° then it certainly holds for some h, we conclude that φ is not valid in $\mathcal{F}$ iff $(*)$ holds for some a and for h°.

We now make one final observation: it is *automatic* that $(\mathcal{F}, h^\circ), b_j \models \beta_j$ for each $1 \leq j \leq n$, since h° is defined precisely to achieve this. We conclude that φ is not valid in $\mathcal{F}$ iff:

$(**)$ there are $a, b_1, \ldots, b_n \in W$ with $\mathcal{F} \models \psi(a, [\![\nu_1]\!]_{h^\circ}, \ldots, [\![\nu_m]\!]_{h^\circ}, b_1, \ldots, b_n)$, where h° is defined as above.

Step 4. The final critical step is to notice that for each atom s, the value $h^\circ(s)$ is *first-order definable with the parameters* $b_1, \ldots, b_n$. We have

$$h^\circ(s) = \{c \in W : \mathcal{F} \models \delta_s(c, b_1, \ldots, b_n)\},$$

where $\delta_s(x, y_1, \ldots, y_n)$ is a certain first-order formula in the frame language, and one that we can explicitly construct. In the example above, we had $h^\circ(s) = \{c \in W : \mathcal{F} \models \exists z(R(b_3, z) \wedge R(z, c))\} \cup \{b_7\}$ — this is definable as $\{c \in W : \mathcal{F} \models \delta_s(c, b_1, \ldots, b_n)\}$, where

$$\delta_s(x, y_1, \ldots, y_n) = \exists z(R(y_3, z) \wedge R(z, x)) \vee x = y_7.$$

Summing up. In the light of $(**)$ and step 4, we see that φ is not valid in $\mathcal{F}$ iff

$$\mathcal{F} \models \exists x \bar{y}\, \theta(x, \bar{y}), \tag{4}$$

where θ denotes the result of replacing each subformula of $\psi(x, P_1, \ldots, P_m, \bar{y})$ of the form $P_i(t)$ (for some $1 \leq i \leq m$ and some variable t) by: the formula obtained from $ST_t(\nu_i)$ by replacing each subformula $S(v)$ (for an atom s and a variable v) by $\delta_s(v/x, y_1, \ldots, y_n)$ (which is the definition of $h^\circ(s)$). By construction, (4) means exactly the same as (∗∗) and is equivalent to φ's failing to be valid in $\mathcal{F}$. Consequently, the negation $\forall x\bar{y}\neg\theta(x, \bar{y})$ of the first-order sentence in (4) is our desired frame correspondent for φ.

We would like to generalise this argument, eventually to the mu-calculus.

3.2 PIA formulas

In [4], van Benthem showed how to generalise steps 3 and 4 to a wider class of modal formulas than boxed atoms, at the cost of ending up with a frame correspondent not in first-order logic but in FO+LFP: first-order logic plus the least and greatest fixed point operators.

What step 3 needs is the existence of a *minimal assignment* that makes a formula β true at a given world y of a Kripke frame, given that there exists at least one assignment making β true at y. As we saw, if β is a boxed atom $\Box^d s$ then there is indeed a minimal assignment to s, namely, $\{w \in W : \mathcal{F} \models R^d(y, w)\}$, where $R^0(y, w)$ is $y = w$ and $R^{d+1}(y, w)$ is $\exists z(R(y, z) \wedge R^d(z, w))$.

[4] studied *first-order* sentences $\varphi(S)$ (for a unary relation symbol S corresponding to the atom s) that admit such a minimal assignment, in the sense that in any first-order structure M there is a minimal $S \subseteq M$ with $M \models \varphi(S)$. It was shown that a sufficient condition for $\varphi(S)$ to admit a minimal assignment is that it has the *intersection property (IP):* namely, that for any M, index set I, and subsets $S_i \subseteq M$ $(i \in I)$, if $M \models \varphi(S_i)$ for each $i \in I$ then $M \models \varphi(\bigcap_{i \in I} S_i)$. The minimal assignment to S that makes $\varphi(S)$ true is then simply $\bigcap\{S \subseteq M : M \models \varphi(S)\}$. It was also proved that $\varphi(S)$ has IP iff it is equivalent to a sentence of the form

$$\forall y(\psi(S, y) \to S(y)),$$

where $\psi(S, y)$ is positive in S. Such sentences have the form 'positive implies atomic', or for short, *'PIA'*.

This is for first-order logic, and no similar characterisation of the modal version of IP was given. Nonetheless, [4] did exhibit a modal analogue of 'PIA implies IP implies minimal assignment exists'. This arises by considering modal formulas $\varphi(s)$ that we will call *semantically PIA formulas,* whose standard translations $ST_x(\varphi(s))$ are equivalent to PIA formulas of the form

$$\forall y(\psi(S, x, y) \to S(y)), \tag{5}$$

for ψ positive in S. Boxed atoms are examples: $ST_x(\Box^d s) \equiv \forall y(R^d(x, y) \to S(y))$, which is of the required form (5). But there are many more. First, any atom s is a semantically PIA formula, since its standard translation $ST_x(\varphi)$ is $S(x)$ — this is equivalent to $\forall y(y = x \to S(y))$, which is of the form (5). Second, it can be verified that the semantically PIA formulas $\varphi(s)$, for a fixed atom s, are closed under $\wedge$ and $\Box$ (though not under $\vee$). Third, if $\varphi(s)$ is semantically PIA and $\pi(s)$ is positive in s then $\pi(s) \to \varphi(s)$ is also semantically PIA. Since for Sahlqvist purposes we would like a syntactically defined class of semantically PIA formulas, we say that a modal formula $\varphi(s)$ is *syntactically PIA* if it is obtained from atoms by applying $\wedge$, $\Box$, and $\pi(s) \to \cdot$, where $\pi(s)$ is positive in s. Boxed atoms are plainly (very) special cases of syntactically PIA formulas. Any syntactically PIA formula, and indeed any semantically PIA formula, admits a minimal assignment to s as required by step 3 of the correspondence proof in section 3.1.

For step 4, we also need that the minimal assignment is *definable* in first-order logic. The minimal S satisfying (5) need not be first-order definable. However, it *is* definable in FO+LFP. This is because the minimal S satisfying (5) (in a frame $\mathcal{F} = (W, R)$, for a given $x \in W$) is the intersection of *all* S satisfying (5). By the Knaster–Tarski theorem, this intersection is the least fixed point of the monotone map $f_{\psi,x} : \wp(W) \to \wp(W)$ given by $f_{\psi,x}(S) = \{a \in W : \mathcal{F} \models \psi(S, x, a)\}$, for $S \subseteq W$. It is therefore defined by the FO+LFP-formula $[\mathrm{LFP}(S, y)\psi](x, y)$.

The astute reader will have noticed that step 3 also required that we can take the union of the minimal assignments to a given atom s from all of the boxed atoms β_j involving s, obtaining a single (definable) minimal assignment that still satisfies all these β_j. This is true for syntactically PIA formulas, for much the same reason that they are closed under $\wedge$, but properly it is a consideration for the 'clause 3' structure of the Sahlqvist formula.

We conclude that we can allow negated syntactically PIA formulas in clause 2 of definition 3.1, if we do not mind the frame correspondent being in FO+LFP instead of first-order logic.

3.3 PIA mu-calculus formulas

The main contribution of the current paper now begins. As suggested in [4], if we are willing to admit frame correspondents in FO+LFP, why not go further and consider formulas of the *modal mu-calculus,* whose standard translations automatically lie in this language? Let us say that a modal mu-calculus formula $\beta(s)$ is *semantically PIA* if its standard translation $ST_x(\beta)$ is equivalent to a FO+LFP-formula of the form

$$\forall y(\psi(S, x, y) \to S(y)), \tag{6}$$

where ψ is positive in S. There will always be a FO+LFP-definable minimal assignment to s making β true at a world a in a frame $\mathcal{F}$, namely, $\{c : \mathcal{F} \models \psi'(a, c)\}$, where $\psi'(x, y) = [LFP(S, y)\psi](x, y)$.

This definition of PIA formula is semantic. As before, we now have the problem of defining a wide syntactic class of semantically PIA mu-formulas. Starting from an atom s and fixed-point variables, we can close under $\wedge$, $\Box$, and $\pi(s) \to \cdot$ as before, where $\pi(s)$ is now a *modal mu-sentence* positive in s. It turns out that we can also close under the greatest fixed point operator ν. This will be seen in section 4. Any sentence $\varphi(s)$ obtained in this way admits a minimal assignment to s that makes φ true at a world x of a frame; the minimal assignment is definable in FO+LFP. So we could allow the negations of such formulas in clause 2 of definition 3.1.

3.4 Clause 1

In step 3 of the correspondence proof, we noted that the negative formulas kept their truth values when we replaced the original assignment h by the minimal one, h°. All that was needed for this was antimonotonicity, which still holds if we allow *positive mu-calculus formulas* in clause 1 of definition 3.1.

3.5 Clause 3

Sahlqvist formulas were defined as the closure of positive formulas and negated boxed atoms under $\vee, \Box$. We have seen how we can generalise boxed atoms (to PIA mu-formulas) and positive formulas (to positive mu-calculus formulas). Now we would like to generalise the 'clause 3' structure: the closure operations $\vee, \Box$.

All we required of these operations was that, when dualised to $\wedge, \Diamond$, they allow lemma 3.2 to be proved. If we include $\vee$ here as well, a form of the lemma involving a disjunction of formulas of the form (3) can be proved. We would like to add μ, and to leverage this powerful operator we would like to have both $\wedge$ and $\vee$ available. (For example, we can already express $\Diamond(p \wedge q)$ using $\wedge$ and $\Diamond$, so we would like to express its 'transitive closure' version $\Diamond^*(p \wedge q)$, by $\sigma_1(p, q) = \mu X((p \wedge q) \vee \Diamond X)$. This requires $\wedge$ and $\vee$.)

It turns out that a disjunctive form of lemma 3.2 can be proved for any formula $\sigma(p_1, \ldots, p_m, q_1, \ldots, q_n)$ built using only $\vee, \Diamond, \mu$, where the formula ψ in (3) is now in FO+LFP of course.

To allow $\wedge$ as well, we have to make restrictions. For example, the standard translation $ST_x(\sigma_2)$ of the formula $\sigma_2(q_1, q_2) = \mu X(q_1 \vee (q_2 \wedge \Diamond X))$, expressing '$q_2$ until q_1', is not equivalent to a disjunction of formulas of the form $\exists y_1 y_2(\psi(x, y, z) \wedge Q_1(y_1) \wedge Q_2(y_2))$ given in (3). A sufficient restriction is to allow $\sigma \wedge \tau$ only if (i) σ and τ have no atoms from $q_1, \ldots, q_n$ (corresponding to the boxed atoms) in common, and (ii) if either has a free fixed point variable then the other is a sentence not involving $q_1, \ldots, q_n$. This restriction allows σ_1 but not σ_2.

Now lemma 3.2 was only a tool for the correspondence proof. What is the effect of the restrictions on $\wedge$ in σ on actual Sahlqvist formulas? The effect of (i) is nil, since we can meet it by using fresh atoms in τ — this doesn't matter since in (1) we substitute formulas for the atoms of σ. The effect of (ii) is that for $\varphi \vee \psi$ to be a Sahlqvist formula, if one of φ, ψ is not a sentence then the other must be a sentence not involving any negated boxed atoms — i.e., a positive sentence.

The 'reason' why lemma 3.2 can be proved for such formulas σ is that they are *completely additive* in each q_k. Formally, if $\mathcal{F}$ is a frame and h_i ($i \in I \neq \emptyset$) are assignments into $\mathcal{F}$ that agree on all atoms other than q_k, and h is the assignment given by $h(p) = \bigcup_{i \in I} h_i(p)$ for each atom p, then for any world a of $\mathcal{F}$ we have $(\mathcal{F}, h), a \models \sigma$ iff $(\mathcal{F}, h_i), a \models \sigma$ for some $i \in I$. The restrictions on $\wedge$ are to ensure that this holds.

Suppose for example that σ only involves the atom q, and $ST_x(\sigma) = \psi(x, Q)$, say. Let $\psi_0(x)$, $\psi_1(x, y)$ denote the result of replacing each subformula $Q(v)$ of ψ by $\perp$ and $v = y$, respectively. Then by complete additivity,

$$ST_x(\sigma) \equiv \psi_0(x) \vee \exists y(\psi_1(x, y) \wedge Q(y)).$$

This form is close enough to (3) for the correspondence proof to work. If σ involves multiple atoms, the argument can be iterated. So we can replace clause 3 of definition 3.1 by a construction allowing (the duals of) $\vee, \Diamond, \mu$, and the restricted $\wedge$ as just explained.

The trouble-maker is clearly $\wedge$. If $\sigma(p_1, \ldots, p_m, q_1, \ldots, q_n)$ does not involve $\wedge$ then we can prove a stronger form of complete additivity. Passing to the dual operations $\wedge, \Box, \nu$, this becomes a strong form of 'complete multiplicativity' analogous to the intersection property (IP), which we will use to show that $\wedge, \Box, \nu$ and $POS \to \cdot$ can be applied to PIA formulas while preserving the existence of a definable minimal assignment.

3.6 Sahlqvist formulas in the mu-calculus

Let us formalise the position we have arrived at. All formulas below are of the modal mu-calculus.

DEFINITION 3.3 [PIA formulas] Let s be an atom. We define the *syntactically PIA formulas* $\beta(s)$ involving only the atom s, as follows.

1. s is a syntactically PIA formula.
2. Any fixed point variable X is a syntactically PIA formula.
3. If $\beta(s), \gamma(s)$ are syntactically PIA formulas then so are $\beta \wedge \gamma$, $\Box\beta$, and $\nu X \beta$ (for any fixed point variable X).
4. If $\beta(s)$ is a syntactically PIA formula and $\pi(s)$ is any modal mu-sentence positive in s and involving no other atoms than s, then $\pi(s) \to \beta(s)$ is a syntactically PIA formula.

In the end we are only interested in syntactically PIA *sentences.* These may not look of the form 'positive implies atomic', but we will see that their standard translations are equivalent to formulas of this form, so we feel the term 'PIA' is justified.

DEFINITION 3.4 We also define the *compound PIA formulas,* which may involve more than one atom:

1. Any syntactically PIA formula is a compound PIA formula.

2. Any fixed point variable X is a compound PIA formula.
3. If φ, ψ are compound PIA formulas then so are $\varphi \wedge \psi$, $\Box\varphi$, and $\nu X\varphi$ (for any fixed point variable X).

Again, we are only interested in compound PIA *sentences.*

DEFINITION 3.5 [Sahlqvist mu-formula]

1. Any positive sentence is a Sahlqvist mu-formula.
2. Any negated compound PIA sentence is a Sahlqvist mu-formula.
3. Any fixed point variable is a Sahlqvist mu-formula.
4. If φ, ψ are Sahlqvist mu-formulas then so are $\varphi \wedge \psi$, $\Box\varphi$, and $\nu X\varphi$ (for any fixed point variable X).
5. If φ, ψ are Sahlqvist mu-formulas, and if one of them is not a sentence then the other is a positive sentence, then $\varphi \vee \psi$ is a Sahlqvist mu-formula.

A *Sahlqvist mu-sentence* is a Sahlqvist mu-formula that is a sentence.

In summary, a Sahlqvist mu-sentence is any sentence obtained by applying $\wedge, \Box$, and ν to fixed point variables, positive sentences, and negated compound PIA sentences; $\vee$ can also be applied so long as if one of the disjuncts is not a sentence then the other is a positive sentence. In the next section we will prove a correspondence theorem for Sahlqvist mu-sentences.

4 Correspondence theorem

This section contains the formal proofs of the paper. We will prove a correspondence theorem for Sahlqvist mu-sentences (theorem 4.14 below). The initial sections contain preliminaries.

4.1 Skeletons

Our main technical tool will be formulas that we call *skeletons,* because they will support the negative formulas and compound PIA formulas (generalising the boxed atoms) in Sahlqvist formulas, as in (1). (In this role, they are analogous to the universal prefix that is extracted in the 'Sahlqvist–van Benthem algorithm' in [7]. Skeletons allow a richer Sahlqvist syntax, including, for example, negative formulas in antecedents — $\Diamond(\neg p \wedge \Box p) \to \cdots$ is fine.) We will also use them to show that PIA and compound PIA formulas really are semantically PIA.

Recall that $\mathcal{P}$ is the ambient set of atoms, and $\mathcal{V}$ the set of fixed point variables.

DEFINITION 4.1 [$\mathcal{Q}$-skeleton] Let $\mathcal{Q} \subseteq \mathcal{P}$ be arbitrary.

1. Any atomic mu-formula (i.e., an atom, a fixed point variable, $\top$, or $\bot$) is a $\mathcal{Q}$-skeleton.

2. If σ, τ are $\mathcal{Q}$-skeletons then so are $\sigma \vee \tau$, $\Diamond\sigma$, and $\mu X \sigma$ (for any fixed point variable X).

3. If σ is a $\mathcal{Q}$-skeleton and τ is a positive sentence involving no atoms from $\mathcal{Q}$, then $\sigma \wedge \tau$ and $\tau \wedge \sigma$ are $\mathcal{Q}$-skeletons.

REMARK 4.2 Any $\mathcal{Q}$-skeleton is a $\mathcal{Q}'$-skeleton for every $\mathcal{Q}' \subseteq \mathcal{Q}$: increasing $\mathcal{Q}$ strengthens the restrictions on $\mathcal{Q}$-skeletons. However, if σ is a $\mathcal{Q}$-skeleton and $\mathcal{Q}'$ is a set of atoms not occurring in σ, a simple induction shows that σ is a $\mathcal{Q} \cup \mathcal{Q}'$-skeleton.

The main semantic property of skeletons is a form of complete additivity, as we will see in proposition 4.4. Fix a frame $\mathcal{F} = (W, R)$.

DEFINITION 4.3 Let $\mathcal{H}$ be a set of assignments into $\mathcal{F}$.

1. Write $\bigcup \mathcal{H}$ for the assignment g given by $g(\xi) = \bigcup\{h(\xi) : h \in \mathcal{H}\}$ for each atom or fixed point variable ξ.

2. Let $\mathcal{Q} \subseteq \mathcal{P}$ be a set of atoms. We say that $\mathcal{H}$ is *$\mathcal{Q}$-variant* if $h(p) = h'(p)$ for all atoms $p \in \mathcal{P} \setminus Q$ and all $h, h' \in \mathcal{H}$. (Important: there are no restrictions on the values of $h \in \mathcal{H}$ on fixed point variables.)

PROPOSITION 4.4 *Fix $\mathcal{Q} \subseteq \mathcal{P}$. Let σ be a $\mathcal{Q}$-skeleton and $\mathcal{H}$ a non-empty $\mathcal{Q}$-variant set of assignments into $\mathcal{F}$. Then $[\![\sigma]\!]_{\bigcup \mathcal{H}} = \bigcup\{[\![\sigma]\!]_h : h \in \mathcal{H}\}$.*

Proof. We prove the proposition by induction on σ. We write $g = \bigcup \mathcal{H}$. If $\sigma \in \mathcal{P} \cup \mathcal{V}$ then $[\![\sigma]\!]_g = g(\sigma) = \bigcup\{h(\sigma) : h \in \mathcal{H}\} = \bigcup\{[\![\sigma]\!]_h : h \in \mathcal{H}\}$. If $\sigma = \bot$, the result is trivial. If $\sigma = \top$, then because $\mathcal{H} \neq \emptyset$ we have $[\![\top]\!]_g = \bigcup\{[\![\top]\!]_h : h \in \mathcal{H}\}$.

We pass to the inductive steps. First suppose that $\sigma = \tau \vee \xi$, where τ, ξ are $\mathcal{Q}$-skeletons. Then $[\![\sigma]\!]_g = [\![\tau]\!]_g \cup [\![\xi]\!]_g$. By the induction hypothesis, $[\![\tau]\!]_g \cup [\![\xi]\!]_g = \bigcup\{[\![\tau]\!]_h : h \in \mathcal{H}\} \cup \bigcup\{[\![\xi]\!]_h : h \in \mathcal{H}\} = \bigcup\{[\![\tau]\!]_h \cup [\![\xi]\!]_h : h \in \mathcal{H}\} = \bigcup\{[\![\sigma]\!]_h : h \in \mathcal{H}\}$.

Next let $\sigma = \Diamond\tau$ for some $\mathcal{Q}$-skeleton τ. Let $w \in W$. Then $w \in [\![\sigma]\!]_g = [\![\Diamond\tau]\!]_g$ iff there is $v \in [\![\tau]\!]_g$ with $R(w, v)$. Inductively, $[\![\tau]\!]_g = \bigcup\{[\![\tau]\!]_h : h \in \mathcal{H}\}$. So the above holds iff there are $h \in \mathcal{H}$ and $v \in [\![\tau]\!]_h$ with $R(w, v)$. This is iff there is $h \in \mathcal{H}$ with $w \in [\![\Diamond\tau]\!]_h = [\![\sigma]\!]_h$: i.e., iff $w \in \bigcup\{[\![\sigma]\!]_h : h \in \mathcal{H}\}$, as required.

Next suppose that $\sigma = \tau \wedge \xi$ for some $\mathcal{Q}$-skeleton τ and positive[1] sentence ξ involving no atom in $\mathcal{Q}$ (the case $\xi \wedge \tau$ is handled similarly). As $\mathcal{H}$ is $\mathcal{Q}$-variant, for each $h \in \mathcal{H}$, g, h agree on all free symbols in ξ, and so $[\![\xi]\!]_g = [\![\xi]\!]_h$ for each $h \in \mathcal{H}$. Now $[\![\sigma]\!]_g = [\![\tau]\!]_g \cap [\![\xi]\!]_g$. By the induction hypothesis, this is equal to $\bigcup\{[\![\tau]\!]_h : h \in \mathcal{H}\} \cap [\![\xi]\!]_g = \bigcup\{[\![\tau]\!]_h \cap [\![\xi]\!]_g : h \in \mathcal{H}\} = \bigcup\{[\![\tau]\!]_h \cap [\![\xi]\!]_h : h \in \mathcal{H}\} = \bigcup\{[\![\sigma]\!]_h : h \in \mathcal{H}\}$.

[1] This assumption is not used here.

Finally, suppose that $\sigma = \mu X \tau$. By monotonicity it is plain that $[\![\sigma]\!]_g \supseteq [\![\sigma]\!]_h$ for each $h \in \mathcal{H}$, so we have $[\![\sigma]\!]_g \supseteq \bigcup\{[\![\sigma]\!]_h : h \in \mathcal{H}\}$. For the converse, we recall that

$$[\![\sigma]\!]_g = \bigcap\{U \subseteq W : [\![\tau]\!]_{g_X^U} \subseteq U\}$$

and

$$\bigcup_{h \in \mathcal{H}} [\![\sigma]\!]_h = \bigcup_{h \in \mathcal{H}} \bigcap\{U \subseteq W : [\![\tau]\!]_{h_X^U} \subseteq U\}.$$

Let $y \in W$ and suppose that $y \notin \bigcup_{h \in \mathcal{H}} [\![\sigma]\!]_h$. Then for each $h \in \mathcal{H}$ there exists $U_h \subseteq W$ such that $[\![\tau]\!]_{h_X^{U_h}} \subseteq U_h$ and $y \notin U_h$. Let $\mathcal{H}' = \{h_X^{U_h} : h \in \mathcal{H}\}$ and $g' = \bigcup \mathcal{H}'$. Clearly, $\mathcal{H}'$ is also $\mathcal{Q}$-variant. So by the induction hypothesis, we obtain $[\![\tau]\!]_{g'} = \bigcup\{[\![\tau]\!]_{h'} : h' \in \mathcal{H}'\}$. As $[\![\tau]\!]_{h_X^{U_h}} \subseteq U_h$ for each $h \in \mathcal{H}$, we have $\bigcup\{[\![\tau]\!]_{h'} : h' \in \mathcal{H}'\} \subseteq \bigcup_{h \in \mathcal{H}} U_h = V$, say. But plainly, $g' = g_X^V$. Thus, we obtained that $[\![\tau]\!]_{g_X^V} \subseteq V$. Now $y \notin V$, as $y \notin U_h$ for each $h \in \mathcal{H}$. Thus, $y \notin \bigcap\{U \subseteq W : [\![\tau]\!]_{g_X^U} \subseteq U\} = [\![\sigma]\!]_g$. □

A related theorem was proved using games in [13, proposition 5.5.4]. We will see that proposition 4.4 has consequences for standard translations of $\mathcal{Q}$-skeletons.

NOTATION 4.5 We will frequently be working with skeletons of the form

$$\sigma(p_1, \ldots, p_m, q_1, \ldots, q_n),$$

and the following notation will be repeatedly useful. We will write $N = \{1, \ldots, n\}$. Fix distinct first-order variables $x, y_1, \ldots, y_n$. For $U \subseteq V \subseteq N$, we will write

$$\sigma_{U/V}(x, y_i, P_1, \ldots, P_m, Q_j \ : \ i \in U, \ j \in N \setminus V) \tag{7}$$

for the FO+LFP-formula obtained from $ST_x(\sigma)$ by replacing every atomic subformula $Q_k(v)$ (where $k \in V$ and v is a variable) by the formula

$$\begin{cases} v = y_k, & \text{if } k \in U, \\ \bot, & \text{otherwise.} \end{cases}$$

Note that $\sigma_{U/V}$ is a FO+LFP-formula, not a mu-formula.

COROLLARY 4.6 *Let $\mathcal{Q} = \{q_1, \ldots, q_n\}$ and let $\sigma(p_1, \ldots, p_m, q_1, \ldots, q_n)$ be a $\mathcal{Q}$-skeleton sentence. Then $ST_x(\sigma)$ is logically equivalent to*

$$\sigma^* = \sigma_{\emptyset/N}(x, P_1, \ldots, P_m) \vee \bigvee_{1 \le k \le n} \exists y_k \big(\sigma_{\{k\}/N}(x, y_k, P_1, \ldots, P_m) \wedge Q_k(y_k)\big).$$

Proof. Let $\mathcal{F} = (W, R)$ be a frame, and take any assignment g into $\mathcal{F}$, and $a \in W$. It is enough to show that

$$a \in [\![\sigma]\!]_g \iff \mathcal{F} \models \sigma^*(a, g(p_1), \ldots, g(p_m), g(q_1), \ldots, g(q_n)). \tag{8}$$

Let $\mathcal{H}$ be the set of all assignments h into $\mathcal{F}$ such that for some $k \in N$:

- $h(q_k) \subseteq g(q_k)$ and $|h(q_k)| \leq 1$,
- $h(q_l) = \emptyset$ for $l \in N \setminus \{k\}$,
- $h(\xi) = g(\xi)$ for every $\xi \in (\mathcal{P} \cup \mathcal{V}) \setminus \mathcal{Q}$.

Note that $\mathcal{H} \neq \emptyset$, $\mathcal{H}$ is $\mathcal{Q}$-variant, and $\bigcup \mathcal{H} = g$. Now we prove (8). The right-hand side holds iff $\mathcal{F} \models \sigma_{\emptyset/N}(a, g(p_1), \dots, g(p_m))$ or there are $k \in N$ and $b \in g(q_k)$ with $\mathcal{F} \models \sigma_{\{k\}/N}(a, b, g(p_1), \dots, g(p_m))$. By definition of $\sigma_{U/V}$ and $\mathcal{H}$, this is iff $a \in [\![\sigma]\!]_h$ for some $h \in \mathcal{H}$. By proposition 4.4, this is iff $a \in [\![\sigma]\!]_{\bigcup \mathcal{H}} = [\![\sigma]\!]_g$, as required. □

Corollary 4.6 will be useful for PIA formulas, but to rewrite Sahlqvist formulas as we did in (1), we need to extend it to formulas that may not be $\{q_1, \dots, q_n\}$-skeletons, but are only $\{q_i\}$-skeletons for $i = 1, \dots, n$. Because of this weaker assumption, we have to settle for a more complicated conclusion, but the family resemblance should be clear.

COROLLARY 4.7 *Suppose that $\sigma(p_1, \dots, p_m, q_1, \dots, q_n)$ is a $\{q_i\}$-skeleton for each $i = 1, \dots, n$. Then $ST_x(\sigma)$ is logically equivalent to*

$$\sigma^N = \exists y_1 \dots y_n \bigvee_{U \subseteq N} \Big(\sigma_{U/N}(x, y_1, \dots, y_n, P_1, \dots, P_m) \wedge \bigwedge_{k \in U} Q_k(y_k)\Big).$$

We remark that if σ is normal in q_k — that is, $\sigma(p_1, \dots, p_m, q_1, \dots, q_{k-1}, \bot/q_k, q_{k+1}, \dots, q_n) \equiv \bot$ — then all disjuncts with $k \notin U$ are equivalent to $\bot$ and can be deleted.

Proof. The proof is by induction on n. The case $n = 0$ is vacuously true, since then, $\sigma^N = \sigma_{\emptyset/\emptyset} = ST_x(\sigma)$. Let $n > 0$ and assume the result for $n-1$. Treating Q_n as a P and applying the inductive hypothesis to the atoms $q_1, \dots, q_{n-1}$, with $N' = \{1, \dots, n-1\}$, shows that $ST_x(\sigma)$ is equivalent to

$$\sigma^{N'} = \exists y_1 \dots y_{n-1} \bigvee_{U \subseteq N'} \Big(\sigma_{U/N'}(x, y_1, \dots, y_{n-1}, \bar{P}, Q_n) \wedge \bigwedge_{k \in U} Q_k(y_k)\Big), \quad (9)$$

where we write $\bar{P}$ for $(P_1, \dots, P_m)$. As σ is a $\{q_n\}$-skeleton, corollary 4.6 tell us that $ST_x(\sigma)$ is also equivalent to

$$\sigma_{\emptyset/\{n\}}(x, \bar{P}, Q_1, \dots, Q_{n-1}) \vee \exists y_n \big(\sigma_{\{n\}/\{n\}}(x, y_n, \bar{P}, Q_1, \dots, Q_{n-1}) \wedge Q_n(y_n)\big).$$

Using (9) and the definitions of $\sigma_{\emptyset/\{n\}}$ and $\sigma_{\{n\}/\{n\}}$, the first disjunct of this is equivalent to

$$\exists y_1 \dots y_{n-1} \bigvee_{\substack{U \subseteq N \\ n \notin U}} \Big(\sigma_{U/N}(x, y_1, \dots, y_{n-1}, \bar{P}) \wedge \bigwedge_{k \in U} Q_k(y_k)\Big),$$

and the second to

$$\exists y_1 \dots y_n \bigvee_{\substack{U \subseteq N \\ n \in U}} \Big(\sigma_{U/N}(x, y_1, \dots, y_n, \bar{P}) \wedge \bigwedge_{k \in U} Q_k(y_k)\Big).$$

$ST_x(\sigma)$ is equivalent to the disjunction of these, and so to σ^N, which completes the induction. □

4.2 Skeletons and PIA formulas

In this section we will prove that any syntactically PIA sentence has a standard translation equivalent to a 'genuine' PIA (positive implies atomic) formula of FO+LFP, and so is semantically PIA. For compound PIA sentences, we will get a conjunction of FO+LFP PIA formulas, one for each atom.

DEFINITION 4.8 Let $\mathcal{Q} \subseteq \mathcal{P}$ and let σ be a $\mathcal{Q}-$skeleton.

1. σ is said to be *normal* if the formula obtained by replacing every free occurrence of every $\xi \in \mathcal{Q} \cup \mathcal{V}$ in σ by $\bot$ is logically equivalent to $\bot$. ($\mathcal{Q}$ is understood tacitly here. Atoms in $\mathcal{P} \setminus \mathcal{Q}$ are not altered in σ.)
2. We write $\sigma^{\mathcal{Q}}$ (the 'dual' of σ) for the formula obtained from $\neg\sigma$ by replacing each free occurrence of each $\xi \in \mathcal{Q} \cup \mathcal{V}$ by $\neg\xi$. Atoms in $\mathcal{P} \setminus \mathcal{Q}$ are unchanged.

The following is as we would expect when taking duals.

LEMMA 4.9 *Let $\sigma, \sigma_1, \sigma_2$ be $\mathcal{Q}$-skeletons. Then*

1. $(\sigma_1 \vee \sigma_2)^{\mathcal{Q}} \equiv \sigma_1^{\mathcal{Q}} \wedge \sigma_2^{\mathcal{Q}}$,
2. $(\Diamond\sigma)^{\mathcal{Q}} \equiv \Box\sigma^{\mathcal{Q}}$,
3. $(\mu X \sigma)^{\mathcal{Q}} \equiv \nu X \sigma^{\mathcal{Q}}$.

Proof. We prove only the last case. Let $\sigma(\bar{p}, \bar{q}, X, \bar{Y})$ be given, where $\bar{p}$ are atoms not in $\mathcal{Q}$, $\bar{q}$ are atoms in $\mathcal{Q}$, and $X, \bar{Y}$ are fixed point variables. Then in the obvious notation, $(\mu X \sigma)^{\mathcal{Q}} = \neg\mu X \sigma(\bar{p}, \neg\bar{q}, X, \neg\bar{Y}) \equiv \nu X \neg\sigma(\bar{p}, \neg\bar{q}, \neg X, \neg\bar{Y}) = \nu X \sigma^{\mathcal{Q}}$. □

This gives us the following alternative view of syntactically PIA formulas. In the lemma, formulas may have free fixed point variables but we do not display them.

LEMMA 4.10 *Let $\beta(s)$ be a syntactically PIA formula. Fix an atom $q \neq s$. Then $\beta(s) \equiv \sigma^{\{q\}}(s/q, s)$, for some normal $\{q\}$-skeleton $\sigma(q, s)$.*

Proof. By induction on β. We have $s = \sigma^{\{q\}}(s/q)$ where $\sigma = q$ (a normal $\{q\}$-skeleton). For a fixed point variable X, $X = \sigma^{\{q\}}(s/q)$ where $\sigma = X$ (again, X is a normal $\{q\}$-skeleton). Suppose that $\beta_1(s) \equiv \sigma_1^{\{q\}}(s/q, s)$ and $\beta_2(s) \equiv \sigma_2^{\{q\}}(s/q, s)$, for normal $\{q\}$-skeletons $\sigma_1(q, s), \sigma_2(q, s)$.

- Let $\sigma(q, s) = \sigma_1(q, s) \vee \sigma_2(q, s)$ — plainly a normal $\{q\}$-skeleton. By lemma 4.9, $\beta_1 \wedge \beta_2 \equiv \sigma_1^{\{q\}}(s/q, s) \wedge \sigma_2^{\{q\}}(s/q, s) \equiv (\sigma_1 \vee \sigma_2)^{\{q\}}(s/q, s) = \sigma^{\{q\}}(s/q, s)$.
- By lemma 4.9, $\Box\beta_1 \equiv \Box\sigma_1^{\{q\}}(s/q, s) \equiv (\Diamond\sigma_1)^{\{q\}}(s/q, s)$, and $\Diamond\sigma_1$ is normal.

- Let $\sigma(q,s)$ be the $\{q\}$-skeleton $\mu X\sigma_1(q,s)$. It is clearly normal. By lemma 4.9, $\nu X\beta_1 \equiv \nu X\sigma_1^{\{q\}}(s/q,s) \equiv \sigma^{\{q\}}(s/q,s)$.
- Finally, suppose that $\pi(s)$ is any sentence positive in s and involving no other atoms than s. Then $\sigma(q,s) = \pi(s) \wedge \sigma_1(q,s)$ is a normal $\{q\}$-skeleton, and $\sigma^{\{q\}}(q,s) \equiv \neg(\pi(s) \wedge \neg\sigma_1^{\{q\}}(q,s)) \equiv \pi(s) \to \sigma_1^{\{q\}}(q,s)$. So $\pi(s) \to \beta_1(s) \equiv \sigma^{\{q\}}(s/q,s)$.

This completes the induction and the proof. □

LEMMA 4.11 *Let s be an atom and $\beta(s)$ a syntactically PIA sentence. Then $ST_x(\beta(s))$ is equivalent to a 'PIA' formula of FO+LFP of the form*

$$\forall y(\xi(x,y,S) \to S(y)),$$

where $\xi(x,y,S)$ is positive in S.

Proof. By lemma 4.10 we have $\beta(s) \equiv \sigma^{\{q\}}(s/q,s)$ for some normal $\{q\}$-skeleton sentence $\sigma(q,s)$. By corollary 4.6, $ST_x(\sigma(q,s)) \equiv \exists y(\xi(x,y,S) \wedge Q(y))$, where $\xi(x,y,S)$ is obtained from $ST_x(\sigma)$ by replacing every subformula $Q(t)$ by $t = y$. (By normality, the disjunct $\sigma_{\emptyset/N}$ in the corollary is equivalent to $\bot$ and we can dispense with it.) So

$$\begin{aligned} ST_x(\beta(s)) &\equiv ST_x(\sigma^{\{q\}}(s/q,s)) \\ &\equiv \neg\exists y(\xi(x,y,S) \wedge \neg S(y)) \equiv \forall y(\xi(x,y,S) \to S(y)). \end{aligned}$$

By definition, $\sigma(q,s)$ is positive in s, so ξ is positive in S. □

We now extend this to compound PIA sentences.

PROPOSITION 4.12 *Let $\varphi(s_1,\ldots,s_m)$ be a compound PIA sentence. Then $ST_x(\varphi)$ is equivalent to a formula of the form*

$$\bigwedge_{1\le k\le m} \forall y(\psi_k(x,y,S_k) \to S_k(y)),$$

where each $\psi_k(x,y,S_k)$ is a $FO+LFP$-formula positive in S_k.

Proof. Much as in lemma 4.10, it can be shown by induction on φ that

$$\varphi(s_1,\ldots,s_m) = \sigma^{\mathcal{Q}}(\beta_1(p_1)/q_1,\ldots,\beta_n(p_n)/q_n)$$

for some n, where $q_1,\ldots,q_n$ are distinct atoms, $\mathcal{Q} = \{q_1,\ldots,q_n\}$, $\sigma(q_1,\ldots,q_n)$ is a normal $\mathcal{Q}$-skeleton sentence, $\beta_1(p_1),\ldots,\beta_n(p_n)$ are syntactically PIA formulas, and $p_1,\ldots,p_n \in \{s_1,\ldots,s_m\}$ are not necessarily distinct.

For each $l = 1,\ldots,n$, let $\chi_l(x,z) = \sigma_{\{l\}/N}(x,z/y_l)$. It follows from corollary 4.6 that

$$ST_x(\sigma^{\mathcal{Q}}(q_1,\ldots,q_n)) \equiv \bigwedge_{1\le l\le n} \forall z\big(\chi_l(x,z) \to Q_l(z)\big),$$

because by normality, the disjunct $\sigma_{\emptyset/N}$ in the corollary is equivalent to $\perp$ and can be omitted. By lemma 4.11, we also have

$$ST_z(\beta_l(p_l)) \equiv \forall y(\xi_l(z, y, P_l) \to P_l(y)),$$

for some FO+LFP-formula $\xi_l(z, y, P_l)$ positive in P_l. For $k \in \{1, \ldots, m\}$, let $L(k) = \{l : 1 \leq l \leq n,\ p_l = s_k\}$. Then

$$\begin{aligned}
ST_x(\varphi) &\equiv ST_x(\sigma^{\mathcal{Q}}(\beta_1(p_1)/q_1, \ldots, \beta_n(p_n)/q_n)) \\
&\equiv \bigwedge_{1 \leq l \leq n} \forall z(\chi_l(x, z) \to ST_z(\beta_l(p_l))) \\
&= \bigwedge_{1 \leq l \leq n} \forall z\big(\chi_l(x, z) \to \forall y[\xi_l(z, y, P_l) \to P_l(y)]\big) \\
&\equiv \bigwedge_{1 \leq l \leq n} \forall y\big(\exists z(\chi_l(x, z) \wedge \xi_l(z, y, P_l)) \to P_l(y)\big) \\
&\equiv \bigwedge_{1 \leq k \leq m} \bigwedge_{l \in L(k)} \forall y\big(\exists z(\chi_l(x, z) \wedge \xi_l(z, y, S_k)) \to S_k(y)\big) \\
&\equiv \bigwedge_{1 \leq k \leq m} \forall y\Big(\underbrace{\Big[\bigvee_{l \in L(k)} \exists z(\chi_l(x, z) \wedge \xi_l(z, y, S_k))\Big]}_{\psi_k(x,y,S_k)} \to S_k(y)\Big),
\end{aligned}$$

as required. Clearly, the indicated $\psi_k(x, y, S_k)$ is positive in S_k. (If $L(k) = \emptyset$ then $\psi_k \equiv \perp$.) $\square$

We conclude from proposition 4.12 that the standard translation of a compound PIA sentence $\varphi(s_1, \ldots, s_m)$ is equivalent to a conjunction of FO+LFP-formulas in PIA form, one for each atom $s_1, \ldots, s_m$.

4.3 Skeletons and Sahlqvist formulas

The definition of Sahlqvist formula is chosen so that we can view Sahlqvist formulas in terms of skeletons, by the following analogue of lemma 4.10.

LEMMA 4.13 *Let φ be a Sahlqvist formula. Then there are a formula $\sigma(p_1, \ldots, p_m, q_1, \ldots, q_n, X_1, \ldots, X_t)$ that is a $\{q_i\}$-skeleton for each $i = 1, \ldots, n$, negative sentences $\nu_1, \ldots, \nu_m$, and compound PIA sentences $\beta_1, \ldots, \beta_n$ (not necessarily distinct), such that*

$$\varphi \equiv \neg\sigma(\nu_1/p_1, \ldots, \nu_m/p_m, \beta_1/q_1, \ldots, \beta_n/q_n, \neg X_1/X_1, \ldots, \neg X_t/X_t). \qquad (10)$$

Proof. By induction on φ. If φ is a positive sentence then $\varphi \equiv \neg\sigma(\neg\varphi/p)$ where $\sigma = p$. If φ is a negated compound PIA sentence $\neg\beta$ then $\varphi \equiv \neg\sigma(\beta/q)$ where $\sigma = q$. If φ is a fixed point variable X, then $\varphi \equiv \neg\sigma(\neg X/X)$ for $\sigma = X$. Assume

(10); then (10) holds with φ replaced by $\Box\varphi$ and σ by $\Diamond\sigma$. Also, taking νX_1 as an example,

$$\begin{array}{ll} & \nu X_1\varphi \\ \equiv & \nu X_1 \neg\sigma(\nu_1/p_1,\ldots,\nu_m/p_m,\beta_1/q_1,\ldots,\beta_n/q_n,\neg X_1/X_1,\ldots,\neg X_t/X_t) \\ \equiv & \neg\mu X_1\sigma(\nu_1/p_1,\ldots,\beta_n/q_n,X_1,\neg X_2/X_2,\ldots,\neg X_t/X_t), \end{array}$$

which is of the form (10).

Suppose in the obvious notation that

$$\begin{array}{rcl} \varphi & \equiv & \neg\sigma(\bar\nu/\bar p,\bar\beta/\bar q,\neg\bar X/\bar X), \\ \varphi' & \equiv & \neg\sigma'(\bar\nu'/\bar p',\bar\beta'/\bar q',\neg\bar X'/\bar X'), \end{array}$$

where $\sigma(\bar p,\bar q,\bar X)$ is a $\{q\}$-skeleton for every q in $\bar q$, and $\sigma'(\bar p',\bar q',\bar X')$ is a $\{q'\}$-skeleton for every q' in $\bar q'$. We can suppose without loss of generality that no atom in $\bar q$ occurs in σ' and no atom in $\bar q'$ occurs in σ. By remark 4.2, σ, σ', and hence $\sigma\vee\sigma'$ are $\{q\}$-skeletons and $\{q'\}$-skeletons for every q in $\bar q$ and q' in $\bar q'$, and clearly, $\varphi\wedge\varphi'\equiv\neg(\sigma\vee\sigma')(\bar\nu/\bar p,\bar\nu'/\bar p',\bar\beta/\bar q,\bar\beta'/\bar q',\neg\bar X/\bar X,\neg\bar X'/\bar X')$ as required. This covers the case $\varphi\wedge\varphi'$.

Now suppose that $\varphi\vee\varphi'$ is a Sahlqvist formula. Certainly, $\varphi\vee\varphi'\equiv\neg(\sigma\wedge\sigma')(\bar\nu/\bar p,\bar\nu'/\bar p',\bar\beta/\bar q,\bar\beta'/\bar q',\neg\bar X/\bar X,\neg\bar X'/\bar X')$. But we need to check that $(\sigma\wedge\sigma')(\bar p\bar p',\bar q\bar q',\bar X\bar X')$ is an $\{x\}$-skeleton for each atom x in $\bar q\bar q'$.

If φ,φ' are both sentences, then so are σ,σ'. For each atom q in $\bar q$ (resp., q' in $\bar q'$), it is plain that σ' (resp. σ) is a positive sentence not involving it. So $\sigma\wedge\sigma'$ is an $\{x\}$-skeleton for each x in $\bar q\bar q'$.

Suppose instead that φ is not a sentence (the other case is similar). Then (see definition 3.5) φ' is a positive sentence and consequently does not involve any negated compound PIA sentences. So we may assume that $\bar q'$ is empty. Now for each q in $\bar q$, σ' is a positive sentence not involving q, so $(\sigma\wedge\sigma')(\bar p\bar p',\bar q,\bar X\bar X')$ is a $\{q\}$-skeleton. This completes the proof. $\Box$

4.4 Sahlqvist correspondence for mu-calculus

We are now ready to prove our main result.

THEOREM 4.14 *Any Sahlqvist mu-sentence $\varphi(s_1,\ldots,s_t)$ has a FO+LFP frame correspondent — a sentence χ_φ of FO+LFP with the property that for any frame $\mathcal{F}$, we have $\mathcal{F}\models\chi_\varphi$ iff φ is valid in $\mathcal{F}$. The correspondent χ_φ can be computed from φ by a simple[2] algorithm.*

Proof. We follow the same steps as in our original account in section 3. Let $\mathcal{F}=(W,R)$ be any Kripke frame.

[2]Well, fairly simple.

Step 1. Assume that φ is not valid in $\mathcal{F}$. This is the case iff there are an assignment h into $\mathcal{F}$ and $a \in W$ with $(\mathcal{F}, h), a \models \neg\varphi$. Now by lemma 4.13,

$$\neg\varphi \equiv \sigma(\nu_1/p_1, \ldots, \nu_m/p_m,\ \beta_1/q_1, \ldots, \beta_n/q_n),$$

where $\sigma(p_1, \ldots, p_m, q_1, \ldots, q_n)$ is a sentence that is a $\{q_i\}$-skeleton for each $1 \leq i \leq n$, and $\nu_1, \ldots, \nu_m$ are negative sentences and $\beta_1, \ldots, \beta_n$ compound PIA sentences written with the atoms $s_1, \ldots, s_t$. So

$$(\mathcal{F}, h), a \models \sigma(\nu_1/p_1, \ldots, \nu_m/p_m, \beta_1/q_1, \ldots, \beta_n/q_n). \tag{11}$$

Step 2. By corollary 4.7, $ST_x(\sigma(p_1, \ldots, p_m, q_1, \ldots, q_n))$ is logically equivalent to

$$\exists y_1 \ldots y_n \bigvee_{U \subseteq N} \Big(\sigma_{U/N}(x, P_1, \ldots, P_m, y_1, \ldots, y_n) \wedge \bigwedge_{k \in U} Q_k(y_k)\Big).$$

So by (11), we see that φ is not valid in $\mathcal{F}$ iff there are an assignment h into $\mathcal{F}$, $a, b_1, \ldots, b_n \in W$, and $U \subseteq N$ with

$$\mathcal{F} \models \sigma_{U/N}(a, [\![\nu_1]\!]_h, \ldots, [\![\nu_m]\!]_h, b_1, \ldots, b_n) \text{ and } \bigwedge_{k \in U} \big(b_k \in [\![\beta_k]\!]_h\big). \tag{12}$$

Step 3. We now plan to replace h by a 'minimal' assignment h°, preserving (12). This assignment will depend uniformly on $b_1, \ldots, b_n$, as before, and it will also depend on U.

Each $\beta_k(s_1, \ldots, s_t)$ $(1 \leq k \leq n)$ is compound PIA, so by proposition 4.12 its standard translation $ST_{y_k}(\beta_k)$ is equivalent to a FO+LFP-formula of the form

$$\bigwedge_{1 \leq l \leq t} \forall y(\psi_l^k(y_k, y, S_l) \to S_l(y)), \tag{13}$$

where each ψ_l^k is positive in S_l. So the last part of (12) says precisely that

$$(\mathcal{F}, h(s_l)) \models \forall y(\psi_l^k(b_k, y, S_l) \to S_l(y)) \tag{14}$$

for each $l = 1, \ldots, t$ and each $k \in U$. This condition is plainly equivalent to $(\mathcal{F}, h(s_l)) \models \bigwedge_{k \in U} \forall y(\psi_l^k(b_k, y, S_l) \to S_l(y))$ for each $1 \leq l \leq t$, and so to:

$$(\mathcal{F}, h(s)) \models \forall y(\rho_U^s(y, b_1, \ldots, b_n, S) \to S(y)) \tag{15}$$

for each atom $s \in \mathcal{P}$, where

$$\rho_U^s(y, y_1, \ldots, y_n, S) = \bigvee \big\{\psi_l^k(y_k, y, S) : k \in U,\ 1 \leq l \leq t,\ s = s_l\big\}. \tag{16}$$

If $s \notin \{s_1, \ldots, s_t\}$ then $\rho_U^s \equiv \bot$.

Now each ρ_U^s is positive in S. So (15) is in PIA form, and a minimal assignment to each s exists. Call this assignment h°. As we said, it depends on $b_1, \ldots, b_n$, and U. For $s \notin \{s_1, \ldots, s_t\}$ we have $h^\circ(s) = \emptyset$.

If we replace h by h° in (12), the condition $b_k \in [\![\beta_k]\!]_{h^\circ}$ for each $k \in U$ is automatic — h° is by definition the minimal assignment that ensures this. Moreover, $h^\circ(s) \subseteq h(s)$ for all atoms s. By antimonotonicity, $[\![\nu_l]\!]_h \subseteq [\![\nu_l]\!]_{h^\circ}$ for each $1 \leq l \leq m$. As $\sigma_{U/N}$ is positive in $P_1, \ldots, P_m$, we have $\mathcal{F} \models \sigma_{U/N}(a, [\![\nu_1]\!]_{h^\circ}, \ldots, [\![\nu_m]\!]_{h^\circ}, b_1, \ldots, b_n)$.

We conclude from (12) that φ is not valid in $\mathcal{F}$ iff there are $a, b_1, \ldots, b_n \in W$, and $U \subseteq N$ such that with the above h°,

$$\mathcal{F} \models \sigma_{U/N}(a, [\![\nu_1]\!]_{h^\circ}, \ldots, [\![\nu_m]\!]_{h^\circ}, b_1, \ldots, b_n). \tag{17}$$

Step 4. Moreover, for each atom s, the minimal assignment $h^\circ(s)$ that satisfies (15) is definable in FO+LFP: it is given by the set of all c in $\mathcal{F}$ that satisfy the FO+LFP-formula $\eta_U^s(c, b_1, \ldots, b_n)$, where

$$\eta_U^s(y, y_1, \ldots, y_n) = [LFP(S, y)\rho_U^s(y, y_1, \ldots, y_n, S)](y, y_1, \ldots, y_n). \tag{18}$$

This is a well formed FO+LFP-formula, since ρ_U^s is positive in S.

Summing up. Let $\omega_U(x, y_1, \ldots, y_n)$ be the formula obtained as follows. We take $\sigma_{U/N}(x, P_1, \ldots, P_m, y_1, \ldots, y_n)$ and replace each atomic subformula $P_l(t)$ ($1 \leq l \leq m$, t a variable) by the formula obtained from $ST_t(\nu_l)$ by replacing each atomic subformula $S(z)$ (for some atom s and variable z) by $\eta_U^s(z/y, y_1, \ldots, y_n)$ from (18) (the parts of η_U^s are given in (16) and (13)). Then (17) is equivalent to

$$\mathcal{F} \models \omega_U(a, b_1, \ldots, b_n),$$

and φ is not valid in $\mathcal{F}$ iff there are $a, b_1, \ldots, b_n \in W$ and $U \subseteq N$ such that this holds. We conclude that the original statement that φ is not valid in $\mathcal{F}$ is equivalent to

$$\mathcal{F} \models \exists x y_1 \ldots y_n \bigvee_{U \subseteq N} \omega_U(x, y_1, \ldots, y_n).$$

Thus we obtain our correspondent χ_φ as the negation of this. □

5 Examples

We will now give a few examples concerning frame correspondents. We explained the algorithm that constructs the correspondents in full detail in section 4, and in spirit in section 3. In the examples, we will take an informal approach true to the spirit of the algorithm. The reader may like to apply the algorithm to the examples following the precise steps of the preceding section.

5.1 Löb's formula, $\Box(\Box p \to p) \to \Box p$

We simply state the correspondence: $\mathcal{F}, x \models \Box(\Box p \to p) \to \Box p$ iff (1) R is transitive from x, and (2) R is conversely well-founded at x. Note that the antecedent $\Box(\Box p \to p)$ is PIA, and we can see that its minimal valuation stated as a fixed-point by our general procedure amounts to the set $\{y : \forall z(R^*yz \to Rxz) \wedge \text{no infinite sequence starts from } y\}$. Substituting this into the consequent gives the above frame-equivalent.

Now that we have PIA forms, we can go back to earlier work on non-first-order correspondence and see what was going on. For instance, the modal axiom $(\Diamond p \wedge \Box(p \to \Box p)) \to p$ discussed in [3] has a PIA conjunct $\Box(p \to \Box p)$ in its antecedent. Its corresponding frame property is easily determined.

5.2 Axioms of propositional dynamic logic (PDL)

Consider the axioms of PDL, treating complex program expressions as new relation symbols. For instance, the characteristic axiom for composition, $[a;b]p \leftrightarrow [a][b]p$, may be viewed as $[c]p \leftrightarrow [a][b]p$. This axiom consists of two implications that are clearly Sahlqvist forms. Computing their frame equivalents via the usual algorithm yields $R_c = R_a \circ R_b$, where $\circ$ is composition of binary relations. Now consider the two axioms for Kleene star: (i) $[a^*]p \to p \wedge [a][a^*]p$, (ii) $p \wedge [a^*](p \to [a]p) \to [a^*]p$. These may be viewed as (i) $[b]p \to p \wedge [a][b]p$, (ii) $p \wedge [b](p \to [a]p) \to [b]p$. Of these, the first is standard first-order Sahlqvist. What it says is that $Id \subseteq R_b$ and $R_a \circ R_b \subseteq R_b$. The second principle has an antecedent that is PIA by the rules of our syntax. Suppressing a precise calculation here, in conjunction with the preceding two inclusions it says that the relation R_b is equal to the reflexive-transitive closure R_a^*.

5.3 $\varphi_1 = \Box^+ s \to s$

Here, $\Box^+ s$ abbreviates $\nu X \Box(s \wedge X)$, which defines the 'transitive closure' of $\Box$. We could treat φ_1 as a classical Sahlqvist formula in a modal signature with the box $\Box^+$ with accessibility relation R^+, calculate its correspondent by the classical method (§3.1) as $\forall x R^+(x,x)$, and then replace $R^+(x,x)$ by its FO+LFP definition $[LFP(S,x,y) \,.\, R(x,y) \vee \exists z(R(x,z) \wedge R^+(z,y))](x,x)$. Note that this requires a binary LFP operation.

Alternatively, we can use our algorithm. Written out in the mu-calculus, φ_1 is $\nu X \Box(s \wedge X) \to s$. It is valid in a frame $\mathcal{F}$ at a world x iff $(\mathcal{F}, h), x \models \nu X \Box(s \wedge X) \to s$ for all assignments h into $\mathcal{F}$.

Let $\mathcal{H}$ be the set of assignments h (into $\mathcal{F}$) with $(\mathcal{F}, h), x \models \nu X \Box(s \wedge X)$. We will show that there is a 'smallest' h° (with minimum $h(s)$) in $\mathcal{H}$. Then φ_1 is valid in $\mathcal{F}$ iff $(\mathcal{F}, h), x \models s$ for all $h \in \mathcal{H}$. Since s is positive, this holds iff $(\mathcal{F}, h^\circ), x \models s$.

We calculate h° using PIA methods. Clearly, $\nu X \Box(s \wedge X) \equiv [\neg \mu X \Diamond(s \vee X)](\neg s/s)$. As $\mu X \Diamond(s \vee X)$ is normal and completely additive in s, its standard translation ST_x at x is equivalent to $\exists v(\lambda(v,x) \wedge S(v))$, where

$$\lambda(v,x) = [LFP(X,x) \,.\, \exists y(R(x,y) \wedge (y = v \vee X(y))](v,x).$$

So

$$ST_x(\nu X \Box(s \wedge X)) \equiv \forall v(\lambda(v,x) \to S(v)). \tag{19}$$

This is in PIA form. The minimal assignment to s with respect to x is given by $LFP(S, v)$ applied to the antecedent $\lambda(v, x)$. This is equivalent to $\lambda(v, x)$, as S does not occur free in λ. ((19) is 'CIA' — 'constant implies atomic'.)

The 'minimal' $h^\circ \in \mathcal{H}$ is now given by $h^\circ(s) = \{v \in \mathcal{F} : \mathcal{F} \models \lambda(v, x)\}$. So φ is valid in $\mathcal{F}$ at x iff $(\mathcal{F}, h^\circ), x \models s$, iff $\mathcal{F} \models \lambda(x/v, x)$. So φ_1 is valid in a frame $\mathcal{F}$ iff $\mathcal{F} \models \forall x \lambda(x/v, x)$: i.e.,

$$\mathcal{F} \models \forall x\big([LFP(X, x) \,.\, \exists y(R(x, y) \wedge (y = x \vee X(y))](x)\big).$$

This is our frame correspondent. It uses only unary least fixed points, as do all correspondents obtained with our algorithm.

5.4 $\varphi_2 = s \to \nu X\big(\Box(X \wedge \neg s') \vee (\Diamond s \wedge \Diamond s')\big)$

This can be checked to conform to definition 3.5, if we replace the initial '$s \to$ ' by '$\neg s \vee$ '. The skeleton associated with φ_1 above was just $p \wedge q$. For φ_2, the skeleton is nontrivial: φ_2 is equivalent to the Sahlqvist mu-formula

$$\neg\sigma(\eta/p, s/q, s'/q'),$$

where (clearly) s, s' are PIA formulas, $\eta = \neg(\Diamond s \wedge \Diamond s')$ is negative, and

$$\sigma(p, q, q') = q \wedge \mu X(p \wedge \Diamond(q' \vee X))$$

is a $\{q\}$-skeleton and a $\{q'\}$-skeleton. It is not a $\{p\}$-skeleton, because in $p \wedge \Diamond(q' \vee X)$, the right-hand conjunct is not a sentence but the left-hand one involves p. The second conjunct of σ is equivalent to a strict form of pUq'. So $\varphi \equiv s \to \neg([\neg(\Diamond s \wedge \Diamond s')]Us')$.

We calculate the frame correspondent of φ_2. We will suppress some parentheses to aid readability. Note that σ is normal in q and q', so (as we mentioned after the statement of corollary 4.7) $ST_x(\sigma)$ is equivalent to the rather simple formula

$$\exists yy'\big(x = y \wedge [LFP(X, x) \,.\, Px \wedge \exists z(Rxz \wedge (z = y' \vee Xz))] \wedge Qy \wedge Q'y'\big). \quad (20)$$

We now take $ST_x(\eta) = \neg(\exists v(Rxv \wedge Sv) \wedge \exists v(Rxv \wedge S'v))$ and replace references to S, S' by the minimal valuations for them, which are $\{y\}, \{y'\}$, respectively. We obtain $\neg(\exists v(Rxv \wedge v = y) \wedge \exists v(Rxv \wedge v = y'))$, which simplifies to $\neg(Rxy \wedge Rxy')$. This is substituted for Px in (20) and the conjuncts $Qy, Q'y'$ are deleted since they will automatically be true under the minimal assignment. We obtain

$$\exists yy'\big(x = y \wedge [LFP(X, x) \,.\, \neg(Rxy \wedge Rxy') \wedge \exists z(Rxz \wedge (z = y' \vee Xz))]\big),$$

and this holds at a world x iff φ_2 is not valid at x. So our frame correspondent for φ_2 expresses the negation of the above for all x, which boils down to:

$$\forall xyy'\big(x = y \to GFP[X, x]\big[\forall z(Rxz \to (z \neq y' \wedge Xz)) \vee (Rxy \wedge Rxy')\big]\big).$$

The correspondent plainly 'says' that for any path $x = x_0 R x_1 R \ldots R x_n = y$ in the frame, with $n > 0$, there is i with $0 \leq i < n$ such that $Rx_i x$ and $Rx_i y$.

This raises some interesting connections with PDL. We do not believe that there is any PDL formula without tests that is valid in the same frames as φ_2, but φ_2 is valid in the same frames as

$$\varphi_3 \quad = \quad p \wedge \langle (?q\,;a)^* \rangle p' \to \langle (?q\,;a)^* \rangle (\Diamond p \wedge \Diamond p'),$$

where q is a new atom and a is a program with accessibility relation R. The idea is roughly that if $(\mathcal{F}, h), x \models p \wedge \langle (?q\,;a)^* \rangle p'$, then there is y with $R^* xy$ at which p' holds, and a path from x to y along which q holds. The minimal values of p, p', q are now x, y, and the path, respectively. The consequent now states that some world t on the path is R-related to worlds satisfying these minimal values of p, p' : i.e., Rtx and Rty.

In general, the minimal value of q (the path) is not unique, and considering automorphisms shows that it is not going to be definable in terms of x, y in any logic at all. So such PDL-formulas seem to be (possibly much) more powerful than Sahlqvist mu-formulas. On the other hand, Sahlqvist mu-formulas allow rather free use of fixed points, and in expressive power may go beyond even PDL-formulas with tests. Consider for example $\mu X \Box X$. As is well known, this defines the well-founded part of any model. This property appears not to be definable in PDL. The exact relationship between the two formalisms is to be the object of further study.

5.5 McKinsey's axiom: $\Box\Diamond p \to \Diamond\Box p$

Of course, not every modal mu-formula, or even every modal formula, has a frame correspondent in FO+LFP. It was mentioned in [4] that McKinsey's axiom $\varphi = \Box\Diamond p \to \Diamond\Box p$ has no such correspondent and that this can be proved using the Löwenheim–Skolem property for LFP (joint work by van Benthem and Goranko).

Here, we give a little more detail of the proof. It is based on [2]; see also [1, theorem 21] and [15, theorem 2.2]. Note first that φ is equivalent to $\Diamond(\Box p \vee \Box\neg p)$. Let $\mathcal{F}$ be the frame whose set of worlds consists of three disjoint parts: a root r; the natural numbers; and the infinite sets X of natural numbers. The accessibility relation R of $\mathcal{F}$ relates r to every X, X to every member of X, and each natural number to itself; these are the only instances of R. It can be verified that φ is valid in $\mathcal{F}$, because for any assignment of p into $\mathcal{F}$, there must be an infinite set X of natural numbers all having the same truth value for p, and $\Box p \vee \Box\neg p$ is consequently true at such an X. Hence $\Diamond(\Box p \vee \Box\neg p)$ is true at the root. Truth of φ at all other worlds of $\mathcal{F}$ is easy to check.

Suppose for contradiction that χ is a (global) frame correspondent of φ in FO+LFP, so that $\mathcal{F} \models \chi$. It follows from the proof of the downward Löwenheim–Skolem property for FO+LFP in [17, theorem 2.4] that there is a countable elementary substructure $\mathcal{F}_0 \preceq \mathcal{F}$ containing all the natural numbers and with $\mathcal{F}_0 \models \chi$, and so φ is valid in $\mathcal{F}_0$. To see that this is impossible, enumerate the sets of natural numbers in $\mathcal{F}_0$ as $X_0, X_1, \ldots$, and select by induction distinct natural numbers $x_0, y_0, x_1, y_1, \ldots$ in such a way that $x_n, y_n \in X_n$ for each n (this is possible because X_n is infinite). Now assign p to $\{x_0, x_1, \ldots\}$. Every set X_n in $\mathcal{F}_0$ contains a point (x_n) satisfying p and a point (y_n) satisfying $\neg p$, so $\Box p \vee \Box\neg p$ is false at every X_n. Hence, φ is false at the root.

6 Coda: proof-theoretic aspects of correspondence arguments

As we observed, our proof gives a constructive algorithm for computing frame equivalents of modal axioms having the required syntax. There is more here than meets the eye. For a start, the equivalents computed by the algorithm need not be the standard formulations one would expect. This is already true for first-order equivalents of basic modal axioms. For instance, the modal axiom $p \to \Box\Diamond p$ gets a computed correspondent $\forall y(Rxy \to \exists z(Ryz \wedge z = x))$, and this only reduces to the natural version $\forall y(Rxy \to Ryx)$ (symmetry) after transformation into a logical equivalent. This 'simplification' phase can be still more drastic for fixed-point formulas. For instance, the description that we gave of the minimal valuation for the antecedent of Löb's Axiom was not the fixed-point produced directly by our general algorithm, but a simplification reached by analyzing that predicate. And likewise, when we substitute that simplified predicate into the consequent of Löb's Axiom, we have to perform one more simplification to see that $\forall y(Rxy \to \forall z(R^*yz \to Rxz))$ is equivalent to $\forall y(Rxy \to \forall z(Ryz \to Rxz))$, and that $\forall y(Rxy \to$ 'no infinite sequence starts from y') is equivalent to 'no infinite sequence starts from x'. What this shows in general is that optimizing correspondence may involve some manipulation in the logic of the correspondence language.

This observation raises another, and more important issue, namely, finding syntactic completeness versions of semantic correspondence results. The point is that correspondence arguments are proofs, but the issue is *where*. [3] already proved the following

> **Theorem.** The correspondence arguments for first-order Sahlqvist axioms in the basic modal language can all be formalized in a weak monadic second-order logic with comprehension (universal instantiation for second-order quantifiers) only for first-order definable sets in the vocabulary $\{R, =\}$ plus unary predicate parameters.

This is a very weak formalism, and it shows that Sahlqvist's Theorem has a very constructive proof. We do not give details of the proof of this result, but the following may be noted. Substituting first-order definable minimal predicates is the mentioned universal instantiation. What also needs to be proved is a syntactic version of the semantic monotonicity of the consequent, using the syntactic positive occurrence of its proposition letters.

What is the corresponding result for our arguments? This time, the formalism can again be the mentioned second-order logic, but now we also allow definable predicates in FO+LFP. Again we have to show that all steps in the proof go through. Note that this is much weaker than using all validities of FO+LFP, whose logic is highly complex and non-axiomatizable. By the way, the same may be true for 'simplification' of correspondents like that for Löb's Axiom in FO+LFP. We seem to need only a small part of all the validities of that language, say just the obvious fixed-point axioms and rules.

In this light, what is the point of the usual modal completeness versions of the Sahlqvist Theorem? What these show is that for extracting purely modal consequences, we do not need the full language of correspondence proofs, but only a purely modal sublanguage. The usual incompleteness proofs in modal logic [12] then show that this only goes so far, as this can fail for relatively simple modal axioms. We do not investigate this here, but we do have one telling example. [9] shows that the following variant of Löb's Axiom is frame-incomplete: $\Box(\Box p \leftrightarrow p) \to \Box p$. This formula defines the same class of frames as Löb's Axiom, but it fails to derive the latter. What is significant here is that it does have PI form for its antecedent, so our algorithm would treat it just like Löb's Axiom itself.

But we end with an observation about strength of modal deduction. Sometimes, the mu-calculus seems strong enough to prove exact analogues of our correspondence arguments! Löb's Axiom itself is a good example. For instance, the second 'simplification' stated above is just the provable fixed-point principle $\Box\mu p(\Box p \leftrightarrow \mu p\Box p)$. Next, here is a result from [5]:

FACT 6.1 *Löb's Logic is equivalently axiomatized by the two principles (a) $\Box p \to \Box\Box p$, (b) $\mu p\Box p$.*

Proof. From Löb's Logic to (a) is a purely modal deduction. Next, (b) is derived as follows. By the fixed-point axiom of the mu-calculus, $\Box\mu p(\Box p \to \mu p\Box p)$. So it suffices to get $\Box\mu p\Box p$. Now Löb's Axiom implies $\Box(\Box\mu p\Box p \to \mu p\Box p) \to \Box\mu p\Box p$, and the antecedent of this is derivable by modal Necessitation from the mu-calculus fixed-point axiom. Conversely, assume (a) and (b). We show that, in the modal logic K4, $\mu p\Box p \to (\Box(\Box q \to q) \to \Box q)$. The derivation rule for smallest fixed-points proves $\mu p\Box p \to \alpha$ for any formula α if $\Box\alpha \to \alpha$ is proved. But K4 proves $\Box(\Box(\Box q \to q) \to \Box q) \to (\Box(\Box q \to q) \to \Box q)$. $\Box$

While we have not given any deep results, we have hopefully shown how correspondence arguments have interesting proof-theoretic aspects.

7 Conclusions and future work

We conclude with a discussion of some possible directions for future work.

Multiple recursions. In this paper we defined Sahlqvist mu-formulas and proved that they have FO+LFP-correspondents. Here, in defining the crucial syntactic notion of a PIA-formula, we allowed only one atom. Can we also allow several atoms, while still obtaining mu-formulas that have FO+ LFP-correspondents? This question is related to the 'inductive formulas' of Goranko and Vakarelov [16], which are more general than Sahlqvist formulas and admit minimal valuations constructed step-by-step for each atom. A generalization of Sahlqvist mu-formulas to allow several atoms in PIA-formulas may require combining our approach with the one of [16].

Strengthening the modal base. In this paper we only consider the basic modal language extended with fixed point operators. However, there is room for further expansions involving hybrid modal languages, or the Guarded Fragment with fixed point operators. Extensions of classical Sahlqvist correspondence to these languages have already been studied in, e.g., [8]. We think our approach can be generalized in the same way.

Fragments of the mu-calculus. One can also look into an opposite direction, at languages weaker than mu-calculus, and examine the consequences of the Sahlqvist correspondence developed in this paper. One obvious candidate is propositional dynamic logic (PDL), which has already played a large role in our examples.[3]

The fixed-point correspondence language. We now turn to the other end of our Sahlqvist correspondence: the logic FO+LFP. It is of course of interest to know how much power of this logic we are really using. In other words, in what subfragment of FO+LFP do the correspondents of Sahlqvist mu-formulas 'land'? For the classical Sahlqvist correspondence this question has been answered by Kracht [18, 7]. But for the modal mu-calculus this question is wide open.[4]

Further questions. Of course one could also ask for analogues of other famous definability results for the mu-calculus, such as the Goldblatt–Thomason theorem, which gives necessary and sufficient condition for a class of frames to be modally definable. Another example is Fine's theorem, which states that every elementarily definable modal logic is canonical. There are different ways to formulate canonicity for modal mu-logics, and a useful framework for this might be the admissible semantics of modal mu-calculus used in [6].

To sum everything up, we hope to have shown that the mu-calculus provides a natural new take on many traditional issues in modal definability, and that there is a lot of interesting syntactic and semantic structure awaiting further exploration.

References

[1] Balbiani, P., V. Shehtman, and I. Shapirovsky, 'Every world can see a Sahlqvist world', in G. Governatori, I. Hodkinson, and Y. Venema, (eds.), *Proc. Advances in Modal Logic*, College Publications, 2006, pp. 69–85.

[3] Fontaine [13, §5.5] characterizes PDL-formulas (with the restriction that these formulas may contain only one atom) as a certain subfragment of the mu-calculus.

[4] Mu-calculus formulas retain all the bisimulation-induced key semantic properties of modal ones, such as preservation under generated subframes, p-morphic images, disjoint unions. Can we find some further syntax restrictions?

[2] Benthem, J. van, 'A note on modal formulas and relational properties', *Journal of Symbolic Logic*, 40 (1975), 1, 55–58.

[3] Benthem, J. van, *Modal logic and classical logic*, Bibliopolis, Naples, 1985.

[4] Benthem, J. van, 'Minimal predicates, fixed-points, and definability', *J. Symbolic Logic*, 70 (2005), 696–712.

[5] Benthem, J. van, 'Modal frame correspondences and fixed-points', *Studia Logica*, 83 (2006), 133–155.

[6] Bezhanishvili, N., and I. Hodkinson, 'Sahlqvist theorem for modal fixed point logic', (2010). Submitted.

[7] Blackburn, P., M. de Rijke, and Y. Venema, *Modal logic*, Tracts in Theoretical Computer Science, Cambridge University Press, Cambridge, UK, 2001.

[8] Cate, B. ten, M. Marx, and P. Viana, 'Hybrid logics with Sahlqvist axioms', *Logic J. IGPL*, 13 (2005), 293–300.

[9] Cresswell, M., 'An incomplete decidable modal logic', *J. Symbolic Logic*, 49 (1984), 520–527.

[10] Ebbinghaus, H-D., and J. Flum, *Finite model theory*, 2nd edn., Perspectives in mathematical logic, Springer-Verlag, New York, 1999.

[11] Esakia, L. L., 'Topological Kripke models', *Soviet Math. Dokl.*, 15 (1974), 147–151.

[12] Fine, K., 'An incomplete logic containing S4', *Theoria*, 40 (1974), 23–29.

[13] Fontaine, G., *Modal fixpoint logic: some model theoretic questions*, Ph.D. thesis, ILLC, Amsterdam, 2010. ILLC Dissertation Series DS-2010-09.

[14] Givant, S., and Y. Venema, 'The preservation of Sahlqvist equations in completions of Boolean algebras with operators', *Algebra Universalis*, 41 (1999), 47–84.

[15] Goldblatt, R., and I. Hodkinson, 'The McKinsey–Lemmon logic is barely canonical', *Australasian J. Logic*, 5 (2007), 1–19.

[16] Goranko, V., and D. Vakarelov, 'Elementary canonical formulae: extending Sahlqvist's theorem', *Ann. Pure. Appl. Logic*, 141 (2006), 180–217.

[17] Grädel, E., 'The decidability of guarded fixed point logic', in J. Gerbrandy, M. Marx, M. de Rijke, and Y. Venema, (eds.), *JFAK. Essays Dedicated to Johan van Benthem on the Occasion of his 50th Birthday*, Vossiuspers, Amsterdam University Press, Amsterdam, 1999. CD-ROM, ISBN 90 5629 104 1.

[18] Kracht, M., ‘How completeness and correspondence theory got married’, in M. de Rijke, (ed.), *Diamonds and Defaults*, Kluwer Academic Publishers, 1993, pp. 175–214.

[19] Sahlqvist, H., ‘Completeness and correspondence in the first and second order semantics for modal logic’, in S. Kanger, (ed.), *Proc. 3rd Scandinavian logic symposium, Uppsala, 1973*, North Holland, Amsterdam, 1975, pp. 110–143.

[20] Sambin, G., and V. Vaccaro, ‘A new proof of Sahlqvist’s theorem on modal definability and completeness’, *J. Symbolic Logic*, 54 (1989), 992–999.

[21] Tarski, A., ‘A lattice-theoretical fixpoint theorem and its applications’, *Pacific Journal of Mathematics*, 5 (1955), 285–309.

Nick Bezhanishvili
Department of Computing, Imperial College London, London SW7 2AZ, UK
nbezhani@doc.ic.ac.uk

Johan van Benthem
Institute for Logic, Language & Computation (ILLC), University of Amsterdam, P.O. Box 94242, 1090 GE Amsterdam, The Netherlands
J.vanBenthem@uva.nl

Ian Hodkinson
Department of Computing, Imperial College London, London SW7 2AZ, UK
i.hodkinson@imperial.ac.uk