Inquisitive Bisimulation

Ivano Ciardelli and Martin Otto*

March 8, 2018

Abstract

Inquisitive modal logic, InqML, is a generalisation of standard Kripke-style modal logic. In its epistemic incarnation, it extends standard epistemic logic to capture not just the information that agents have, but also the questions that they are interested in. Technically, InqML fits within the family of logics based on team semantics. From a model-theoretic perspective, it takes us a step in the direction of monadic second-order logic, as inquisitive modal operators involve quantification over sets of worlds. We introduce and investigate the natural notion of bisimulation equivalence in the setting of InqML. We compare the expressiveness of InqML and first-order logic in the context of relational structures with two sorts, one for worlds and one for information states. We characterise inquisitive modal logic, as well as its multi-agent epistemic (S5-like) variant, as the bisimulation invariant fragment of first-order logic over various natural classes of two-sorted structures. These results crucially require non-classical methods in studying bisimulation and first-order expressiveness over non-elementary classes of structures, irrespective of whether we aim for characterisations in the sense of classical or of finite model theory.

*Ivano Ciardelli's research was financially supported by the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 680220). Martin Otto's research was partially funded by DFG grant OT 147/6-1: *Constructions and Analysis in Hypergraphs of Controlled Acyclicity*; it was also greatly supported through his participation in a programme on *Logical Structure in Computation* at the Simons Institute in Berkeley in 2016, which is gratefully acknowledged. Parts of the present work were presented in [8].

Contents

1 Introduction

The recently developed framework of *inquisitive logic* [9, 7, 5, 3] can be seen as a generalisation of classical logic which encompasses not only statements, but also questions. One reason why this generalisation is interesting is that it provides a novel perspective on the logical notion of *dependency*, which plays an important rôle in applications (e.g., in database theory) and which has recently received attention in the field of *dependence logic* [27]. Indeed, dependency is nothing but a facet of the fundamental logical relation of entailment, once this is extended so as to apply not only to statements, but also to questions [4]. This connection explains the deep similarities existing between systems of inquisitive logic and systems of dependence logic (see [30, 4, 3, 31]). A different rôle for questions in a logical system comes from the setting of modal logic: once the notion of a modal operator is suitably generalised, questions can be embedded under modal operators to produce new statements that have no "standard" counterpart. This approach was first developed in [10] in the setting of epistemic logic. The resulting *inquisitive epistemic logic* models not only the information that agents have, but also the issues that they are interested in, i.e., the information that they would like to obtain. Modal formulae in inquisitive epistemic logic can express not only that an agent knows that p (by the formula $\Box p$) but also that she knows *whether* p ($\Box ?p$) or that she *wonders whether* p (by the formula $\boxplus ?p$) — a statement that cannot be expressed without the use of embedded questions. As shown in [10], several key notions of epistemic logic generalise smoothly to questions: besides common knowledge we now have *common issues*, the issues publicly entertained by the group; and besides publicly announcing a statement, agents can now also publicly ask a question, which typically results in new common issues. Thus, inquisitive epistemic logic may be seen as one step in extending modal logic from a framework to reason about information and information change, to a richer framework which also represents a higher stratum of cognitive phenomena, in particular issues that may be raised in a communication scenario.

Of course, like standard modal logic, inquisitive modal logic provides a general framework that admits various interpretations, each suggesting corresponding constraints on models. E.g., an interpretation of InqML as a logic of action is suggested in [3]. In that interpretation, a modal formula $\Box ?p$ expresses that whether a certain fact p will come about is pre-determined independently of the agent's choices, while $\boxplus ?p$ expresses that whether p will come about is fully determined by the agent's choices.

From the perspective of mathematical logic, inquisitive modal logic is a natural generalisation of standard modal logic. There, the accessibility relation of a Kripke model associates with each possible world $w \in W$ a set $\sigma(w) \subseteq W$ of possible worlds, namely, the worlds accessible from w; any formula φ of modal logic is semantically associated with a set $|\varphi|_{\mathbb{M}} \subseteq W$ of worlds, namely, the set of worlds where it is true; modalities then express relationships between these sets: for instance, $\Box\varphi$ expresses the fact that $\sigma(w) \subseteq |\varphi|_{\mathbb{M}}$. In the inquisitive setting, the situation is analogous, but both the entity $\Sigma(w)$ attached to a possible world and the semantic extension $[\varphi]_{\mathbb{M}}$ of a formula are sets *of sets* of worlds, rather than simple sets of worlds. Inquisitive modalities still express relationships between these two objects: $\Box\varphi$ expresses the fact that $\bigcup \Sigma(w) \in [\varphi]_{\mathbb{M}}$, while $\boxplus\varphi$ expresses the fact that $\Sigma(w) \subseteq [\varphi]_{\mathbb{M}}$.

In this manner, inquisitive logic leads to a new framework for modal logic that can be viewed as a generalisation of the standard framework. Clearly, this raises the question of whether and how the classical notions and results of modal logic carry over to this more general setting. In this paper we address this question for the fundamental notion of *bisimulation* and for two classical results revolving around this notion, namely, the Ehrenfeucht-Fraïssé theorem for modal logic, and van Benthem style characterisation theorems [15, 29, 26, 21]. A central topic of this paper is the rôle of *bisimulation invariance* as a unifying semantic feature that distinguishes modal logics from classical predicate logics. As in many other areas, from temporal logics and process logics to knowledge representation in AI and database applications, so also in the inquisitive setting we find that the appropriate notion of bisimulation invariance allows for precise model-theoretic characterisations of the expressive power of modal logic in relation to first-order logic.

Our first result is that the right notion of inquisitive bisimulation equivalence $\sim$, with finitary approximation levels $\sim^n$, supports a counterpart of the classical Ehrenfeucht–Fraïssé characterisations for first-order logic or for basic modal logic. This result establishes an exact correspondence between the expressive power of InqML and the *finite* approximation levels of inquisitive bisimulation equivalence: if two points are behaviourally different in a way that can be detected within a finite number of steps, then the difference between them is witnessed by an InqML formula, and vice versa. The result is non-trivial in our setting because of some subtle issues stemming from the interleaving of first- and second-order features in inquisitive modal logic.

Theorem 1.1 (inquisitive Ehrenfeucht–Fraïssé theorem)**.**
Over finite vocabularies, the finite levels $\sim^n$ of inquisitive bisimulation equivalence correspond to the levels of InqML*-equivalence up to modal nesting depth n.*

In order to compare InqML with classical first-order logic, we define a class of two-sorted relational structures, and show how such structures encode models for InqML. With respect to such relational structures we find not only a “standard translation” of InqML into two-sorted first-order logic, but also a van Benthem style characterisation of InqML as the bisimulation-invariant fragment of (two-sorted) first-order logic over several natural classes of models. These results are technically interesting, and they are not available on the basis of classical techniques, because the relevant classes of two-sorted models are non-elementary (in fact, first-order logic is not compact over these classes, as we show). Our techniques yield characterisation theorems both in the setting of arbitrary inquisitive models, and in restriction to just finite ones — i.e. both in the sense of classical model theory and in the sense of finite model theory.

Theorem 1.2. *Inquisitive modal logic can be characterised as the $\sim$-invariant fragment of first-order logic* FO *over natural classes of (finite or arbitrary) relational inquisitive models.*

We go on to extend these results from the basic inquisitive modal setting to the setting of inquisitive epistemic logic — the inquisitive counterpart of multi-agent S5. This setting is technically more challenging due to the additional S5-type constraints on models.

Theorem 1.3. *Inquisitive epistemic logic (in a multi-agent setting) can be characterised as the $\sim$-invariant fragment of* FO *over natural classes of (finite or arbitrary) relational inquisitive epistemic models.*

Beside the conceptual development and the core results themselves, we think that also the methodological aspects of the present investigations have some intrinsic value. Just as inquisitive logic models cognitive phenomena at a level strictly above that of standard modal logic, so the model-theoretic analysis moves up from the level of ordinary first-order logic to a level strictly between first- and second-order logic. This level is realised by first-order logic in a two-sorted framework that incorporates second-order objects in the second sort in a controlled fashion. This leads us to substantially generalise a number of notions and techniques developed in the model-theoretic analysis of modal logic ([15, 21, 11, 22], among others).

2 Inquisitive modal logic

In this section we provide an essential introduction to inquisitive modal logic, InqML [3]. For further details and proofs, we refer to §7 of [3].

2.1 Foundations of inquisitive semantics

Usually, the semantics of a logic specifies truth-conditions for the formulae of the logic. In modal logics these truth-conditions are relative to possible worlds in a Kripke model. However, this approach is limited in an important way: while suitable for statements, it is inadequate for questions. To overcome this limitation, inquisitive logic interprets formulae not relative to states of affairs (possible worlds), but relative to states of information. Following a tradition that goes back to the work of Hintikka [18], information states are modelled extensionally as sets of worlds, namely, those worlds which are compatible with the given information.[1]

[1] An analogous step from single worlds to sets of worlds (or, depending on the setting, from assignments to sets of assignments) is taken in recent work on independence-friendly logic [19, 20] and dependence logic [27, 28, 1, 14, 30, 31], where sets of worlds are referred to as *teams*. Although they originated independently and were developed for different purposes, inquisitive logic and dependence logic are tightly related. For detailed discussion of this connection, see [4, 3].

Definition 2.1. [information states]
An *information state* over a set of worlds W is a subset $s \subseteq W$.

Rather than specifying when a sentence is *true* at a world w, inquisitive semantics specifies when a sentence is *supported* by an information state s: intuitively, for a statement α this means that the information available in s implies that α is true; for a question μ, it means that the information available in s settles μ.

If t and s are information states and $t \subseteq s$, this means that t holds at least as much information as s: we say that t is an *extension* of s. If t is an extension of s, everything that is supported at s will also be supported at t. This is a key feature of inquisitive semantics, and it leads naturally to the notion of an *inquisitive state* (see [6, 25, 10]).

Definition 2.2. [inquisitive states]
An *inquisitive state* over a set of possible worlds W is a non-empty set of information states $\Pi \subseteq \wp(W)$ that is *downward closed* in the sense that

- $s \in \Pi$ implies $t \in \Pi$ for all $t \subseteq s$ (downward closure).

The downward closure condition requires that Π be closed under extensions of information states, i.e. robust under any strengthening of the available information.

2.2 Inquisitive modal models

A Kripke frame can be thought of as a set W of worlds together with a map σ that equips each world with a set of worlds $\sigma(w)$ — the set of worlds that are accessible from w — i.e., an information state.

Similarly, an inquisitive modal frame consists of a set W of worlds together with an *inquisitive assignment*, a map $\Sigma : W \to \wp\wp(W)$ that assigns to each world a corresponding inquisitive state $\Sigma(w)$, i.e., a set of information states closed under subsets. An inquisitive modal model is an inquisitive frame equipped with a propositional assignment (a valuation function for propositional atoms).

Definition 2.3. [inquisitive modal models]
An inquisitive modal frame is a pair $\mathbb{F} = (W, \Sigma)$, where W is a set, whose elements are referred to as *worlds*, and $\Sigma \colon W \to \wp\wp(W)$ assigns to each world $w \in W$ an inquisitive state $\Sigma(w)$.

An inquisitive modal model for a set $\mathcal{P}$ of propositional atoms is a pair $\mathbb{M} = (\mathbb{F}, V)$ where $\mathbb{F}$ is an inquisitive modal frame, and $V \colon \mathcal{P} \to \wp(W)$ is a propositional assignment.

A world-(or state-)pointed inquisitive modal model is a pair consisting of a model $\mathbb{M}$ and a distinguished world (or state) in $\mathbb{M}$.

With an inquisitive modal model $\mathbb{M}$ we can always associate a standard Kripke model $\mathfrak{K}(\mathbb{M})$ having the same set of worlds and modal accessibility map $\sigma : W \to \wp(W)$ induced by the inquisitive map Σ according to

$$\begin{array}{rcl} \sigma : W & \longrightarrow & \wp(W) \\ w & \longmapsto & \sigma(w) := \bigcup \Sigma(w). \end{array}$$

A natural interpretation for inquisitive modal models is the epistemic one, developed in [10, 2]. In this interpretation, the map Σ is taken to describe not only an agent's *knowledge*, as in standard epistemic logic, but also an agent's *issues*. The agent's knowledge state at w, $\sigma(w) = \bigcup \Sigma(w)$, consists of all the worlds that are compatible with what the agent knows. The agent's inquisitive state at w, $\Sigma(w)$, consists of all those information states where the agent's issues are settled. This interpretation is particularly interesting in the multi-modal setting, where a model comes with multiple state maps Σ_a, one for each agent a in a set $\mathcal{A}$. Moreover, this specific interpretation suggests some constraints on the maps Σ_a, analogous to the usual S5 constraints on Kripke models.

Definition 2.4. [inquisitive epistemic models]
An *inquisitive epistemic frame* for a set $\mathcal{A}$ of agents is a pair $\mathbb{F} = (W, (\Sigma_a)_{a \in \mathcal{A}})$, where each map $\Sigma_a : W \to \wp\wp(W)$ assigns to each world w an inquisitive state $\Sigma_a(w)$ in accordance with the following constraints (where $\sigma_a(w) = \bigcup \Sigma_a(w)$):

- $w \in \sigma_a(w)$ (factivity);
- $v \in \sigma_a(w) \;\Rightarrow\; \Sigma_a(v) = \Sigma_a(w)$ (full introspection).

An inquisitive epistemic model consists of an inquisitive epistemic frame together with a propositional assignment $V : \mathcal{P} \to \wp(W)$.

It is easy to verify that the Kripke frame associated with an inquisitive epistemic frame is an $S5$ frame, i.e., the accessibility maps σ_a correspond to accessibility relations $R_a := \{(v, w) : v \in \sigma_a(w)\}$ that are equivalence relations on W.

Example 2.5. Consider a model with four worlds, $w_{pq}, w_{p\overline{q}}, w_{\overline{p}q}, w_{\overline{pq}}$, where the subscript indicate the propositional valuation at each world. The inquisitive state map Σ is as follows, where $S^{\downarrow}$ indicates the closure of the set $S \subseteq \wp(W)$ under subsets.

$$\Sigma(w_{pq}) = \Sigma(w_{p\overline{q}}) = \{\{w_{pq}\}, \{w_{q\overline{q}}\}\}^{\downarrow}$$
$$\Sigma(w_{\overline{p}q}) = \Sigma(w_{\overline{pq}}) = \{\{w_{\overline{p}q}, w_{\overline{qq}}\}\}^{\downarrow}$$

This model is depicted in Figure 1. At a world w, the epistemic state $\sigma(w)$ of the agent consists of those worlds included in the same dashed area as w; the solid blocks inside this area are the maximal elements of the inquisitive state $\Sigma(w)$ — i.e. the maximal states in which the issue is resolved.

At worlds w_{pq} and $w_{p\overline{q}}$, the agent's knowledge state is $\{w_{pq}, w_{p\overline{q}}\}$: that is, the agent knows that p is true, but not whether q is true. Moreover, in order to settle the agent's issues it is necessary and sufficient to reach an extension

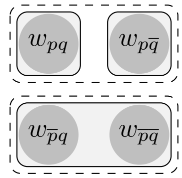

Figure 1: A single-agent inquisitive epistemic model

of the current state which settles whether q. In short, then, these are worlds where the agent knows that p and wonders whether q.

At worlds $w_{\overline{p}q}$ and $w_{\overline{pq}}$, the agent's knowledge state is $\{w_{\overline{p}q}, w_{\overline{qq}}\}$: that is, the agent knows that $\neg p$, but not whether q. However, at these worlds no further information is needed to resolve the agent's issues. Thus, these are worlds where the agent knows that $\neg p$ and does not have any remaining issues.

2.3 Inquisitive modal logic

The syntax of inquisitive modal logic InqML is given by:

$$\varphi ::= p \mid \bot \mid (\varphi \wedge \varphi) \mid (\varphi \to \varphi) \mid (\varphi \mathbin{\vee\!\!\!\vee} \varphi) \mid \Box\varphi \mid \boxplus\varphi$$

We treat negation and disjunction as defined connectives (syntactic shorthands) according to:

- $\neg\varphi := \varphi \to \bot$
- $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$

In this sense, the above syntax includes standard propositional formulae in terms of atoms and connectives $\wedge$ and $\to$ together with the defined $\neg$ and $\vee$. As we will see, the semantics for such formulae will be essentially the same as in standard propositional logic.

In addition to standard connectives, our language contains a new connective, $\mathbin{\vee\!\!\!\vee}$, called *inquisitive disjunction*. We may read formulae built up by means of this connective as propositional questions. E.g., we read the formula $p \mathbin{\vee\!\!\!\vee} \neg p$ as the question *whether or not* p, and we abbreviate this formula as $?p$.

Finally, our language contains two modalities, which are allowed to embed both statements and questions. As we shall see, both these modalities coincide with a standard Kripke box modality when applied to statements, but crucially differ when applied to questions. In particular, under an epistemic interpretation $\Box?p$ expresses the fact that the agent knows whether p, while $\boxplus?p$ expresses (roughly) the fact that the agent is interested in the issue whether p.

The syntax of inquisitive epistemic logic is defined analogously, except that modalities are indexed by agents; that is, for every agent $a \in \mathcal{A}$ we have two corresponding modalities $\Box_a$ and $\boxplus_a$, which are interpreted based on the state map Σ_a associated with the agent.[2]

As mentioned above, the semantics of InqML is given in terms of support in an information state, rather than truth at a possible world.[3]

Definition 2.6. [semantics of InqML]
Let $\mathbb{M} = (W, \Sigma, V)$ be an inquisitive modal model, $s \subseteq W$:

- $\mathbb{M}, s \models p \iff s \subseteq V(p)$
- $\mathbb{M}, s \models \bot \iff s = \emptyset$
- $\mathbb{M}, s \models \varphi \wedge \psi \iff \mathbb{M}, s \models \varphi$ and $\mathbb{M}, s \models \psi$
- $\mathbb{M}, s \models \varphi \to \psi \iff \forall t \subseteq s : \mathbb{M}, t \models \varphi \Rightarrow \mathbb{M}, t \models \psi$
- $\mathbb{M}, s \models \varphi \mathbin{\vee\!\!\!\vee} \psi \iff \mathbb{M}, s \models \varphi$ or $\mathbb{M}, s \models \psi$
- $\mathbb{M}, s \models \Box\varphi \iff \forall w \in s : \mathbb{M}, \sigma(w) \models \varphi$
- $\mathbb{M}, s \models \boxplus\varphi \iff \forall w \in s\ \forall t \in \Sigma(w) : \mathbb{M}, t \models \varphi$

As an illustration, consider the support conditions for the formula $?p := p \mathbin{\vee\!\!\!\vee} \neg p$: this formula is supported by a state s in case p is true at all worlds in s (i.e., if the information available in s implies that p is true) or in case p is false at all worlds in s (i.e., if the information available in s implies that p is false). Thus, $?p$ is supported precisely by those information states that settle whether or not p is true.

The following two properties hold generally in InqML.

Proposition 2.7.

- *persistency: if $\mathbb{M}, s \models \varphi$ and $t \subseteq s$, then $\mathbb{M}, t \models \varphi$;*
- *semantic ex-falso: $\mathbb{M}, \emptyset \models \varphi$ for all $\varphi \in$ InqML.*

The first principle says that support is preserved as information increases, i.e., as we move from a state to an extension of it. The second principle says that the empty set of worlds — the inconsistent information state — vacuously supports every formula. Together, these principles imply that the support set $[\varphi]_{\mathbb{M}} := \{s \subseteq W : \mathbb{M}, s \models \varphi\}$ of a formula is downward closed and non-empty, i.e., it is an inquisitive state.

Although the primary notion of our semantics is support at an information state, truth at a world is retrieved by defining it as support with respect to singleton states.

[2]In [10, 2] the modalities $\Box_a$ and $\boxplus_a$ are denoted K_a and E_a, and read as "know" and "entertain" respectively.

[3]This means that InqML fits within the quickly growing family of logics based on a *team semantics*. See Footnote 1 and the references therein.

Definition 2.8. [truth]
φ is *true* at a world w in a model $\mathbb{M}$, denoted $\mathbb{M}, w \models \varphi$, in case $\mathbb{M}, \{w\} \models \varphi$.

Spelling out Definition 2.6 in the special case of singleton states, we see that standard connectives have the usual truth-conditional behaviour. For modal formulae, we find the following truth-conditions.

Proposition 2.9 (truth conditions for modal formulae)**.**

- $\mathbb{M}, w \models \Box\varphi \iff \mathbb{M}, \sigma(w) \models \varphi$
- $\mathbb{M}, w \models \boxplus\varphi \iff \forall t \in \Sigma(w) : \mathbb{M}, t \models \varphi$

Notice that truth in InqML cannot be given a direct recursive definition, as the truth conditions for modal formulae $\Box\varphi$ and $\boxplus\varphi$ depend on the support conditions for φ — not just on its truth conditions.

For many formulae, support at a state just boils down to truth at each world. We refer to these formulae as *truth-conditional*.[4]

Definition 2.10. [truth-conditional formulae]
We say that a formula φ is *truth-conditional* if for all models $\mathbb{M}$ and information states s: $\mathbb{M}, s \models \varphi \iff \mathbb{M}, w \models \varphi$ for all $w \in s$.

Following [3], we view truth-conditional formulae as statements, and non-truth-conditional formulae as questions. The next proposition identifies a large class of formulae which are truth-conditional.

Proposition 2.11. *Atomic formulae, $\bot$, and all formulae of the form $\Box\varphi$ and $\boxplus\varphi$ are truth-conditional. Furthermore, the class of truth-conditional formulae is closed under all connectives except for $\mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}}$.*

Using this fact, it is easy to see that all formulae of standard modal logic, i.e., formulae which do not contain $\mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}}$ or $\boxplus$, receive exactly the same truth-conditions as in standard modal logic.

Proposition 2.12. *If φ is a formula not containing $\mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}}$ or $\boxplus$, then $\mathbb{M}, w \models \varphi \iff \mathfrak{K}(\mathbb{M}), w \models \varphi$ in standard Kripke semantics.*

As long as questions are not around, the modality $\boxplus$ also coincides with $\Box$, and with the standard box modality. That is, if φ is truth-conditional, then

$$M, w \models \Box\varphi \iff M, w \models \boxplus\varphi \iff M, v \models \varphi \text{ for all } v \in \sigma(w).$$

Thus, the two modalities coincide on statements. However, they come apart when they are applied to questions. For an illustration, consider the formulae $\Box?p$ and $\boxplus?p$ in the epistemic setting: $\Box?p$ is true iff the knowledge state of the agent, $\sigma(w)$, settles the question $?p$; thus, $\Box?p$ expresses the fact that the agent knows whether p. By contrast, $\boxplus?p$ is true iff any information state $t \in \Sigma(w)$, i.e., any state that settles the agent's issues, also settles $?p$; thus $\boxplus?p$ expresses that finding out whether p is part of the agent's goals.

[4] In team semantic terminology (e.g., [27, 31]), truth-conditional formulae are called *flat*.

Example 2.13. Consider again the model of Example 2.5. The agent's knowledge state at world w_{pq} is $\sigma(w_{pq}) = \{w_{pq}, w_{p\overline{q}}\}$. Since $\{w_{pq}, w_{p\overline{q}}\}$ does not support $?q$ we have $\mathbb{M}, w \models \neg\Box ?q$. On the other hand, since the agent's inquisitive state is $\Sigma(w_{pq}) = \{\{w_{pq}\}, \{w_{p\overline{q}}\}\}^{\downarrow}$, and since each element in this state supports $?q$, we do have $\mathbb{M}, w_{pq} \models \boxplus ?q$. This witnesses that, at world w_{pq}, the agent does not know whether q ($\neg\Box ?p$), but she's interested in finding out ($\boxplus ?q$). By contrast, one can check that at world $w_{\overline{p}q}$ we have $\mathbb{M}, w_{\overline{p}q} \models \neg\Box ?q \wedge \neg\boxplus ?q$, witnessing that at this world, the agent is neither informed about whether q, nor interested in finding out.

3 Inquisitive bisimulation

An inquisitive modal model can be seen as a structure with two sorts of entities, worlds and information states, which interact with each other. On the one hand, an information state s is completely determined by the worlds that it contains; on the other hand, a world w is determined by the atoms it makes true and the information states which lie in $\Sigma(w)$. Taking a more behavioural perspective, we can look at an inquisitive modal model as a model where two kinds of transitions are possible: from an information state s, we can make a transition to a world $w \in s$, and from a world w, we can make a transition to an information state $s \in \Sigma(w)$. This suggests a natural notion of bisimilarity, together with its natural finite approximations of n-bisimilarity for $n \in \mathbb{N}$. As usual, these notions can equivalently be defined either in terms of back-and-forth systems or in terms of strategies in corresponding bisimulation games. We start from the latter due to its more immediate and intuitive appeal to the underlying dynamics of a "probing" of behavioural equivalence.

The inquisitive bisimulation game is played by two players, **I** and **II**, who act as challenger and defender of a similarity claim involving a pair of worlds w and w' or information states s and s' over two models $\mathbb{M} = (W, \Sigma, V)$ and $\mathbb{M}' = (W', \Sigma', V')$. We denote world-positions as (w, w') and state-positions as (s, s'), where $w \in W, w' \in W'$ and $s \in \wp(W), s' \in \wp(W')$, respectively. The game proceeds in rounds that alternate between world-positions and state-positions. Playing from a world-position (w, w'), **I** chooses an information state in the inquisitive state associated to one of these worlds ($s \in \Sigma(w)$ or $s' \in \Sigma'(w')$) and **II** must respond by choosing an information state on the opposite side, which results in a state-position (s, s'). Playing from a state-position (s, s'), **I** chooses a world in either state ($w \in s$ or $w' \in s'$) and **II** must respond by choosing a world from the other state, which results in a world-position (w, w'). A round of the game consists of four moves leading from a world-position to another.

In the bounded version of the game, the number of rounds is fixed in advance. In the unbounded version, the game is allowed to go on indefinitely. Either player loses when stuck for a move. The game ends with a loss for **II** in any world-position (w, w') that shows a discrepancy at the atomic level, i.e., such that w and w' disagree on the truth of some $p \in \mathcal{P}$. All other plays, an in particular infinite runs of the unbounded game, are won by **II**.

Definition 3.1. [bisimulation equivalence]
Two world-pointed models $\mathbb{M}, w$ and $\mathbb{M}', w'$ are *n-bisimilar*, $\mathbb{M}, w \sim^n \mathbb{M}', w'$, if **II** has a winning strategy in the n-round game starting from (w, w'). $\mathbb{M}, w$ and $\mathbb{M}', w'$ are *bisimilar*, denoted $\mathbb{M}, w \sim \mathbb{M}', w'$, if **II** has a winning strategy in the unbounded game starting from (w, w').

Two state-pointed models $\mathbb{M}, s$ and $\mathbb{M}', s'$ are (n-)bisimilar, denoted $\mathbb{M}, s \sim \mathbb{M}', s'$ (or $\mathbb{M}, s \sim^n \mathbb{M}', s'$), if every world in s is (n-)bisimilar to some world in s' and vice versa.

Two models $\mathbb{M}$ and $\mathbb{M}'$ are *globally bisimilar*, denoted $\mathbb{M} \sim \mathbb{M}'$, if every world in $\mathbb{M}$ is bisimilar to some world in $\mathbb{M}'$ and vice versa.

These notions generalise naturally to the multi-modal setting with inquisitive assignments $(\Sigma_a)_{a \in \mathcal{A}}$ for a set $\mathcal{A}$ of agents; at a world-position, player **I** also gets the choice of which agent to probe.

Now let us turn to the static perspective on inquisitive bisimulations. One natural way to define a bisimulation between two models $\mathbb{M}$ and $\mathbb{M}'$ is as a relation which pairs up both the worlds and the states of these two models in such a way as to guarantee a winning strategy in the unbounded bisimulation game. This leads to the following definition.

Definition 3.2. [bisimulation relations]
Let $M = (W, \Sigma, V)$ and $M' = (W', \Sigma', V')$ be two inquisitive modal models. A non-empty relation $Z \subseteq W \times W' \cup \wp(W) \times \wp(W')$ is called a bisimulation in case the following constraints are satisfied:

- atom equivalence: if wZw' then for all $p \in \mathcal{P}$, $w \in V(p) \iff w' \in V'(p)$
- state-to-world back&forth: if sZs' then
  - for all $w \in s$ there is some $w \in s'$ s.t. wZw'
  - for all $w' \in s'$ there is some $w \in s$ s.t. wZw'
- world-to-state back&forth: if wZw' then
  - for all $s \in \Sigma(w)$ there is some $s' \in \Sigma'(w')$ s.t. sZs'
  - for all $s' \in \Sigma'(w')$ there is some $s \in \Sigma(w)$ s.t. sZs'

It is then routine to check that bisimilarity can be characterised in terms of the existence of a bisimulation relation.

Proposition 3.3. *Let $\mathbb{M}, x$ and $\mathbb{M}', x'$ be two world- or state-pointed models. $\mathbb{M}, x \sim \mathbb{M}', x' \iff$ there exists a bisimulation Z such that xZx'.*

Alternatively, we can view an inquisitive bisimulation as a relation which is defined exclusively on the worlds of the two models. We will call such a relation a world-bisimulation. In order to define it, let us first fix a way to lift a binary relation between two sets to a relation between the corresponding powersets.

Definition 3.4. The *lifting* of a relation $Y \subseteq W \times W'$ to information states is the relation $\overline{Y} \subseteq \wp(W) \times \wp(W')$ linking information states s and s' iff

- for all $w \in s$ there is a $w' \in s'$ s.t. wYw'
- for all $w' \in s'$ there is a $w \in s$ s.t. wYw'

Definition 3.5. [world-bisimulation]
Let $M = (W, \Sigma, V)$ and $M' = (W', \Sigma', V')$ be two inquisitive modal models. A non-empty relation $Y \subseteq W \times W'$ is called a *world-bisimulation* in case the following constraints are satisfied whenever wYw':

- atom equivalence:
  - $\forall p \in \mathcal{P}$: $w \in V(p) \iff w' \in V'(p)$
- back&forth:
  - for all $s \in \Sigma(w)$ there is a $s' \in \Sigma'(w')$ s.t. $s\overline{Y}s'$
  - for all $s' \in \Sigma'(w')$ there is a $s \in \Sigma(w)$ s.t. $s\overline{Y}s'$

Bisimulations and world-bisimulations are tightly connected, as the following proposition brings out. The straightforward proof is omitted.

Proposition 3.6. *If Z is a bisimulation between two models $\mathbb{M}$ and $\mathbb{M}'$, then its restriction to worlds, $Z_\bullet := Z \cap (W \times W')$, is a world-bisimulation. Conversely, if Y is a world-bisimulation, then $Y \cup \overline{Y}$ is a bisimulation.*

If Z is a bisimulation, then Z is included in $Z_\bullet \cup \overline{Z_\bullet}$, but not necessarily identical to it. Thus, a bisimulation is not uniquely determined by its restriction $Z_\bullet$ to worlds. Rather, given a world-bisimulation Y, the bisimulation $Y \cup \overline{Y}$ is the largest among the bisimulations Z with $Z_\bullet = Y$.

Corollary 3.7. *Two world-pointed models $\mathbb{M}, w$ and $\mathbb{M}', w'$ are bisimilar iff there exists a world-bisimulation Y such that wYw'. Two state-pointed models $\mathbb{M}, s$ and $\mathbb{M}', s'$ are bisimilar iff there exists a world-bisimulation Y such that $s\overline{Y}s'$.*

We now turn to the finite levels of bisimilarity.

Definition 3.8. Let $\mathbb{M}$ and $\mathbb{M}'$ be two inquisitive modal models. A back-and-forth system of height n is a family $(Z_i)_{i \leq n}$ of non-empty relations $Z_i \subseteq W \times W' \cup \wp(W) \times \wp(W')$ satisfying the following constraints for each $i \leq n$:

- atom equivalence: if wZ_iw' then for all $p \in \mathcal{P}$, $w \in V(p) \iff w' \in V'(p)$
- state-to-world back&forth: if sZ_is' then
  - for all $w \in s$ there is some $w \in s'$ s.t. wZ_iw'
  - for all $w' \in s'$ there is some $w \in s$ s.t. wZ_iw'
- world-to-state back&forth: if $i > 0$ and wZ_iw' then
  - for all $s \in \Sigma(w)$ there is some $s' \in \Sigma'(w')$ s.t. $sZ_{i-1}s'$

– for all $s' \in \Sigma'(w')$ there is some $s \in \Sigma(w)$ s.t. $sZ_{i-1}s'$

It is straightforward to check that n-bisimilarity can be characterised in terms of back&forth systems as follows.

Proposition 3.9. *Let $\mathbb{M}, x$ and $\mathbb{M}', x'$ be two world- or state-pointed models. $\mathbb{M}, x \sim^n \mathbb{M}', x'$ if, and only if, there exists a back&forth system $(Z_i)_{i \leqslant n}$ such that $xZ_n x'$.*

Analogously to what we did for full bisimilarity, it is also possible to give a purely world-based notion of back&forth-system of height n as a family of relations $(Y_i)_{i \leq n} \subseteq W \times W'$. As expected, n-bisimilarity can then be characterised in terms of the existence of such a system, in a way analogous to the one given by Corollary 3.7. We leave the details to the reader.

4 An Ehrenfeucht–Fraïssé theorem

The crucial rôle of these notions of equivalence for the model theory of inquisitive modal logic is brought out in a corresponding Ehrenfeucht–Fraïssé theorem.

Using the standard notion of the modal depth of a formula, we denote as $\textsc{InqML}_n$ the class of InqML-formulae of depth up to n. It is easy to see that the semantics of any formula in $\textsc{InqML}_n$ is preserved under n-bisimilarity; as a consequence, all of inquisitive modal logic is preserved under full bisimilarity. The following analogue of the classical Ehrenfeucht–Fraïssé theorem shows that, for finite sets $\mathcal{P}$ of atomic propositions, n-bisimilarity coincides with logical indistinguishability in $\textsc{InqML}_n$, which we denote as $\equiv^n_{\textsc{InqML}}$:

$$\mathbb{M}, s \equiv^n_{\textsc{InqML}} \mathbb{M}', s' \;:\Longleftrightarrow\; \begin{cases} \mathbb{M}, s \models \varphi \;\;\Leftrightarrow\;\; \mathbb{M}', s' \models \varphi \\ \qquad\quad \text{for all } \varphi \in \textsc{InqML}_n. \end{cases}$$

Theorem 4.1 (Ehrenfeucht–Fraíssé theorem for InqML)**.**
Over any finite set of atomic propositions $\mathcal{P}$, for any $n \in \mathbb{N}$ and inquisitive state-pointed modal models $\mathbb{M}, s$ and $\mathbb{M}', s'$:

(i) $\mathbb{M}, s \sim^n \mathbb{M}', s' \iff \mathbb{M}, s \equiv^n_{\textsc{InqML}} \mathbb{M}', s'$

(ii) $\mathbb{M}, w \sim^n \mathbb{M}', w' \iff \mathbb{M}, w \equiv^n_{\textsc{InqML}} \mathbb{M}', w'$

Notice that item (ii) of the theorem follows from item (i) by taking s and s' to be singleton states. As usual, the crucial implication of the theorem, from right to left, follows from the existence of *characteristic formulae* for $\sim^n$-classes of pointed models — and it is here that the finiteness of $\mathcal{P}$ is crucial. Notice, however, that while we can expect a formula to uniquely characterise the $\sim^n$ class of a world, we cannot expect a formula to uniquely define the $\sim^n$-class of an information state, for this would conflict with the persistency property of the logic (Proposition 2.7): if a formula is supported at $\mathbb{M}, s$, it must also be supported at $\mathbb{M}, s'$ for all $s' \subseteq s$ even when $\mathbb{M}, s \not\sim^n \mathbb{M}, s'$. However, the next proposition shows that $\textsc{InqML}_n$-formulae characterise the $\sim_n$-class of an information state up to persistency.

Proposition 4.2 (characteristic formulae for $\sim^n$-classes)**.**
Let $\mathbb{M}, w$ be a world-pointed model and $\mathbb{M}, s$ a state-pointed model over a finite set of atomic propositions $\mathcal{P}$. There are InqML*-formulae $\chi^n_{\mathbb{M},w}$ and $\chi^n_{\mathbb{M},s}$ of modal depth n such that:*

(i) $\mathbb{M}', w' \models \chi^n_{\mathbb{M},w} \iff \mathbb{M}', w' \sim^n \mathbb{M}, w$

(ii) $\mathbb{M}', s' \models \chi^n_{\mathbb{M},s} \iff \mathbb{M}', s' \sim^n \mathbb{M}, t$ *for some* $t \subseteq s$

These results can be extended straightforwardly to a multi-modal inquisitive setting with a finite set $\mathcal{A}$ of agents.

Proof. By simultaneous induction on n, we define formulae $\chi^n_{\mathbb{M},w}$ and $\chi^n_{\mathbb{M},s}$ together with auxiliary formulae $\chi^n_{\mathbb{M},\Pi}$ for all worlds w, information states s and inquisitive states Π over $\mathbb{M}$. Given two inquisitive states Π and Π' in models $\mathbb{M}$ and $\mathbb{M}'$, we write $\mathbb{M}, \Pi \sim^n \mathbb{M}', \Pi'$ if every state $s \in \Pi$ is n-bisimilar to some state $s' \in \Pi'$, and vice versa. Dropping reference to the fixed $\mathbb{M}$, we let:

$$\begin{aligned}
\chi^0_w &= \bigwedge\{p \colon w \in V(p)\} \wedge \bigwedge\{\neg p \colon w \notin V(p)\} \\
\chi^n_s &= \bigvee\{\chi^n_w \colon w \in s\} \\
\chi^n_\Pi &= \mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}}\!\!\!\!\;\; \{\chi^n_s \colon s \in \Pi\} \\
\chi^{n+1}_w &= \chi^n_w \wedge \boxplus\chi^n_{\Sigma(w)} \wedge \bigwedge\{\neg\boxplus\chi^n_\Pi \colon \Pi \subseteq \Sigma(w),\, \Pi \not\sim^n \Sigma(w)\}
\end{aligned}$$

These formulae are of the required modal depth; the conjunctions and disjunctions in the definition are well-defined since, for a given n, there are only finitely many distinct formulae of the form χ^n_w, and analogously for χ^n_s or χ^n_Π (indeed, it is easy to check that, for finite $\mathcal{P}$, InqML_n is finite up to logical equivalence). We can then prove by simultaneous induction on n that these formulae satisfy the following properties:

1. $\mathbb{M}', w' \models \chi^n_{\mathbb{M},w} \iff \mathbb{M}', w' \sim^n \mathbb{M}, w$
2. $\mathbb{M}', s' \models \chi^n_{\mathbb{M},s} \iff \mathbb{M}', s' \sim^n \mathbb{M}, t$ for some $t \subseteq s$
3. $\mathbb{M}', s' \models \chi^n_{\mathbb{M},\Pi} \iff \mathbb{M}', s' \sim^n \mathbb{M}, s$ for some $s \in \Pi$

First let us show that, if claim (1) holds for a certain $n \in \mathbb{N}$, then the claims (2) and (3) hold for n as well.

For claim (2), suppose $\mathbb{M}', s' \models \chi^n_{\mathbb{M},s}$, that is, suppose $\mathbb{M}', s' \models \bigvee\{\chi^n_{\mathbb{M},w} \colon w \in s\}$. This requires that for any $w' \in s'$ we have $\mathbb{M}', w' \models \chi^n_{\mathbb{M},w}$ for some $w \in s$. By (1), this means that any world in s' is n-bisimilar to some world in s. Letting t be the set of worlds in s that are n-bisimilar to some world in s', we have $t \subseteq s$ and $\mathbb{M}', s' \sim_n \mathbb{M}, t$. Conversely, suppose $\mathbb{M}', s' \sim_n \mathbb{M}, t$ for some $t \subseteq s$. Then every $w' \in s'$ is n-bisimilar to some $w \in s$. By (1), this means that $\mathbb{M}', w' \models \chi^n_{\mathbb{M},w}$, which implies $\mathbb{M}', w' \models \chi^n_{\mathbb{M},s}$. Since this holds for any $w' \in s'$, and since $\chi^n_{\mathbb{M},s}$ is a truth-conditional formula (by Proposition 2.11), it follows that $\mathbb{M}', s' \models \chi^n_{\mathbb{M},s}$.

For claim (3), suppose $\mathbb{M}', s' \models \chi^n_{\mathbb{M},\Pi}$. This implies $\mathbb{M}', s' \models \chi^n_{\mathbb{M},s}$ for some $s \in \Pi$. By claim (2) we have $\mathbb{M}', s' \sim_n \mathbb{M}, t$ for some $t \subseteq s$. Since Π is downward closed, $t \in \Pi$. Conversely, suppose $\mathbb{M}', s' \sim_n \mathbb{M}, t$ for some $t \in \Pi$. By (2), $\mathbb{M}', s' \models \chi^n_{\mathbb{M},t}$, and since $t \in \Pi$, also $\mathbb{M}', s' \models \chi^n_{\mathbb{M},\Pi}$.

Next, we use these facts to show that claim (1) holds for all $n \in \mathbb{N}$, by induction on n. The claim $\mathbb{M}', w' \models \chi^0_{\mathbb{M},w} \Leftrightarrow \mathbb{M}', w' \sim_0 \mathbb{M}, w$ follows immediately from the definition of $\chi^0_{\mathbb{M},w}$. Now assume that claim (1), and thus also claims (2) and (3), hold for n, and let us consider the claim for $n+1$.

For the right-to-left direction, suppose $\mathbb{M}', w' \sim_{n+1} \mathbb{M}, w$. We want to show that $\mathbb{M}', w' \models \chi^{n+1}_{\mathbb{M},w}$. This amounts to showing that: (i) $\mathbb{M}', w' \models \chi^n_{\mathbb{M},w}$; (ii) $\mathbb{M}', w' \models \boxplus\chi^n_{\mathbb{M},\Sigma(w)}$; (iii) $\mathbb{M}', w' \models \neg \boxplus \chi^n_{\mathbb{M},\Pi}$ when $\Pi \subseteq \Sigma(w)$ and $\Pi \not\rightleftharpoons_n \Sigma(w)$. Let us show each in turn.

(i) $\mathbb{M}', w' \sim_{n+1} \mathbb{M}, w$ implies $\mathbb{M}', w' \sim_n \mathbb{M}, w$, so by the induction hypothesis $\mathbb{M}', w' \models \chi^n_{\mathbb{M},w}$.

(ii) Take $s' \in \Sigma'(w')$. Since $\mathbb{M}', w' \sim_{n+1} \mathbb{M}, w$ we must have $\mathbb{M}', s' \sim_n \mathbb{M}, s$ for some $s \in \Sigma(w)$. By the induction hypothesis, $\mathbb{M}', s' \models \chi^n_{\mathbb{M},\Sigma(w)}$. This holds for all $s' \in \Sigma'(w')$, and so $\mathbb{M}', w' \models \boxplus\chi^n_{\mathbb{M},\Sigma(w)}$.

(iii) Suppose for a contradiction that for some $\Pi \subseteq \Sigma(w)$, $\Pi \not\rightleftharpoons_n \Sigma(w)$ and $\mathbb{M}', w' \models \boxplus\chi^n_{\mathbb{M},\Pi}$. This means that every $s' \in \Sigma'(w')$ supports $\chi^n_{\mathbb{M},\Pi}$ and thus, by our induction hypothesis, is n-bisimilar to some $s \in \Pi$. Since $\Pi \subseteq \Sigma(w)$ and $\Pi \not\rightleftharpoons_n \Sigma(w)$, there must be a state $t \in \Sigma(w)$ which is not n-bisimilar to any $s \in \Pi$. But since any state $s' \in \Sigma'(w')$ is n-bisimilar to some $s \in \Pi$, this means that t is not n-bisimilar to any $s' \in \Sigma'(w')$. Since $t \in \Sigma(w)$, this contradicts the assumption that $\mathbb{M}', w' \sim_{n+1} \mathbb{M}, w$.

This establishes the right-to-left direction of the claim. For the converse, suppose $\mathbb{M}', w' \models \chi^{n+1}_{\mathbb{M},w}$. To prove $\mathbb{M}', w' \sim_{n+1} \mathbb{M}, w$, we must show that: (i) w' and w coincide on atomic formulae; (ii) any $s' \in \Sigma'(w')$ is n-bisimilar to some $s \in \Sigma(w)$; and (iii) any $s \in \Sigma(w)$ is n-bisimilar to some $s' \in \Sigma'(w')$.

(i) Since $\chi^n_{\mathbb{M},w}$ is a conjunct of $\chi^{n+1}_{\mathbb{M},w}$, by the induction hypothesis we have $\mathbb{M}', w' \sim_n \mathbb{M}, w$, which implies that w and w' make true the same atomic formulae.

(ii) Since $\boxplus\chi^n_{\mathbb{M},\Sigma(w)}$ is a conjunct of $\chi^{n+1}_{\mathbb{M},w}$, $\mathbb{M}', w' \models \boxplus\chi^n_{\mathbb{M},\Sigma(w)}$. This implies that any $s' \in \Sigma'(w')$ supports $\chi^n_{\mathbb{M},\Sigma(w)}$. By ind. hypothesis, this means that any $s' \in \Sigma'(w')$ is n-bisimilar to some $s \in \Sigma(w)$.

(iii) Let Π be the set of states in $\Sigma(w)$ which are n-bisimilar to some $s' \in \Sigma'(w')$. Now, consider any $s' \in \Sigma'(w')$. We have already seen that s' is n-bisimilar to some state $s \in \Sigma(w)$, which must then be in Π by definition. By induction hypothesis, the fact that s' is n-bisimilar to some state in Π implies $\mathbb{M}', s' \models \chi^n_{\mathbb{M},\Pi}$. And since this is true for each $s' \in \Sigma'(w')$, we have $\mathbb{M}', w' \models \boxplus\chi^n_{\mathbb{M},\Pi}$. Now suppose towards a contradiction that some $s \in \Sigma(w)$ were not n-bisimilar to any state in $\Sigma'(w')$. Then, s would not be n-bisimilar to any state in Π either. This would mean that $\Pi \not\rightleftharpoons_n \Sigma(w)$, which means that $\neg \boxplus \chi^n_{\mathbb{M},\Pi}$ is a conjunct of $\chi^{n+1}_{\mathbb{M},w}$. But then, since $\mathbb{M}', w' \models \chi^{n+1}_{\mathbb{M},w}$, we should have $\mathbb{M}', w' \models \neg \boxplus \chi^n_{\mathbb{M},\Pi}$ contrary to what we found above.

This completes the proof of Proposition 4.2. $\square$

We can now use the properties of characteristic formulae to prove the non-trivial direction of Theorem 4.1.

Proof of Theorem 4.1. We only prove the left-to-right direction, since the converse follows from the observation that InqML-formulae of depth up to n are invariant under n-bisimilarity. For item (i), suppose $\mathbb{M}, s \not\rightleftharpoons_n \mathbb{M}', s'$: then either of the states s and s' is not n-bisimilar to any subset of the other. Without loss of generality, suppose it is s'. By the property of the formula $\chi^n_{\mathbb{M},s}$ we have $\mathbb{M}, s \models \chi^n_{M,s}$ but $\mathbb{M}', s' \not\models \chi^n_{\mathbb{M},s}$. Since the modal depth of $\chi^n_{\mathbb{M},s}$ is n, this shows that $\mathbb{M}, s \not\equiv^n_{\text{InqML}} \mathbb{M}', s'$. The proof for item (ii) is immediate: if $\mathbb{M}, w \not\rightleftharpoons^n \mathbb{M}', w'$ then $\mathbb{M}, w \models \chi^n_{\mathbb{M},w}$ but $\mathbb{M}', w' \not\models \chi^n_{\mathbb{M},w}$; since the modal depth of $\chi^n_{\mathbb{M},w}$ is n, it follows that $\mathbb{M}, w \not\equiv^n_{\text{InqML}} \mathbb{M}', w'$. $\square$

Let us say that a class $\mathcal{C}$ of world-pointed (state-pointed) models is *defined* by a formula φ if $\mathcal{C}$ is the set of world-pointed models where φ is true (state-pointed models in which φ is supported). Theorem 4.1 then yields the following characterisation of InqML-definable classes.

Corollary 4.3. *A class $\mathcal{C}$ of world-pointed models is definable in* InqML *if and only if it is closed under $\sim^n$ for some $n \in \mathbb{N}$. A class $\mathcal{C}$ of state-pointed models is definable in* InqML *if and only if it is both downward closed and closed under $\sim^n$ for some $n \in \mathbb{N}$.*

Proof. If a class $\mathcal{C}$ of world- or state-pointed models is defined by a formula φ of depth n, then since φ is invariant under n-bisimilarity, $\mathcal{C}$ is closed under $\sim^n$. Moreover, if $\mathcal{C}$ is a class of state-pointed models, it is downward closed by the persistency of InqML (Proposition 2.7).

Conversely, suppose that $\mathcal{C}$ is a class of world-pointed models closed under $\sim^n$. Using Proposition 4.2 it is easy to show that $\mathcal{C}$ is defined by the formula $\chi^{\mathcal{C}}_n := \bigvee\{\chi^n_{\mathbb{M},w} \colon (\mathbb{M}, w) \in \mathcal{C}\}$. Notice that the disjunction is well-defined, since for a given n there are only finitely many distinct formulae of the form $\chi^n_{\mathbb{M},w}$. Similarly, if $\mathcal{C}$ is a class of state-pointed models which is both downward closed and closed under $\sim^n$, it follows from Proposition 4.2 that $\mathcal{C}$ is defined by the inquisitive disjunction $\chi^{\mathcal{C}}_n = \mathbin{\rotatebox[origin=c]{180}{$\bigwedge$}}\!\!\!\bigvee\{\chi^n_{\mathbb{M},s} \colon (\mathbb{M}, s) \in \mathcal{C}\}$. $\square$

Remark 4.4. *Notice that the construction of characteristic formulae does not use the modality $\Box$. This implies that $\Box$ can be eliminated from the language of* InqML *without loss of expressive power. This was proved in a more direct way in [3], where it is shown that a formula $\Box\varphi$ can always be turned into an equivalent $\Box$-free formula. However, this translation is not schematic, i.e., there is no $\Box$-free formula $\psi(p)$ such that for every φ, $\Box\varphi \equiv \psi(\varphi)$.*

5 Interlude: InqML and neighbourhood semantics

In neighbourhood semantics for modal logic (see [24] for a recent overview), modal formulae are interpreted with respect to *neighbourhood models*, which are defined as triples $\mathbb{M} = (W, \Sigma, V)$ where W is a set of worlds, $V : \mathcal{P} \to \wp(W)$ is a propositional valuation, and $\Sigma : W \to \wp\wp(W)$, called a *neighbourhood map*, is a function which assigns to each world a set of information states. The standard language of modal logic is interpreted on such models by means of the standard truth-conditional clauses for connectives, and the following clause for modalities:

$$\mathbb{M}, w \models_{\text{nhd}} \Box\varphi \iff |\varphi|_M \in \Sigma(w)$$

where $|\varphi|_M$ is the set of worlds in $\mathbb{M}$ where φ is true. A class of neighbourhood models which is particularly well-studied is that of *monotonic* neighbourhood models [16], which are characterised by the fact that, for all worlds w, the set $\Sigma(w)$ is upward-closed, i.e., closed under supersets.

Clearly, an inquisitive modal model is a special case of neighbourhood model: it is a neighbourhood model such that $\Sigma(w)$ is non-empty and downward closed, i.e., closed under subsets. That is, inquisitive modal models are neighbourhood models which have exactly the opposite monotonicity property than monotonic neighbourhood models have.

In spite of this similarity in models, however, there are profound differences between InqML and neighbourhood semantics, in terms of the logics that arise from these approaches, their expressive power and the induced notions of equivalence.

These differences arise from the way in which the neighbourhood function is used to interpret modal formulae. In neighbourhood semantics, to interpret $\Box\varphi$ we check whether the interpretation of φ is a neighbourhood. The clause for the main modality of InqML, $\boxplus$, is very different: just as in Kripke semantics, we have to check whether φ holds in all successors of the given world — only, these successors are now information states rather than worlds. As a consequence of this, whereas neighbourhood semantics gives rise to non-normal modal logics, the logic of the $\boxplus$ modality in InqML is normal: it validates the K axiom, as well as distribution over conjunction and the necessitation rule.[5]

Besides giving rise to very different modal logics, InqML and neighbourhood semantics are also different, and in fact incomparable, in terms of their expressive power. To see that neighbourhood semantics can draw distinctions that InqML cannot draw, consider the formula $\Box\top$. In neighbourhood semantics, this expresses the property of having the whole universe as a neighbourhood:

$$\mathbb{M}, w \models_{\text{nhd}} \Box\top \iff W \in \Sigma(w)$$

[5] In this discussion, we have set aside the modality $\Box$ of InqML for simplicity since, as remarked above, this modality is definable from $\boxplus$ and the connectives. However, as shown in [3], $\Box$ is also a normal modality in InqML.

This property is clearly not invariant under inquisitive bisimulations (indeed, it is not preserved under disjoint unions!). Therefore, it is not expressible by an InqML-formula.

To see that InqML can also draw distinctions that neighbourhood semantics cannot draw, consider the formula $\boxplus ?p$. This formula expresses the fact that at every neighbourhood of the evaluation world, the truth-value of p is constant.

$$\mathbb{M}, w \models \boxplus ?p \iff \forall s \in \Sigma(w) : s \subseteq |p|_{\mathbb{M}} \text{ or } s \subseteq |\neg p|_M$$

We claim that this property is not expressible in neighbourhood semantics. To see this, consider two models $\mathbb{M}_1$ and $\mathbb{M}_2$ with the same universe $W = \{v, u, u'\}$ and the same valuation $V(p) = \{v\}$. The two models differ in their neighbourhood map, which are both constant, with values

$$\{\{v\}, \{u\}\}^{\downarrow} \text{ for } \Sigma_1 \text{ in } \mathbb{M}_1 \quad \text{versus} \quad \{\{v, u\}\}^{\downarrow} \text{ for } \Sigma_2 \text{ in } \mathbb{M}_2.$$

Given any $w \in W$, we have $\mathbb{M}_1, w \models \boxplus ?p$ but $\mathbb{M}_2, w \not\models \boxplus ?p$. However, it is not hard to show by induction that the set $\{v, u\}$ is not the truth-set of any formula in neighbourhood semantics. Using this fact, it follows easily that in this semantics we have $\mathbb{M}_1, w \models_{\text{nhd}} \varphi \iff \mathbb{M}_2, w \models_{\text{nhd}} \varphi$ for all formulae φ. Hence, the property expressed by $\boxplus ?p$ in InqML is not expressed by any formula in neighbourhood semantics.

Clearly, since InqML and neighbourhood semantics are sensitive to different features of a model, the appropriate notion of bisimilarity is also different in these two contexts. E.g., consider again the above models $\mathbb{M}_1$ and $\mathbb{M}_2$: according to the notion of bisimilarity $\sim^{\mathrm{N}}$ appropriate for neighbourhood semantics [17], the relation $R = \{(v, v), (u, u), (u, u'), (u', u), (u', u')\}$ is a bisimulation. This implies that $\mathbb{M}_1, v \sim^{\mathrm{N}} \mathbb{M}_2, v$. By contrast, a single round of the inquisitive bisimulation game suffices to show that $\mathbb{M}_1, v \not\sim \mathbb{M}_2, v$ in our setting.

Conversely, under our notion of bisimulation, a point w in a model $\mathbb{M}$ is always fully bisimilar to its copy in the disjoint union $\mathbb{M} \uplus \mathbb{M}'$. Clearly, the same cannot be true in neighbourhood semantics, given that in this semantics modal formulae are not in general preserved under disjoint unions.

A notion of bisimulation which is much closer to the one we study here is found in the literature on monotonic neighbourhood models [16]. In terms of the bisimulation game, the difference between the two notions can be described as follows. Starting from a world-position (w, w'), Player **I** picks a state s in either $\Sigma(w)$ or $\Sigma'(w')$; Player **II** responds with a state s' on the opposite side. At this point, the two games come apart: in our version of the game, **I** can choose a world from either s or s', while in the version given in [16], **I** is required to pick a world from s'. Imposing such a restriction in our setting would trivialise the game, providing **II** with a universal winning strategy: always pick $s' = \emptyset$.[6]

Interestingly, however, one can show that due to the downward-closure of $\Sigma(w)$, in our setting it would not make a difference (in terms of the resulting

[6] Notice that since $\Sigma(w)$ and $\Sigma'(w')$ are non-empty and downward closed, they always contain $\emptyset$.

notion of bisimilarity) if Player **I** were required to pick a world from the state s that he himself selected in the world-to-state phase. Thus, we could equivalently have presented the game in a form which is the mirror image of the game used in monotonic neighbourhood frames. Clearly, this symmetry reflects the opposite monotonicity constraints that these two logics place on the neighbourhood map.

6 Relational inquisitive models

In the remainder of this paper, our aim will be to compare the expressive power of inquisitive modal logic with that of first-order logic. This, however, is not quite as straightforward as for ordinary modal logic. A standard Kripke model can be identified naturally with a relational structure with a binary accessibility relation R and a unary predicate P_i for the interpretation of each atomic sentence $p_i \in \mathcal{P}$. By contrast, an inquisitive modal model also needs to encode the inquisitive state map $\Sigma : W \to \wp\wp(W)$. This map can be identified with a binary relation $E \subseteq W \times \wp(W)$. In order to view this as part of a relational structure, however, we need to adopt a two-sorted perspective, and view W and $\wp(W)$ as domains of two distinct sorts. This leads to the following notion.

6.1 Relational inquisitive modal models

Definition 6.1. [relational models] For atomic propositions $\mathcal{P} = \{p_i \colon i \in I\}$, a *relational inquisitive modal model* over $\mathcal{P}$ is a relational structure

$$\mathfrak{M} = (W, S, E, \varepsilon, (P_i)_{i \in I})$$

where W, S are non-empty sets related by $E, \varepsilon \subseteq W \times S$, and, for $i \in I$, $P_i \subseteq W$.

With $s \in S$ we associate the set $\underline{s} := \{w \in W \colon w \,\varepsilon\, s\} \subseteq W$ and require the following conditions, which enforce resemblance with inquisitive modal models:

- extensionality: if $\underline{s} = \underline{s}'$, then $s = s'$.
- local powerset: if $s \in S$ and $t \subseteq \underline{s}$, there is an $s' \in S$ such that $\underline{s}' = t$.
- non-emptiness: for every w, $E[w] \neq \emptyset$.
- downward closure: if $s \in E[w]$ and $\underline{s'} \subseteq \underline{s}$, then $t \in E[w]$.

Multi-modal variants are analogously defined, with a relation $E_a \subseteq W \times S$ to encode the inquisitive assignments Σ_a for agent $a \in \mathcal{A}$.[7]

By extensionality, the second sort S of such a relational model can always be identified with a domain of sets over the first sort, namely, $\{\underline{s} \colon s \in S\} \subseteq \wp(W)$. In the following, we will assume this identification and view a relational model as a structure $\mathfrak{M} = (W, S, E, \in, (P_i))$ where $S \subseteq \wp(W)$ and $\in$ is the actual

[7]See Section 9 for the additional constraints on relational inquisitive epistemic models, that are subject to additional constraints of Definition 2.4, viz., factivity and full introspection.

membership relation. We shall therefore also specify relational models by just $\mathfrak{M} = (W, S, E, (P_i))$ when the fact that $S \subseteq \wp(W)$ and the natural interpretation of ε are understood. Notice that, given this identification, the downward closure condition can be stated more simply as: if $s \in E[w]$ and $t \subseteq s$, then $t \in E[w]$.

Notice that a relational model $\mathfrak{M}$ induces a corresponding Kripke structure $\mathfrak{K}(\mathfrak{M}) = (W, R, (P_i)_{i \in I})$, where $R \subseteq W \times W$ is the relation defined as follows:

$$wRw' \iff \text{for some } s \in S : wEs \text{ and } w' \in s,$$

and correspondingly we have $R[w] := \{w' \colon wRw'\} = \bigcup E[w]$ as the natural relational encoding of the map $\sigma \colon w \mapsto \bigcup \Sigma(w)$.

In addition to the above conditions, we might impose other constraints on a relational model $\mathfrak{M}$: in particular, we may require S to be the full powerset of W, or to resemble the powerset from the local perspective of each world $w \in W$.

Definition 6.2. A relational model $\mathfrak{M} = (W, S, E, (P_i))$ is

- *full* if $S = \wp(W)$;
- *locally full* if for all $w \in W$, $\wp(R[w]) \subseteq S$ (i.e., $\wp(\bigcup E[w]) \subseteq S$).

Notice in particular that $\mathfrak{M}$ is locally full whenever S is closed under arbitrary unions, but need not be full if $\bigcup_{w \in W} R[w] \neq W$.

6.2 Relational encoding of inquisitive modal models

The connection between inquisitive modal models and their relational counterparts is not one-to-one. In one direction, a relational model $\mathfrak{M} = (W, S, E, (P_i))$ uniquely determines an inquisitive modal model $\mathfrak{M}^* = (W, \Sigma, V)$ where $\Sigma(w) = E[w]$ and $V(p_i) = P_i$. Notice that the non-emptiness and downward closure conditions on E guarantee that $\mathfrak{M}^*$ is indeed an inquisitive modal model. The passage from relational to plain inquisitive modal models immediately gives semantics to InqML over relational inquisitive models, and supports all natural notions like bisimulation equivalence over these. Since this passage obliterates information about the second sort S, there are in general many different relational models $\mathfrak{M}$ that determine the same inquisitive modal model $\mathbb{M}$. That is, a given inquisitive modal model may have different relational counterparts. Let us call such counterparts the *relational encodings* of $\mathbb{M}$.

Definition 6.3. A *relational encoding* of an inquisitive modal model $\mathbb{M}$ is a relational model $\mathfrak{M}$ with $\mathfrak{M}^* = \mathbb{M}$.

Clearly, two relational counterparts of $\mathbb{M}$ must coincide in terms of their components W, E and P_i. They may, moreover, be partially ordered in terms of inclusion of their second sort domains.

Definition 6.4. Let $\mathfrak{M}_1$ and $\mathfrak{M}_2$ be two relational encodings of $\mathbb{M}$ with second sort domains S_1 and S_2. We say that the encoding $\mathfrak{M}_1$ is more parsimonious than the encoding $\mathfrak{M}_2$ in case $S_1 \subseteq S_2$.

Given an inquisitive modal model $\mathbb{M}$, there are three salient relational encodings of it. The first yields the most parsimonious relational counterpart of $\mathbb{M}$; the second yields the most parsimonious counterpart of $\mathbb{M}$ that is locally full; and the third yields the unique counterpart of $\mathbb{M}$ that is full.

Definition 6.5. [relational encodings] Given an inquisitive modal model $\mathbb{M} = (W, \Sigma, V)$, we define three relational encodings $\mathfrak{M}^{[\cdots]}(\mathbb{M})$ of $\mathbb{M}$, each based on W, and with $wEs \Leftrightarrow s \in \Sigma(w)$, $w \varepsilon s \Leftrightarrow w \in s$ and $P_i = V(p_i)$. The encodings differ in the second sort domain S:

- for $\mathfrak{M}^{\mathsf{rel}}(\mathbb{M})$,the minimal encoding of $\mathbb{M}$:
  $S := \mathrm{image}(\Sigma)$;
- for $\mathfrak{M}^{\mathsf{lf}}(\mathbb{M})$, the minimal locally full encoding of $\mathbb{M}$:
  $S := \{s \subseteq \sigma(w) \colon w \in W\}$;
- for $\mathfrak{M}^{\mathsf{full}}(\mathbb{M})$, the unique full encoding of $\mathbb{M}$:
  $S := \wp(W)$.

To encode state-pointed models $\mathbb{M}, s$ we augment the corresponding S by $\wp(s)$. These definitions generalise in a natural way to the multi-modal case.

6.3 Relational models and first-order logic

A relational inquisitive model supports a two-sorted first-order language having two relation symbols $\mathtt{E}$ and ϵ corresponding to the relations E and ε, respectively, and predicate symbols $\mathtt{P}_i$ for $i \in I$. We use $\mathtt{w}, \mathtt{v}$ as variables for the first sort, and $\mathtt{s}, \mathtt{t}$ as variables for the second sort.[8] Moreover, we make use of two defined binary predicates. The first is simply inclusion, defined in the natural way in terms of ϵ:

$$\mathtt{s} \subseteq \mathtt{t} := \forall \mathtt{w}(\epsilon(\mathtt{w}, \mathtt{s}) \rightarrow \epsilon(\mathtt{w}, \mathtt{t})).$$

The second defined predicate, $\mathtt{e}(\mathtt{w}, \mathtt{t})$, corresponds to the relation $R[w] = t$ (i.e., the relational encoding of the graph of the map σ):

$$\mathtt{e}(\mathtt{w}, \mathtt{t}) := \forall \mathtt{v}(\epsilon(\mathtt{v}, \mathtt{t}) \leftrightarrow \exists \mathtt{s}(\mathtt{E}(\mathtt{w}, \mathtt{s}) \wedge \epsilon(\mathtt{v}, \mathtt{s}))).$$

In terms of this language we can define a pair of standard translations $\mathrm{ST}_{\mathtt{w}}(\varphi)$ of $\mathrm{ST}_{\mathtt{s}}(\varphi)$ of a formula, which capture its truth conditions in a world and its support conditions in an information state, respectively. The definitions of $\mathrm{ST}_{\mathtt{w}}$ and $\mathrm{ST}_{\mathtt{s}}$ are intertwined. In the definitions below we also appeal to an auxiliary map $\mathrm{ST}_{\mathtt{t}}$: the clauses defining $\mathrm{ST}_{\mathtt{t}}$ are understood to be analogous to those defining $\mathrm{ST}_{\mathtt{s}}$, with the roles of $\mathtt{s}$ and $\mathtt{t}$ exchanged.

Definition 6.6. [standard translation of InqML to first-order logic]

[8] Notice that we use different fonts to distinguish object language symbols ($\mathtt{E}, \mathtt{w}, \mathtt{s}, \dots$), in typewriter font, from the corresponding notation for semantic objects ($E, w, s, \dots$), in regular font.

- $\mathrm{ST}_{\mathtt{w}}(p_i) = \mathtt{P}_i(\mathtt{w})$
  $\mathrm{ST}_{\mathtt{s}}(p_i) = \forall \mathtt{w}(\epsilon(\mathtt{w}, \mathtt{s}) \to \mathrm{ST}_{\mathtt{w}}(p_i))$
- $\mathrm{ST}_{\mathtt{w}}(\bot) = \bot$
  $\mathrm{ST}_{\mathtt{s}}(\bot) = \forall \mathtt{w}(\epsilon(\mathtt{w}, \mathtt{s}) \to \mathrm{ST}_{\mathtt{w}}(\bot))$
- $\mathrm{ST}_{\mathtt{w}}(\varphi \wedge \psi) = \mathrm{ST}_{\mathtt{w}}(\varphi) \wedge \mathrm{ST}_{\mathtt{w}}(\psi)$
  $\mathrm{ST}_{\mathtt{s}}(\varphi \wedge \psi) = \mathrm{ST}_{\mathtt{s}}(\varphi) \wedge \mathrm{ST}_{\mathtt{s}}(\psi)$
- $\mathrm{ST}_{\mathtt{w}}(\varphi \mathbin{\vee\!\!\!\vee} \psi) = \mathrm{ST}_{\mathtt{w}}(\varphi) \vee \mathrm{ST}_{\mathtt{w}}(\psi)$
  $\mathrm{ST}_{\mathtt{s}}(\varphi \mathbin{\vee\!\!\!\vee} \psi) = \mathrm{ST}_{\mathtt{s}}(\varphi) \vee \mathrm{ST}_{\mathtt{s}}(\psi)$
- $\mathrm{ST}_{\mathtt{w}}(\varphi \to \psi) = \mathrm{ST}_{\mathtt{w}}(\varphi) \to \mathrm{ST}_{\mathtt{w}}(\psi)$
  $\mathrm{ST}_{\mathtt{s}}(\varphi \to \psi) = \forall \mathtt{t}(\mathtt{t} \subseteq \mathtt{s} \to (\mathrm{ST}_{\mathtt{t}}(\varphi) \to \mathrm{ST}_{\mathtt{t}}(\psi)))$
- $\mathrm{ST}_{\mathtt{w}}(\boxplus\varphi) = \forall \mathtt{s}(\mathtt{E}(\mathtt{w}, \mathtt{s}) \to \mathrm{ST}_{\mathtt{s}}(\varphi))$
  $\mathrm{ST}_{\mathtt{s}}(\boxplus\varphi) = \forall \mathtt{w}(\epsilon(\mathtt{w}, \mathtt{s}) \to \mathrm{ST}_{\mathtt{w}}(\varphi))$
- $\mathrm{ST}_{\mathtt{w}}(\Box\varphi) = \forall \mathtt{s}(\mathtt{e}(\mathtt{w}, \mathtt{s}) \to \mathrm{ST}_{\mathtt{s}}(\varphi))$
  $\mathrm{ST}_{\mathtt{s}}(\Box\varphi) = \forall \mathtt{w}(\epsilon(\mathtt{w}, \mathtt{s}) \to \mathrm{ST}_{\mathtt{w}}(\varphi))$

It is straightforward to verify that the truth-conditions and support-conditions of φ in a model $\mathbb{M}$ correspond, respectively, to the satisfaction conditions for $\mathrm{ST}_{\mathtt{w}}(\varphi)$ and $\mathrm{ST}_{\mathtt{s}}(\varphi)$ in any locally full relational encoding of $\mathbb{M}$.

Proposition 6.7. *Let $\mathbb{M}$ be an inquisitive modal model, $\mathfrak{M}$ a locally full encoding of $\mathbb{M}$, and $\varphi \in$ InqML. For all worlds $w \in W$ and all states $s \in S$:*

(i) $\mathbb{M}, w \models \varphi \iff \mathfrak{M} \models \mathrm{ST}_{\mathtt{w}}(\varphi)[w]$

(ii) $\mathbb{M}, s \models \varphi \iff \mathfrak{M} \models \mathrm{ST}_{\mathtt{s}}(\varphi)[s]$

The assumption that the encoding $\mathfrak{M}$ be locally full is crucial for this result. This is because, if a model is not locally full, then for some $w \in W$ it could be that the state $\sigma(w) = \bigcup \Sigma(w)$ which is involved in determining the truth condition of $\Box\varphi$ is not represented in $\mathfrak{M}$. If so, there will be no state $s \in S$ satisfying $\epsilon(w, s)$, which means that $\mathrm{ST}_{\mathtt{w}}(\Box\varphi)$ will come out as vacuously true at w, regardless of whether or not $\mathbb{M}, w \models \Box\varphi$.

However, even when the encoding $\mathfrak{M}$ is not locally full, preservation still holds for all $\Box$-free formulae, as one can easily verify.

Proposition 6.8. *Let $\mathbb{M}$ be an inquisitive modal model, $\mathfrak{M}$ a relational encoding of $\mathbb{M}$, and $\varphi \in$ InqML a $\Box$-free formula. Then for all worlds $w \in W$ and all states $s \in S$:*

(i) $\mathbb{M}, w \models \varphi \iff \mathfrak{M} \models \mathrm{ST}_{\mathtt{w}}(\varphi)[w]$

(ii) $\mathbb{M}, s \models \varphi \iff \mathfrak{M} \models \mathrm{ST}_{\mathtt{s}}(\varphi)[s]$

Recall that, by Remark 4.4, any formula φ of InqML is equivalent to some $\Box$-free formula φ^*. Combining this with the previous proposition, we have the following corollary.

Corollary 6.9. *Let $\mathbb{M}$ be an inquisitive modal model, $\mathfrak{M}$ a relational encoding of $\mathbb{M}$, and $\varphi \in$ INQML. Then there are first-order formulae $\varphi^\star_{\mathtt{w}} := \mathrm{ST}_{\mathtt{w}}(\varphi^*)$ and $\varphi^\star_{\mathtt{s}} := \mathrm{ST}_{\mathtt{s}}(\varphi^*)$ such that for all worlds $w \in W$ and all states $s \in S$:*

(i) $\mathbb{M}, w \models \varphi \iff \mathfrak{M} \models \varphi^\star_{\mathtt{w}}[w]$

(ii) $\mathbb{M}, s \models \varphi \iff \mathfrak{M} \models \varphi^\star_{\mathtt{s}}[s]$

The corollary allows us to view INQML as a syntactic fragment of first-order logic, INQML $\subseteq \mathsf{FO}$, over the class of all relational inquisitive models, just as standard modal logic ML may be regarded as a fragment $\mathsf{ML} \subseteq \mathsf{FO}$ over Kripke models.

Importantly, however, the class of relational inquisitive modal models is not first-order definable in this framework, since the local powerset condition involves a second-order quantification. In other words, we are dealing with first-order logic over non-elementary classes of intended models.

We remark that all the considerations of this section admit straightforward variations for the multi-modal inquisitive setting, where models are equipped with a family $(\Sigma_a)_{a \in \mathcal{A}}$ of inquisitive assignments, indexed by a set $\mathcal{A}$ of agents.

7 Bisimulation invariance

7.1 Bisimulation invariance as a semantic constraint

Regarding INQML as a fragment of first-order logic (over relational models, in any one of the above classes), we may think of downward closure and $\sim$-invariance as characteristic semantic features of this fragment. The core question is to which extent INQML may express all $\sim$-invariant properties of worlds or information states that are FO-expressible.

In other words, over which classes $\mathcal{C}$ of models, if any, can INQML be characterised as the bisimulation invariant fragment of first-order logic? In short, for what classes $\mathcal{C}$ do we have

$$\text{INQML} \equiv \mathsf{FO}/{\sim} \qquad (\dagger)$$

just as $\mathsf{ML} \equiv \mathsf{FO}/{\sim}$ by van Benthem's theorem?

A comprehensive answer to this question in the basic scenario of (finite or general) relational inquisitive models and locally full relational models will be given in Section 8, while the corresponding issues for classes of inquisitive epistemic models are treated in Section 9. Before that, we discuss some of the underlying model-theoretic concerns more generally, and in particular stress the connection with the all-important classical rôle of first-order compactness, as well as the rôle of non-classical model-theoretic techniques in our present context where we deal with first-order logic over non-elementary classes of structures.

7.2 Bisimulation invariance and compactness

The inquisitive Ehrenfeucht–Fraïssé theorem, Theorem 4.1, implies $\sim$-invariance for all of INQML. By Corollary 4.3 it further implies *expressive completeness* of $\textsc{InqML}_n$ for any $\sim^n$-invariant property of world-pointed models and any downward-closed $\sim^n$-invariant property of state-pointed models. In order to prove (†) in restriction to some particular class $\mathcal{C}$ of relational inquisitive models, it is thus sufficient to show that, for any formula $\varphi(x) \in \mathsf{FO}$ in a single free variable, $\sim$-invariance of $\varphi(x)$ over $\mathcal{C}$ implies $\sim^n$-invariance of $\varphi(x)$ over $\mathcal{C}$ for some finite n.

This may be viewed as a *compactness principle* for $\sim$-invariance of first-order properties, which is non-trivial in the non-elementary setting of relational inquisitive models.

Observation 7.1. *(a) For any class $\mathcal{C}$ of relational inquisitive models, the following are equivalent:*

(i) $\textsc{InqML} \equiv \mathsf{FO}/{\sim}$ *for world properties over $\mathcal{C}$;*

(ii) for FO*-properties of world-pointed models, $\sim$-invariance over $\mathcal{C}$ implies $\sim^n$-invariance over $\mathcal{C}$ for some n.*

(b) Similarly, the following are equivalent:

(i) $\textsc{InqML} \equiv \mathsf{FO}/{\sim}$ *for downward-closed state properties over $\mathcal{C}$;*

(ii) for downward closed FO*-properties of state-pointed models, $\sim$-invariance over $\mathcal{C}$ implies $\sim^n$-invariance over $\mathcal{C}$ for some n.*

Remark 7.2. *The expressive completeness assertion for downward closed state properties (as expressed in (b) above) implies the corresponding assertion for world properties in (a).*

Proof. With the $\sim$-invariant world property expressed by $\varphi(\mathtt{w}) \in \mathsf{FO}$ associate the state property defined by $\varphi'(\mathtt{s}) = \forall \mathtt{w}(\mathtt{w} \in \mathtt{s} \rightarrow \varphi(\mathtt{w}))$, which is downward closed and also $\sim$-invariant. If the state property defined by $\varphi'(\mathtt{s})$ is equivalent to the support semantics of $\psi \in \textsc{InqML}$, then the specialisation of the semantics of ψ to singleton states $\{w\}$ is equivalent to $\varphi(\mathtt{w})$. $\square$

Interestingly, first-order logic does not satisfy compactness in restriction to the (non-elementary) class of relational inquisitive models, as the following example illustrates.

Example 7.3. There is a first-order formula φ in a single free variable of the second sort (information state) which over any relational inquisitive model says of an element s that there are no infinite R-paths included in s. This is because the local power set condition implies that the entire power set $\wp(s)$ is represented in the second sort of the relational model. So the standard monadic second-order formalisation of our property (cf. Example 7.4 below for details) becomes first-order expressible in this setting.

Even more importantly, over the class of *full* relational models, violations of compactness can even be exhibited for $\sim$-invariant formulae.

Example 7.4. Over full relational inquisitive models, the absence of infinite R-paths from the designated world w (i.e., well-foundedness of the converse of R at w) is a first-order definable and $\sim$-invariant property of worlds that is not preserved under $\sim^n$ for any n, hence not expressible in InqML. In particular, two-sorted first-order logic FO and even its $\sim$-invariant fragment $\mathsf{FO}/\sim$ violates compactness over full relational models.

Proof of the claim in Example 7.4. The observation refers to the following first-order property of worlds w in full relational models, which is $\sim$-invariant and (as a well-foundedness assertion) incompatible with compactness:

$$\mathbb{P}(w) := \text{ there is no infinite } R\text{-path from } w.$$

As this property is not $\sim^n$-invariant for any finite level n, it cannot be expressed in InqML, by Corollary 4.3. It is first-order definable over full relational models, because those afford the full expressive power of monadic second-order quantification over the first sort, W, via first-order quantification over the second sort $S = \wp(W)$. The following MSO-formula, which defines $\mathbb{P}(w)$ over the underlying Kripke frame, can therefore be expressed in two-sorted first-order logic over full relational models:

$$\neg\exists X\big(x \in X \wedge \forall y\big(y \in X \rightarrow \exists z(z \in X \wedge Ryz)\big)\big).$$

The violation of compactness is witnessed by the combination of the first-order formula $\varphi(\mathtt{w})$ for $\mathbb{P}$ together with formulae $\psi_n(\mathtt{w})$, for $n \in \mathbb{N}$, saying that any path of length n from w can be extended to length $n+1$. Then every finite subset is satisfiable in a full relational model, but the whole set is not. $\square$

A corresponding well-foundedness property can also also be captured in first-order logic over full relational inquisitive epistemic models for two agents and using one basic proposition. It suffices to describe analogous path properties for paths formed by a strict alternation of R_a- and R_b-edges on a path that alternates between worlds where p is true and where p is false, for some atomic proposition p and two distinct agents $a, b \in \mathcal{A}$. We therefore also get the following.

Remark 7.5. *Compactness fails for $\sim$-invariant first-order formulae also over full relational encodings in the epistemic setting.*

This shows that the analogues of our Theorems 1.2 and 1.3 fail for the class of full relational models or full relational epistemic models: over these classes, there are properties that are FO-definable, $\sim$-invariant, but not definable in InqML. Although this is in sharp contrast with our Theorems 1.2 and 1.3, their failure over full relational models is not too surprising: over such models, FO, unlike InqML, has access to full-fledged monadic second-order quantification.

7.3 A non-classical route to expressive completeness

In all our characterisation theorems to be treated in the following sections, we establish a semantic correspondence

$$\text{InqML} \equiv \mathsf{FO}/\sim \qquad (\dagger)$$

$$
\begin{tikzcd}
\mathfrak{M}, w \arrow[r, dash, "\sim^n"] \arrow[d, dash, "\sim"'] & \mathfrak{M}', w' \arrow[d, dash, "\sim"] \\
\hat{\mathfrak{M}}, \hat{w} \arrow[r, dash, "\equiv_q"'] & \hat{\mathfrak{M}}', \hat{w}'
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
\mathfrak{M}, s \arrow[r, dash, "\sim^n"] \arrow[d, dash, "\sim"'] & \mathfrak{M}', s' \arrow[d, dash, "\sim"] \\
\hat{\mathfrak{M}}, \hat{s} \arrow[r, dash, "\equiv_q"'] & \hat{\mathfrak{M}}', \hat{s}'
\end{tikzcd}
$$

Figure 2: Generic upgrading patterns.

Note that (†) is an assertion about equal expressive power between two systems presented in very different style: while InqML is based on concrete syntax with clearly defined semantics, $\mathsf{FO}/\!\sim$ is defined in terms of the semantic constraint of $\sim$-invariance.[9] As the semantics of InqML is obviously $\sim$-invariant, its semantically faithful translations into FO as discussed in Section 6.3 show the semantic inclusion from left to right in (†). The essence of the equivalence, therefore, is the *expressive completeness* claim for InqML, which corresponds to the inclusion from right to left in (†): that InqML can express any first-order definable property of worlds (or downward closed property of information states) that is invariant under inquisitive bisimulation. By Observation 7.1, this expressive completeness assertion corresponds to a compactness phenomenon that relates $\sim$-invariance to $\sim^n$-invariance for some finite level n, in the non-classical context of non-elementary classes of relational inquisitive models. To establish (†) over a class $\mathcal{C}$ we need to show that, e.g., a first-order formula $\varphi(\mathtt{w})$ whose semantics as a world property is invariant under $\sim$ over the class $\mathcal{C}$, is in fact invariant under $\sim^n$ as a world property over $\mathcal{C}$, for some finite $n = n(\varphi)$ depending on φ. For this there is a general approach that has been successful in a number of similar investigations, starting from an elementary and constructive proof in [21] of van Benthem's characterisation theorem [29] and its finite model theory version due to Rosen [26] (for ramifications of this method, see also [22, 11] and [23]). This approach involves an *upgrading* of a sufficiently high finite level $\sim^n$ of bisimulation equivalence to a finite target level $\equiv_q$ of elementary equivalence, where q is the quantifier rank of φ. Concretely, and in the case of properties of worlds, this amounts to finding, for any world-pointed relational model $\mathfrak{M}, w$ a fully bisimilar pointed model $\hat{\mathfrak{M}}, \hat{w}$ with the property that, if $\mathfrak{M}, w \sim^n \mathfrak{M}', w'$, then $\hat{\mathfrak{M}}, \hat{w} \equiv_q \hat{\mathfrak{M}}', \hat{w}'$. The diagram in Figure 2 shows how $\sim$-invariance of φ, together with its nature as a first-order formula of quantifier rank q, entails its $\sim^n$-invariance: one simply chases the diagram through its lower rung to check that, for φ that is preserved both under $\sim$ and under $\equiv_q$, then $\mathfrak{M}, w \models \varphi$ iff $\mathfrak{M}', w' \models \varphi$.

The reasoning for properties of information states is analogous, using a corresponding upgrading for state-pointed models. At the technical level, we shall mostly restrict the explicit discussion to the more familiar world-pointed scenario, and only mention the necessary variations for the state-pointed case where relevant.

[9] $\sim$-invariance is easily seen to be undecidable as a property of first-order formulae, hence not a syntactic fragment in any reasonable sense.

Any upgrading of the kind we just discussed involves an interesting tension between the very distinct levels of expressiveness of InqML-formulae and FO-formulae. While the latter can, for instance, distinguish worlds w.r.t. finite branching degrees of the accessibility relation R or w.r.t. short cycles that R may form in the vicinity of a world, no $\sim$-invariant logic can. The challenge is to overcome this discrepancy in bisimilar companion structures, using the malleability up to $\sim$ of relational inquisitive models (within the respective class!) — and, for instance, to boost all multiplicities and lengths of all cycles beyond what can be distinguished in FO_q (FO up to quantifier rank q).

In the next sections, we show how to achieve the required upgradings for various classes $\mathcal{C}$ of relational models to establish our two main lines of characterisation theorems:

(1) for the class of all (finite) relational models, as well as the class of (finite) locally full models, in Section 8, leading up to Theorem 1.2;

(2) for the class of (finite) relational epistemic models, as well as the class of (finite) locally full epistemic models, in Section 9, leading up to Theorem 1.3.

For (1), we use a variation on an upgrading technique from [21] to instantiate the above more general idea; this is based on an inquisitive analogue of partial tree unfoldings; after this pre-processing, the models involved support locality arguments for first-order Ehrenfeucht–Fraïssé games (in effect we shall deviate slightly form the generic picture in Figure 2 by interleaving $\sim$-preserving pre-processing steps and $\equiv_q$-preserving steps).

The classes of models in (2), on the other hand, do not allow for simple partial unfoldings and require a more sophisticated analysis; in particular, some features of monadic second-order logic need to be taken more seriously, features that come into play through the presence of the second sort.

Essentially disjoint unions and localilty. Inquisitive bisimulation between world- or state-pointed inquisitive models is robust under the augmentation of the set of worlds by disconnected sets of new worlds. This phenomenon is well known from ordinary bisimulation between Kripke structures. But whereas the disjoint union of two Kripke models is again a Kripke model, any two relational inquisitive models will necessarily share at least the empty information state in their second sorts (they will share no other elements in either sort if the underlying sets of worlds are disjoint).

Disregarding the universally shared empty information state, which plays a trivial rôle in every semantic respect, we shall therefore speak of *essentially disjoint* unions of models (or of subsets of their domains) if the corresponding unions are disjoint with the exception of the necessary identification between the empty information states in the second sort.

Towards the assessment of the expressive power of FO over relevant classes of relational inquisitive models, which are not elementary, we cannot rely on classical compactness arguments. Instead we invoke *locality arguments* based on the local nature of first-order logic over relational structures, in terms of *Gaifman distance* (cf. Definition 9.13). In the setting of inquisitive relational models, Gaifman distance is graph distance in the undirected bi-partite graph on the sets W of worlds and S of states with edges between any pair linked by E or ε; the *ℓ-neighbourhood* $N^{\ell}(w)$ of a world w consists of all worlds or states at distance up to ℓ from w in this sense, and $N^{\ell}(s)$ is similarly defined. It is easy to see that if $\mathfrak{M}, w$ is a world-pointed relational model and $\ell \neq 0$ is even, the restriction of this model to $N^{\ell}(w)$, denoted $\mathfrak{M} \restriction N^{\ell}(w), w$, is also a world-pointed relational model. Note for this that the empty information state $\emptyset$ is always present in $E[u]$ for all worlds u and hence also in $N^{\ell}(w)$ for $\ell \geqslant 1$, so that the inquisitive state $\{\emptyset\}$ is assigned to all worlds on the cut-off of $\mathfrak{M} \restriction N^{\ell}(w)$.

But the presence of the empty information state $\emptyset \in S$ might seem to spoil any locality-based arguments because it trivialises the distance measure. In fact, the Gaifman diameter of any relational inquisitive model is easily seen to be bounded by 4, as $\emptyset \in S$ is E-related to every world, so that also every information state has distance at most 2 from $\emptyset$. On the other hand, $\emptyset \in S$ plays a trivial rôle not only w.r.t. bisimulation, where it only occurs as a dead end, but also w.r.t. FO expressiveness: the relational model $\mathfrak{M}$ with $\emptyset$ in its second sort is uniformly FO-interpretable in the structure $\mathfrak{M}^{\circ}$ obtained from $\mathfrak{M}$ by dropping $\emptyset$ from the second sort. This observation will play a crucial rôle in some of the technical arguments based on FO locality.

8 Characterisation theorem for InqML

Our aim is to show the following.

THEOREM 1.2. *Let $\mathcal{C}$ be either of the following classes of relational models: the class of all models; of finite models; of locally full models; of finite locally full models. Over each of these classes,* INQML $\equiv \mathsf{FO}/{\sim}$*, i.e., a property of world-pointed models is definable in* INQML *over $\mathcal{C}$ if and only if it is both* FO*-definable over $\mathcal{C}$ and $\sim$-invariant over $\mathcal{C}$. Similarly, a property of state-pointed models is definable in* INQML *over $\mathcal{C}$ if and only if it is* FO*-definable, downward closed, and $\sim$-invariant over $\mathcal{C}$.*

As discussed in connection with Observation 7.1 above, the expressive completeness claim reduces to the following.

Proposition 8.1 (compactness property for $\sim/\sim^{n}$-invariance)**.** *Let $\mathcal{C}$ be either of the following classes of relational models: the class of all models; of finite models; of locally full models; or of finite locally full models. Over each of these classes, any first-order formula whose semantics is $\sim$-invariant is in fact $\sim^{n}$-invariant over that class for some finite level $n \in \mathbb{N}$.*

8.1 Partial unfolding and stratification

Theorem 1.2 boils down to the compactness property expressed in Proposition 8.1 for the relevant classes of relational models. To show this property we make use of a process of *stratification*, comparable to tree-like unfoldings in standard modal logic. We first attend to the relevant constructions for world-pointed models, which will then also support Proposition 8.1 in the case of state-pointed models. The necessary variation will be outlined at the end of this section.

Definition 8.2. We say that a relational inquisitive model $\mathfrak{M}$ is *stratified* if its two domains W and S consist of essentially disjoint[10] strata $(W_i)_{i\in\mathbb{N}}$ and $(S_i)_{i\in\mathbb{N}}$ s.t.

(i) $W = \dot{\bigcup} W_i$ and

(ii) $E[w] \subseteq S_i$ for all $w \in W_i$ and $S_i \subseteq \wp(W_{i+1})$.

For an even number $\ell \neq 0$ and a world w, we say that $\mathfrak{M}$ is *stratified to depth ℓ from w* if $\mathfrak{M} \restriction N^\ell(w)$ is stratified with $W_0 = \{w\}$. $\mathfrak{M}$ is said to be *stratified to depth ℓ* from an information state $s \in S$, if $\mathfrak{M} \restriction N^{\ell+1}(s)$ is stratified with $W_0 = \emptyset$, $S_0 = \wp(s)$.

We note that no non-trivial stratified model can be full.

Proposition 8.3. *Any world-pointed relational inquisitive model $\mathfrak{M}, w$ is bisimilar to a stratified one. For even $\ell \neq 0$, any finite $\mathfrak{M}, w$ is bisimilar to a finite model that is stratified to depth ℓ from w. Similarly for state-pointed relational inquisitive models $\mathfrak{M}, s$. If $\mathfrak{M}$ is locally full, the (ℓ-)stratified target model can be chosen to be locally full, too.*

Proofsketch. The underlying process of partial unfolding is similar to the well-known tree unfolding of Kripke structures, but leaves quite some flexibility as to the choice of the second sort. The stratified domains of the fully stratified target model $\mathfrak{M}', w' \sim \mathfrak{M}, w$ (or $\mathfrak{M}', s' \sim \mathfrak{M}, s$) will consist of $\mathbb{N}$-tagged copies of worlds and information states from $\mathfrak{M}$, so that $W' \subseteq W \times \mathbb{N}$ and $S' \subseteq S \times \mathbb{N}$. In the world-pointed case, let $w' := (w, 0)$. We take $W'_0 := \{(w, 0)\}$. For any $n \in \mathbb{N}$, we then choose a downward closed set $S_n \supseteq \bigcup_{(u,n)\in W'_n} E[u]$ and we let:

$$S'_n := S_n \times \{n\},$$
$$W'_{n+1} := \textstyle\bigcup_{s\in S_n} s \times \{n+1\}.$$

If we define E', ε', and the P'_i as $E' = \{((u,n),(s,n)) \colon (u,s) \in E\}$, $\varepsilon' = \{((u,n+1),(s,n)) \colon u \in s\}$, and $P'_i = \{(u,n) \colon u \in P_i\}$, it is easy to verify that $\mathfrak{M}', w' \sim \mathfrak{M}, w$. In order to maintain finiteness, the unfolding process can be truncated at any stage n if we replace the above W'_{n+1} by W and correspondingly put S instead of S'_{n+1} and augment E' by all of E. The resulting $\mathfrak{M}', w'$ still is fully bisimilar to $\mathfrak{M}, w$, is finite if $\mathfrak{M}$ is, and is stratified to depth $2n$. With the straightforward maximal choice for the S'_n, viz. $S'_n := S \times \{n\}$, the (full or truncated) unfolding process preserves local fullness, too.

In the state-pointed case, we start out by setting $W'_0 := \emptyset$, $S'_0 := \wp(s) \times \{0\}$ and we then proceed inductively as above. □

[10]See discussion at the end of Section 7 for the notion of 'essential disjointness': here the W_i are disjoint and the S_i are disjoint up to $\emptyset$, which they necessarily share.

$$\begin{tikzcd}
\mathfrak{M}, w \arrow[r, dash, "\sim^n" description] \arrow[d, dash, "\sim" description] & \mathfrak{M}', w' \arrow[d, dash, "\sim" description] \\
\mathfrak{M}_1, w \arrow[r, dotted, dash, "\sim^n" description] \arrow[d, dash, "\equiv_q" description] & \mathfrak{M}'_1, w' \arrow[d, dash, "\equiv_q" description] \\
\mathfrak{M}_0, w \arrow[r, dotted, dash, "\sim^n" description] \arrow[d, dash, "\sim" description] & \mathfrak{M}'_0, w' \arrow[d, dash, "\sim" description] \\
\mathfrak{M} \restriction N^\ell(w), w \arrow[r, dash, "\sim" description] & \mathfrak{M}' \restriction N^\ell(w'), w'
\end{tikzcd}$$

Figure 3: Upgrading pattern for Theorem 1.2/Proposition 8.1.

Observation 8.4. *For relational models $\mathfrak{M}$ and $\mathfrak{M}'$ that are stratified to depth ℓ for some even $\ell \neq 0$, and for $n \geqslant \ell/2$:*

$$\begin{aligned} & \mathfrak{M} \restriction N^\ell(w), w \sim^n \mathfrak{M}' \restriction N^\ell(w'), w' \\ \Rightarrow \; & \mathfrak{M} \restriction N^\ell(w), w \sim \mathfrak{M}' \restriction N^\ell(w'), w'. \end{aligned}$$

Analogously for state-pointed models that are stratified to depth ℓ, in restriction to the $(\ell+1)$-neighbourhoods of their distinguished states.

This is because, due to stratification and cut-off, the n-round game exhausts all possibilities in the unbounded game.

Proof of Theorem 1.2. We present the upgrading argument for the case of world-pointed models, which is closer to the classical intuition. The version for state-pointed models, which is formally the stronger, will be discussed below. Let $\mathcal{C}$ be any one of the classes in the theorem and let $\varphi(x) \in \mathsf{FO}_q$ be $\sim$-invariant as a world property over $\mathcal{C}$. We want to show that φ is $\sim^n$-invariant over $\mathcal{C}$ for $n = 2^q$, where q is the quantifier rank of φ. The upgrading argument is sketched in Figure 3. Towards its ingredients, consider a world-pointed relational model $\mathfrak{M}, w$ in $\mathcal{C}$. Since φ is $\sim$-invariant, we can, by Proposition 8.3, assume w.l.o.g. that $\mathfrak{M}, w$ is stratified to depth $\ell = n$. We define two world-pointed models $\mathfrak{M}_0, w$ and $\mathfrak{M}_1, w$ as follows. Each of the $\mathfrak{M}_i$ consists of an essentially disjoint union of the following constituents: both models contain q distinct isomorphic copies of $\mathfrak{M}$ as well as of $\mathfrak{M} \restriction N^\ell(w)$. In addition, $\mathfrak{M}_0$ contains a copy of $\mathfrak{M} \restriction N^\ell(w)$ with the distinguished world w, while $\mathfrak{M}_1$ contains a copy of $\mathfrak{M}$ with the distinguished world w.[11]

$$\begin{aligned} \mathfrak{M}_0, w &:= \; q \otimes \mathfrak{M} \;\oplus\; \mathfrak{M} \restriction N^\ell(w), w \;\oplus\; q \otimes \mathfrak{M} \restriction N^\ell(w) \\ \mathfrak{M}_1, w &:= \; q \otimes \mathfrak{M} \;\oplus\; \mathfrak{M}, w \;\oplus\; q \otimes \mathfrak{M} \restriction N^\ell(w) \end{aligned}$$

[11]See discussion at the end of Section 7 for the notion of 'essential disjointness'.

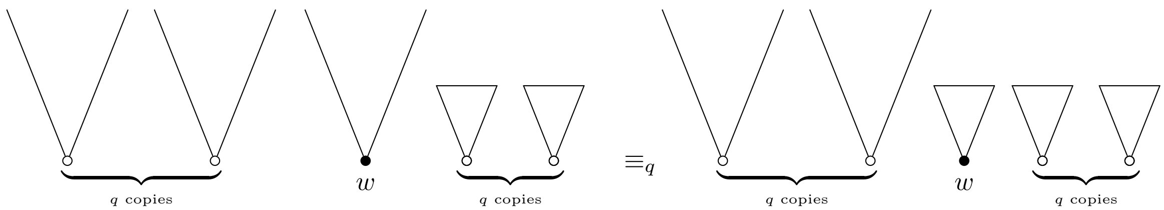

Figure 4: The structures $\mathfrak{M}_1^\circ, w$ and $\mathfrak{M}_2^\circ, w$ in the game argument.

Using a locality-based Ehrenfeucht-Fraïssé game argument for FO we can show that

$$(*) \quad \mathfrak{M}_0, w \equiv_q \mathfrak{M}_1, w.$$

As essentially disjoint sums the $\mathfrak{M}_i$ are disjoint sums up to the identifications of the empty information states $\emptyset$. It is easily checked that the q-equivalence claim in $(**)$ is insensitive to whether the empty information state, which is uniformly present in the second sort of any relational inquisitive model, is represented in the second sort or not. So we may as well work with actual disjoint unions after the removal of the empty information state from the second sort of every component structure. The correspondingly modified structures, which are not admissible as relational inquisitive models as they violate downward closure, are denoted as $\mathfrak{M}_i^\circ$.[12] We therefore argue for

$$(**) \quad \mathfrak{M}_0^\circ, w \equiv_q \mathfrak{M}_1^\circ, w$$

in order to show $(*)$. The diagram in Figure 4 suggests the arrangement, with open cones for copies of $\mathfrak{M}^\circ$ and truncated cones for $\mathfrak{M}\restriction N^\ell(w)^\circ$ and with filled circles for the distinguished worlds.

At the heart of this claim is the following game argument.

We argue that the second player has a winning strategy in the classical q-round Ehrenfeucht–Fraïssé game over the two structures in $(**)$ starting in the position with a single pebble on the distinguished world w on either side. Indeed, the second player can force a win by maintaining the following invariant w.r.t. the game positions $(\mathbf{u}; \mathbf{u}')$ for $\mathbf{u} = (u_0, u_1, \dots, u_m)$ with $u_0 = w$ in $\mathfrak{M}_0^\circ$ and $\mathbf{u}' = (u_0', u_1', \dots, u_m')$ with $u_0' = w$ in $\mathfrak{M}_1^\circ$ after round m, for $m = 0, \dots, q$, for $\ell_m := 2^{q-m}$:

> $\mathbf{u}$ and $\mathbf{u}'$ are partitioned into clusters of matching sub-tuples such that the distance between separate clusters is greater than ℓ_m and

[12] A formal argument can be based either on the observation that $\mathfrak{M}$ and $\mathfrak{M}^\circ$ are uniformly FO-interpretable in one another, or that for every $\varphi(x) \in$ FO there is a $\varphi^\circ(x) \in$ FO (even of the same quantifier rank) such that $\mathfrak{M}, w \models \varphi$ iff $\mathfrak{M}^\circ, w \models \varphi^\circ$.

corresponding clusters are in isomorphic configurations of isomorphic component structures of $\mathfrak{M}_0^\circ$ and $\mathfrak{M}_1^\circ$ or in isomorphic configurations in $\mathfrak{M}_0^\circ \restriction N^\ell(w)$ and $\mathfrak{M}_1^\circ \restriction N^\ell(w)$.

This condition is satisfied at the start of the game, for $m = 0$. The second player can maintain this condition through a round, say in the step from m to $m+1$, as follows. Suppose the first player puts a pebble in position $u = u_{m+1}$ in $\mathfrak{M}_0^\circ$ or $u' = u'_{m+1}$ in $\mathfrak{M}_1^\circ$ at distance up to ℓ_{m+1} of one of the level m clusters (it cannot fall within distance ℓ_{m+1} of two distinct clusters, since the distance between two distinct clusters from the previous level is greater than $\ell_m = 2\ell_{m+1}$); then this new position joins a sub-cluster of that cluster and its match is found in an isomorphic position relative to the matching cluster. If the first player puts the new pebble in a position $u = u_{m+1}$ in $\mathfrak{M}_0^\circ$ or $u' = u'_{m+1}$ in $\mathfrak{M}_1^\circ$ at distance greater than ℓ_{m+1} of each one of the level m clusters, this position will form a new cluster and can be matched with an isomorphic position in one of the as yet unused component structures on the opposite side.

This argument restricts naturally to the scenarios of (finite or general) locally full relational inquisitive structures, because stratification (to some depth) according to Proposition 8.3 preserves local fullness, and so does restriction to some even depth and the formation of essentially disjoint sums.

Given any two pointed models $\mathfrak{M}, w \sim^n \mathfrak{M}', w'$ in any of the relevant classes $\mathcal{C}$, we see that a first-order formula φ of quantifier rank q that is preserved under $\sim$, is preserved by chasing the diagram in Figure 3 along the path through the auxiliary models, which are all in $\mathcal{C}$. The expressive completeness claim for Theorem 1.2, i.e. expressibility of φ in InqML over $\mathcal{C}$, now follows from Corollary 4.3: indeed, φ is logically equivalent over $\mathcal{C}$ to the disjunction over the characteristic formulae $\chi^n_{\mathbb{M},w}$ for all $\mathfrak{M}, w \in \mathcal{C}$ that satisfy φ. $\square$

The case of state properties. For Proposition 8.1 in the case of state properties, we can similarly upgrade the situation $\mathfrak{M}, s \sim^n \mathfrak{M}', s'$ in companion structures through passage to truncations of fully bisimilar models that are stratified to depth ℓ from their distinguished states s. Assuming w.l.o.g. that $\mathfrak{M}, s$ is itself stratified to depth $\ell = 2^q$, we define as before the following essentially disjoint unions

$$\begin{array}{rl} \mathfrak{M}_0, s := & q \otimes \mathfrak{M} \;\oplus\; \mathfrak{M} \restriction N^{\ell+1}(s), s \;\oplus\; q \otimes \mathfrak{M} \restriction N^{\ell+1}(s) \\ \mathfrak{M}_1, s := & q \otimes \mathfrak{M} \;\oplus\; \qquad \mathfrak{M}, s \qquad \;\oplus\; q \otimes \mathfrak{M} \restriction N^{\ell+1}(s) \end{array}$$

and we find that $\mathfrak{M}_0, s \equiv_q \mathfrak{M}_1, s$. We do the same for $\mathfrak{M}', s'$. The rest of the argument for Proposition 8.1 is completed with the straightforward analogue of Figure 3 for the relevant state-pointed models and with $\mathfrak{M} \restriction N^{\ell+1}(s), s \sim \mathfrak{M}' \restriction N^{\ell+1}(s'), s'$ in the bottom rung.

9 Characterisation theorem for epistemic InqML

9.1 Inquisitive epistemic logic and relational models

Inquisitive epistemic frames, or inquisitive S5 frames, are defined as multi-modal inquisitive frames $\mathbb{F} = (W, (\Sigma_a)_{a\in\mathcal{A}})$ that meet the following conditions (see Definition 2.4 for an equivalent formulation):

1. the accessibility relations $R_a = \{(w, w') \in W \times W : w' \in \sigma_a(w)\}$ induced by the maps $\sigma_a : w \mapsto \sigma_a(w) = \bigcup \Sigma_a(w)$, are equivalence relations.
2. Σ_a is constant on each R_a-equivalence class.

In other words, the induced Kripke frame $\mathfrak{K}(\mathbb{F}) = (W, (R_a)_{a\in\mathcal{A}})$ is required to be an $S5$ frame, and for each agent a the inquisitive assignment Σ_a is required to factor w.r.t. R_a-equivalence. In the following we use the term *a-classes* to speak of equivalence classes w.r.t. R_a in inquisitive $S5$ frames as well as in their underlying Kripke $S5$ frames. We denote the a-class of a world w as $[w]_a$. Note that the a-classes are of the form $[w]_a = \sigma_a(w)$ and that on $[w]_a$ the inquisitive assignment Σ_a takes a constant value in $\wp\wp([w]_a) \subseteq \wp\wp(W)$.

The expansion of an inquisitive epistemic frame by a propositional assignment $V : \mathcal{P} \to \wp(W)$ is an inquisitive epistemic model.

Similarly, a relational inquisitive model $\mathfrak{M} = (W, S, (E_a)_{a\in\mathcal{A}}, (P_i)_{i\in I})$ is called an epistemic model, or an S5 model, if it satisfies the following conditions for each $a \in \mathcal{A}$:

- $R_a := \{(w, w') \mid w' \in \bigcup E_a[w]\}$ is an equivalence relation;
- $E_a[w] = E_a[w']$ whenever $(w, w') \in R_a$.

As before, we say that $\mathfrak{M}$ is *full* if $S = \wp(W)$ and *locally full* if, for every world w and each agent a, $\wp(\sigma_a(w)) = \wp([w]_a) \subseteq S$. Notice that, given the downward closure condition, local fullness amounts to the requirement that all the a-classes $[w]_a$ be represented in the second sort.

As in the general setting, we can consider different relational encodings of an inquisitive S5 model, differing only in the choice of the second sort domain. In this setting, locally full relational encodings are a particularly natural choice, since the a-classes $[w]_a$ feature as distinguished second-order objects in inquisitive epistemic $S5$ frames.

9.2 The characterisation theorem

In the remainder of the paper we aim to prove the following characterisation theorem.

THEOREM 1.3. *Let $\mathcal{C}$ be either of the following classes of relational models: the class of all inquisitive epistemic models; of all finite inquisitive epistemic models; of all locally full, or of all finite locally full inquisitive epistemic models. Over each of these classes,* INQML $\equiv \mathsf{FO}/{\sim}$*, i.e., a property of world-pointed models* *is definable in* InqML *over* $\mathcal{C}$ *if and only if it is both* FO*-definable over* $\mathcal{C}$ *and* $\sim$*-invariant over* $\mathcal{C}$*. Similarly, a property of state-pointed models is definable in* InqML *over* $\mathcal{C}$ *if and only if it is* FO*-definable, downward closed, and* $\sim$*-invariant over* $\mathcal{C}$*.*

The expressive completeness claim in the theorem again reduces to the compactness property for $\sim/\sim^\ell$-invariance expressed in the following proposition.

Proposition 9.1 (compactness property for $\sim/\sim^n$-invariance in $S5$ models)**.** *Let* $\varphi(x)$ *be a first-order formula in a single free variable (for worlds or states) whose semantics is* $\sim$*-invariant over one of these classes of relational encodings of inquisitive models: relational inquisitive* $S5$*, finite relational inquisitive* $S5$*, relational locally full inquisitive* $S5$*, or finite relational locally full inquisitive* $S5$*. Then* $\varphi(x)$ *is in fact* $\sim^n$*-invariant over that class, for some finite level* $n \in \mathbb{N}$*.*

The proof of Theorem 1.3 via this proposition will eventually be given at the end of Section 9.5, based on a development of suitable techniques in the following sections. For that development we again focus on the world-pointed scenario, which hold the crucial technical content and may be more familiar from the usual epistemic perspective. The state-pointed version is then obtained, essentially by reduction to the world-pointed case, as discussed at the end of this section. The core argument once more is an upgrading of $\sim^n$-equivalence to $\equiv_q$-equivalence in relational encodings of suitable bisimilar companion structures.

To deal with various uniform constructions of bisimilar companions we use the notion of (globally bisimilar) coverings as expounded e.g. in [22, 11, 23], whose natural adaptation to the setting of multi-modal inquisitive epistemic models is the following.

Definition 9.2. A *bisimilar covering* of an inquisitive structure $\mathbb{M} = (W, (\Sigma_a), V)$ by an inquisitive structure $\hat{\mathbb{M}} = (\hat{W}, (\hat{\Sigma}_a), \hat{V})$ is a map

$$\pi\colon \hat{\mathbb{M}} \longrightarrow \mathbb{M},$$

based on a surjection $\pi\colon \hat{W} \to W$ with natural induced maps $\pi\colon \wp(\hat{W}) \to \wp(W)$ and $\pi\colon \wp\wp(\hat{W}) \to \wp\wp(W)$, such that

- π is compatible with the inquisitive and propositional assignments, i.e. a homomorphism, in the sense that the following diagrams commute

$$\begin{tikzcd} \hat{W} \arrow[r, "\hat{\Sigma}_a"] \arrow[d, "\pi"] & \wp\wp(\hat{W}) \arrow[d, "\pi"] \\ W \arrow[r, "\Sigma_a"] & \wp\wp(W) \end{tikzcd} \qquad \begin{tikzcd} \mathcal{P} \arrow[r, "\hat{V}"] \arrow[dr, "V"'] & \wp(\hat{W}) \arrow[d, "\pi"] \\ & \wp(W) \end{tikzcd}$$

- induces a global bisimulation $\pi\colon \hat{\mathbb{M}} \sim \mathbb{M}$.

This covering is *finite* if its fibres $\pi^{-1}(w) \subseteq \hat{W}$ are finite for all $w \in W$.

We note that an inquisitive bisimulation as in the definition may be thought of as the *lifting*, in the sense of Definition 3.4, of the graph of the world-based homomorphism π.

Towards the desired upgrading argument for Proposition 9.1, a partial unfolding of the global pattern of equivalence classes $[w]_a$ in an inquisitive $S5$ model will be achieved in finite bisimilar coverings in Section 9.5. That covering construction requires a pre-processing at the local level of individual equivalence classes $[w]_a$ w.r.t. the locally constant inquisitive assignment $\Sigma_a \colon [w]_a \to \wp\wp([w]_a)$. This preliminary local transformation concerns a modification of inquisitive assignments Σ_a within $[w]_a$, for individual agents $a \in \mathcal{A}$ and individual a-classes $[w]_a$. The following section is devoted to this a-local pre-processing.

9.3 A local upgrading argument

Local a-structures: a mono-modal situation. It is instructive to consider the very basic case of $S5$ frames and structures with a single agent a, which means that we are dealing with a mono-modal inquisitive frame or structure and need not refer to a. As in the classical modal situation of mono-modal $S5$, this case is rather degenerate also in the inquisitive setting. In particular, up to bisimulation, we may assume that a mono-modal inquisitive $S5$ structure has but one equivalence class, because, in a disjoint union of such classes, disconnected pieces are not mutually accessible through moves in the bisimulation game. Since Σ is constant across the single equivalence class of the distinguished world, moreover, bisimulation equivalence collapses to its second level $\sim^1$. Provided there are only finitely many basic propositions (as there certainly will be in any individual logical formula with finitary syntax), there are just finitely many distinct bisimulation types of such pointed structures, each one characterised by the propositional type of its distinguished world, the collection of propositional types realised overall in its worlds and the collections of propositional types realised in its inquisitive assignment, which is constant across the set of its worlds.

The natural restriction-cum-reduct of any multi-modal inquisitive $S5$ structure to any one of its a-classes, $\mathbb{M} \restriction [w]_a = ([w]_a, \Sigma_a \restriction [w]_a, V \restriction [w]_a)$, which ignores the inquisitive structure induced by agents $b \neq a$, falls into one of these $\sim^1$ equivalence classes. The constant function $\Sigma_a \restriction [w]_a$ may moreover be identified with its constant value $\Sigma_a(w) \in \wp\wp([w]_a)$. All the remaining complexity of multi-agent inquisitive $S5$ structures arises from the overlapping of a-classes for different $a \in \mathcal{A}$, i.e. from the manner in which various $\mathbb{M} \restriction [u]_a$ overlap. In order to fix and analyse this manner in which the substructures on individual a-classes are stitched together, we extract *local structures* $\mathbb{M} \restriction [w]_a$ with a colouring of their worlds induced by their bisimulation types in relation to the global structure $\mathbb{M}$. For different purposes, we use different levels of granularity for this colouring. The finest one of interest is the colouring based on the full bisimulation-type in $\mathbb{M}$, with colour set $C = W/{\sim}$, the set of $\sim$-equivalence classes of worlds in $\mathbb{M}$. Coarser colourings are induced by the $\sim^n$-type, for a fixed finite level n, with colour set $C_n = W/{\sim^n}$. These may also be obtained via the natural projection $\eta\colon C \to C_n$ that identifies all $\sim$-types that fall into the same $\sim^n$-class.

Formally, a colouring of worlds by any colour set C can be encoded as a propositional assignment, for new atomic propositions $c \in C$. Unlike propositional assignments in general, these propositions will be mutually exclusive, so as to partition the set of worlds. To avoid conceptual overhead, we encode colourings as functions

$$\begin{array}{rcl} \rho\colon W & \to & C \\ w & \mapsto & \rho(w) \end{array} \qquad \text{instead of} \qquad \begin{array}{rcl} V\colon C & \to & \wp(W) \\ c & \mapsto & V(c) = \{w\colon \rho(w) = c\} \end{array}$$

and correspondingly write, e.g., $(W, (\Sigma_a)_{a\in\mathcal{A}}, \rho_\sim)$ for the inquisitive epistemic model with the colouring $\rho_\sim\colon W \to C$ that assigns to each world its $\sim$-class:

$$\begin{array}{rrcl} \rho_\sim\colon & W & \longrightarrow & C := W/{\sim} \\ & w & \longmapsto & [w]_\sim = \{u \in W\colon \mathbb{M}, u \sim \mathbb{M}, w\} \end{array}$$

The coarser colouring based on $\sim^n$-types is correspondingly formalised in the inquisitive epistemic model $(\mathbb{M}, (\Sigma_a)_{a\in\mathcal{A}}, \rho_{\sim^n})$, with the colouring $\rho_{\sim^n}\colon W \to C_n = W/{\sim^n}$. Clearly $\rho_\sim$ refines $\rho_{\sim^n}$, and $\rho_{\sim^m}$ refines $\rho_{\sim^n}$ for $m > n$. In particular, all levels refine $\rho_{\sim^0}$, which determines each world's original propositional type as induced by V in $\mathbb{M}$.

Definition 9.3. For an a-class $[w]_a$ in an inquisitive $S5$ structure $\mathbb{M}$, let the associated *local a-structure* be the mono-modal inquisitive $S5$ structure $\mathbb{M}\restriction[w]_a$ with the propositional assignment induced by $\rho_\sim$ in $\mathbb{M}$, which colours each world in $[w]_a$ according to its $\sim$-type in $\mathbb{M}$.

For $n \in \mathbb{N}$, the *local a-structure of granularity n*, $\mathbb{M}_n \restriction [w]_a$, is the mono-modal inquisitive $S5$ structure with propositional assignment induced by $\rho_{\sim^n}$ in $\mathbb{M}$, which colours each world in $[w]_a$ according to its $\sim^n$-type in $\mathbb{M}$.

Interestingly, a local a-structure can be modified considerably without changing its bisimulation type. We now want to achieve local a-structures

- whose inquisitive assignment Σ_a is as simple as possible within its bisimulation class; this leads to the notion of *simplicity* in Definition 9.6;
- which realise every bisimulation type of worlds with high multiplicity; this leads to the notion of *richness* in Definition 9.4.

The idea for *richness* is a quantitative one, viz., that all bisimulation types that are realised in any information state in $\Sigma_a(w)$ will be realised with high multiplicity in some superset that is also in $\Sigma_a(w)$. It turns out that any finite level of richness can be achieved in finite bisimilar coverings (Lemma 9.5), which essentially just put a fixed number of copies of every world.

The idea of *simplicity*, on the other hand, imposes qualitative constraints on the structure of the family of information states and requires a modification of the local inquisitive assignment, which is not described as a covering. It is meant to *normalise* the downward closed family of information states in $\Sigma_a(w)$ within the boolean algebra $\wp([w]_a)$ of subsets of $[w]_a$ as much as possible while preserving the bisimulation type. It may also be thought of as a local saturation condition imposed on the representation, rather than the epistemic content, of inquisitive assignments.

Definition 9.4. A local a-structure $\mathbb{M}{\restriction}[w]_a$ in $\mathbb{M}$ is *K-rich* for some $K \in \mathbb{N}$ if, for every information state $s \in \Sigma_a(w)$ there is some s' such that $s \subseteq s' \in \Sigma_a(w)$ in which every bisimulation type[13] $c \in W/{\sim}$ that occurs at all in $[w]_a$ occurs with multiplicity at least K: $|c \cap s'| \geqslant K$ or $|c \cap s'| = 0$.

An inquisitive $S5$ structure $\mathbb{M} = (W, (\Sigma_a), V)$ is called *K-rich* if each one of its local a-structures is K-rich, for all $a \in \mathcal{A}$.

Lemma 9.5. *For any $K \in \mathbb{N}$, any inquisitive $S5$ structure $\mathbb{M} = (W, (\Sigma_a), V)$ admits a finite bisimilar covering that is K-rich.*

Proof. It suffices to take the natural product $\mathbb{M} \times [K]$ of $\mathbb{M}$ with the standard K-element set $[K] = \{1, \ldots, K\}$. This results in a model whose set of worlds is $W \times [K]$, whose propositional assignment is induced by the projection $\pi \colon W \times [K] \to W$, and whose inquisitive assignments is $\Sigma_a(w, m) := \{s \subseteq W \times [K] \colon \pi(s) \in \Sigma_a(w)\}$. It is easily checked that π induces a bisimulation and constitutes a bisimilar covering in the sense of Definition 9.2 by $\mathbb{M} \times [K]$, which is K-rich. $\Box$

Towards the desired notion of simplicity observe that bisimilarity is rather robust under changes in the actual composition of the inquisitive assignments $\Sigma_a(w)$ in an inquisitive $S5$ structure $\mathbb{M}$. Let $S := \Sigma_a(w) \subseteq \wp([w]_a)$. By the $S5$ nature of $\mathbb{M}$, Σ_a has constant value S on $[w]_a$ and S is a downward closed collection of subsets of $[w]_a$ such that $w \in \bigcup S = [w]_a$; the only additional invariant imposed on S by $\sim$ is the associated colour set for the colouring $\rho = \rho_\sim$ of worlds by their $\sim$-types in $C = W/{\sim}$:

$$\begin{aligned} \rho(S) &:= \{\rho(s) \colon s \in S\} \subseteq \wp(C) \quad \text{where} \\ \rho(s) &:= \{\rho(v) \colon v \in s\} = \{c \in W/{\sim} \colon c \cap s \neq \emptyset\}. \end{aligned}$$

Note that $\rho(S)$ is downward closed in $\wp(C)$. Moreover, the collection of bisimulation types in the set of worlds $[w]_a$ is fully determined by $\rho(S)$. Up to bisimulation, we may therefore replace S by any inquisitive epistemic assignment $\hat{S} \supseteq S$ for which $\rho_\sim(\hat{S}) = \rho_\sim(S)$. The assignment $\hat{S}$ we choose is the *maximal* one that is compatible with the $\sim$-type of the worlds in $[w]_a$. This $\hat{S}$ is generated, by downward closure, from maximal subsets of $[w]_a$ that realise colour combinations in $\rho_\sim(S)$. More specifically, for every $s \in S$, there is a unique such maximal information state

$$\hat{s} := \bigcup \{c \cap [w]_a \colon c \cap s \neq \emptyset\} = \{v \in [w]_a \colon \rho_\sim(v) \in \rho_\sim(s)\}$$

[13] Recall that these bisimulation types are bisimulation types in $\mathbb{M}$, which are imported into the local a-structure through $\rho_\sim$.

that realises the same colours: $\rho_\sim(\hat{s}) = \rho_\sim(s)$. Putting $\hat{\Sigma}_a(v) := \hat{S}$ for all $v \in [w]_a$, we obtain a variant of the local a-structure that is not only locally bisimilar but compatible with the (colour-coded) relationship of worlds in the local a-structure with the surrounding inquisitive $S5$ structure $\mathbb{M}$, since it preserves the local bisimulation type of each individual world. In other words, the $S5$ structure $\hat{\mathbb{M}}$ obtained from $\mathbb{M}$ by modifying the local a-structure $\mathbb{M} \restriction [w]_a$ by replacing $\Sigma_a(w)$ with $\hat{\Sigma}_a(w)$ preserves global bisimulation equivalence, and indeed the bisimulation type of each individual world, i.e. is consistent with $\rho_\sim$.

In a local a-structure based on a colouring $\rho\colon [w]_a \to C$ (e.g. $\rho_\sim$ in $\mathbb{M}\restriction[w]_a$ or $\rho_{\sim^n}$ in $\mathbb{M}_n \restriction [w]_a$), we call an information state s *ρ-saturated* or *colour-saturated* if it is a union of full colour classes, i.e. if

$$s = \rho^{-1}(\rho(s)).$$

Definition 9.6. A local a-structure $\mathbb{M}\restriction[w]_a$ is *simple* if the inquisitive assignment $\Sigma_a(w)$ for $w \in [w]_a$ is generated by downward closure by a collection of $\rho_\sim$-saturated subsets of $[w]_a$. An inquisitive $S5$ structure $\mathbb{M} = (W, (\Sigma_a), V)$ is called *simple* if each one of its local a-structures is simple.

Observation 9.7. *For a simple local a-structure $\mathbb{M} \restriction [w]_a$, the inquisitive assignment $\Sigma_a \restriction [w]_a$ is fully determined by the colouring $\rho_\sim \restriction [w]_a$ and the local $\sim^1$-type, in the sense that the following are equivalent for any $s \subseteq [w]_a$:*

(i) $s \in \Sigma_a(w)$;

(ii) $\rho(s) \in \rho(\Sigma_a(w))$.

We have established the following.

Observation 9.8. *Every inquisitive $S5$ structure $\mathbb{M} = (W, (\Sigma_a), V)$ is globally bisimilar to a simple inquisitive $S5$ structure $\hat{M} = (W, (\hat{\Sigma}_a), V)$ with the same underlying basic modal $S5$ structure (i.e., the same set W of worlds, the same equivalence relations R_a induced by σ_a or $\hat{\sigma}_a$, and the same propositional assignment V) such that for every $w \in W$, $a \in \mathcal{A}$ and information state $s \in \Sigma_a(w)$ and its $\rho_\sim$-saturated companion $\hat{s} = \bigcup\{c \cap [w]_a \colon c \cap s \neq \emptyset\} \in \hat{\Sigma}_a(w)$:*

(i) $\mathbb{M}, w \sim \hat{\mathbb{M}}, w$;

(ii) $\Sigma_a(w) \subseteq \hat{\Sigma}_a(w)$ *and* $\mathbb{M}, s \sim \hat{\mathbb{M}}, s \sim \hat{\mathbb{M}}, \hat{s}$.

Lemma 9.9. *Let $\mathbb{M}$ be an inquisitive $S5$ structure, $\rho = \rho_\sim$ the colouring by $\sim$-types in $C = M/{\sim}$, with the associated lifts $\rho(s) := \{\rho(w) \colon w \in s\}$ to the level of information states and $\rho(\Pi) := \{\rho(s) \colon s \in \Pi\}$ to the level of inquisitive states. Then the bisimulation type of a world $w \in W$ uniquely determines $\rho(\Sigma_a(w))$ for all $a \in \mathcal{A}$, in the sense that*

$$M, w \sim M, w' \;\Rightarrow\; \rho_\sim(\Sigma_a(w)) = \rho_\sim(\Sigma_a(w')).$$

If $\mathbb{M}$ is simple, and $\mathbb{M}, w \sim \mathbb{M}, w'$, then the Boolean algebras of all $\rho_\sim$-saturated subsets of $[w]_a$ and $[w']_a$ are linked be a unique isomorphism that preserves the colouring, which moreover preserves membership in the inquisitive assignments $\Sigma_a(w)$ and $\Sigma'_a(w')$.

$$\begin{tikzcd}
\mathfrak{M}, w \arrow[r, dash, "\sim^n"] \arrow[d, dash, "\sim"] & \mathfrak{M}', w' \arrow[d, dash, "\sim"] \\
\hat{\mathfrak{M}}, \hat{w} \arrow[r, dash, "\equiv_q"] & \hat{\mathfrak{M}}', \hat{w}'
\end{tikzcd}$$

Figure 5: Upgrading idea in the epistemic setting, for word-pointed models.

Proof. Clearly every colour-saturated information state is uniquely determined by the colour classes that contribute to it, and in a simple structure every $\rho(s)$ for $s \in \Sigma_a(w)$ is represented by a unique colour-saturated member in $\Sigma_a(w)$. The isomorphism that arises from $\mathbb{M}, w \sim \mathbb{M}, w'$ is uniquely determined by the condition that it must relate colour-saturated subsets s and s' precisely if $\rho_\sim(s) = \rho_\sim(s')$. $\square$

Concerning the structure of the Boolean algebra of $\rho_\sim$-saturated subsets it is also useful to remark, for later use, that any two distinct members must differ in at least a full colour class (and this means by a large set of worlds if the underlying local a-structures are rich). Also note that the Boolean algebras in question are finite if the set C of colours is finite, which is in particular the case for finite $\mathbb{M}$.

For the combination of simplicity and richness it is instructive to note that richness is preserved under the transformation that achieves simplicity, and that conversely, an application of the covering by a direct product with the K-element set $[K]$, as we used above to achieve K-richness in Lemma 9.5, does preserve simplicity. In other words, either order of application of the two transformations will do to obtain globally bisimilar companions that are both rich and simple, and finite if starting from a finite inquisitive $S5$ structure.

Observation 9.10. *Every inquisitive $S5$ structure $\mathbb{M}$ admits, for $K \in \mathbb{N}$, finite bisimilar coverings that are K-rich. As a consequence, any inquisitive $S5$ structure $\mathbb{M}$ is globally bisimilar to an inquisitive $S5$ structure that is both K-rich and simple, and finite if $\mathbb{M}$ is.*

In order to deal with the upgrading issue for inquisitive $S5$ structures towards Proposition 9.1, we need to boost a sufficient level of inquisitive bisimulation equivalence $\sim^n$ to a target level of first-order equivalence $\equiv_q$ as in the diagram in Figure 5. If bisimilar companions as in the diagram are available for any pair of n-bisimilar pointed inquisitive structures, whenever $n = n(q)$ is sufficiently large in terms of q, then any first-order formula of quantifier rank q that is preserved under inquisitive bisimulation must be preserved under inquisitive n-bisimulation: the detour via the lower extensions in the diagram demonstrates this, as φ is preserved both under $\sim$ and under $\equiv_q$.

The above preparation sets the stage to view a relational inquisitive epistemic model as a network of interconnected local a-structures, which are linked by shared worlds and whose internal local structure appears manageable despite the higher-order nature of the inquisitive assignment. The desired upgrading argument for Proposition 9.1 according to Figure 5, calls for a strategy in the q-round first-order Ehrenfeucht–Fraïssé game over the two-sorted relational models that encode two suitable pre-processed n-bisimilar models. Their nature as conglomerates of interconnected local a-structures, allows us to largely divide concerns strategically as follows:

Strictly local concerns. moves that involve inquisitive assignments and, more generally, information states in the second sort, can be dealt with locally: since the relevant information states are local to individual local a-structures and of a monadic second-order nature over the corresponding sets of worlds $[w]_a$, we shall look to locally maintain levels of monadic second-order (MSO) equivalence. This is facilitated by the above pre-processing that can guarantee sufficient levels of similarity between the boolean algebras of $\rho_{\sim^n}$-saturated information states.

Local concerns regarding global connectivity. moves in the first sort that touch on new worlds, and may challenge their connectivity in the surrounding model, can be dealt with at the level of the underlying $S5$ Kripke structures, where the availability of suitable responses is governed by Gaifman locality properties of first-order logic, provided the global connectivity pattern of these structures has been suitably pre-processed to locally unclutter the link structure between local a-structures. The relevant pre-processing for this can be based on corresponding covering techniques from [22, 11] as discussed in Section 9.5 below.

The following section provides some technical background on game-based arguments in the relevant MSO and FO-contexts for these two aspects.

9.4 From the Ehrenfeucht–Fraïssé toolbox

MSO-equivalence in the local game. We first turn to the simple and rich local a-structures and their relational encodings. Consider a mono-modal inquisitive $S5$ structure $\mathbb{M} = (W, \Sigma, V)$ with constant inquisitive assignment $\Sigma \colon w \mapsto S \subseteq \wp(W)$ and a propositional assignment V induced by a disjoint C-colouring $\rho \colon W \to C$ relationally encoded by its colour classes

$$P_c = \{w \in W \colon \rho(w) = c\}.$$

By the $S5$ nature of the frame, $W = \bigcup S$ and, by the properties of a colouring, $W = \dot{\bigcup}_{c \in C} P_c$. In light of Observation 9.10 we may also assume that $\mathbb{M}$ is simple and K-rich w.r.t. the given colouring, at some granularity $\sim^n$ and for a suitable level K to be determined. In locally full relational encodings, which in restriction to any local a-structure encode the full powerset, first-order logic has access to every subset of the first sort via quantification over the second sort. In other words, locally in each individual a-structure, we are really dealing with full monadic second-order logic MSO over the first sort, which then also fully covers first-order expressiveness across both sorts.

To prepare for an analysis of the two-sorted relational encodings of local a-structures up to certain levels of first-order equivalence, we therefore consider levels of MSO equivalence in terms of its Ehrenfeucht–Fraïssé game. To compare the sizes of sets up to a critical value d (think of d as a threshold beyond which precise distinctions cease to matter), we write $|P| =_d |P'|$ if $|P| = |P'|$ or $|P|, |P'| \geqslant d$. For tuples $\mathbf{P} = (P_1, \dots, P_k)$ of subsets $P_i \subseteq W$ and $\mathbf{P}' = (P'_1, \dots, P'_k)$ of subsets $P'_i \subseteq W'$, the equivalence

$$[\![\mathbf{P}]\!] =_d [\![\mathbf{P}']\!]$$

means that $|\zeta(\mathbf{P})| =_d |\zeta(\mathbf{P}')|$ for every boolean term ζ. Here *boolean terms* refer to terms in the functional language of Boolean algebras, with binary operations for union and intersection and a unary operation for complementation. The following is then folklore but we indicate the straightforward proof.

Lemma 9.11. *For sets W, W' with tuples $\mathbf{P} = (P_1, \dots, P_k)$ of subsets $P_i \subseteq W$ and $\mathbf{P}' = (P'_1, \dots, P'_k)$, $P'_i \subseteq W'$, and for $d = 2^q$:*

$$[\![\mathbf{P}]\!] =_d [\![\mathbf{P}']\!] \quad \Rightarrow \quad (W, \mathbf{P}) \equiv_q^{\mathsf{MSO}} (W', \mathbf{P}').$$

Proofsketch. The proof is by induction on q. For the induction step assume that $[\![\mathbf{P}]\!] =_{2d} [\![\mathbf{P}']\!]$ and suppose w.l.o.g. that the first player proposes a subset $P \subseteq W$ so that the second player needs to find a response $P' \subseteq W'$ such that $[\![\mathbf{P}P]\!] =_d [\![\mathbf{P}'P']\!]$. Decompose P and its complement $\bar{P}$ into their intersections with the atoms of the boolean algebra generated by $\mathbf{P}$ in $\wp(W)$. Then each part of these partitions of P and $\bar{P}$ can be matched with a subset of the corresponding atom of the boolean algebra generated by $\mathbf{P}'$ in $\wp(W')$ in such a manner that the parts $\zeta(\mathbf{P}) \cap P$ and $\zeta(\mathbf{P}) \setminus P$ match their counterparts $\zeta(\mathbf{P}') \cap P'$ and $\zeta(\mathbf{P}') \setminus P'$ in the sense of $=_d$. This just uses the assumption that $|\zeta(\mathbf{P})| =_{2d} |\zeta(\mathbf{P}')|$. $\square$

It is well known (and easy to show by natural composition arguments for strategies) that $\equiv_q^{\mathsf{MSO}}$ is compatible with (arbitrary, not just binary) disjoint unions. For a representation of $W = \dot{\bigcup}_{c \in C} P_c$ and $W' = \dot{\bigcup}_{c \in C} P'_c$ as disjoint unions and for tuples $\mathbf{P} = (P_1, \dots, P_k)$ of subsets of W and $\mathbf{P}' = (P'_1, \dots, P'_k)$ of W' with corresponding restrictions $\mathbf{P}^c := (P_1 \cap P_c, \dots, P_k \cap P_c)$ to the subsets of $P_c \subseteq W$ and $\mathbf{P}'^c := (P'_1 \cap P'_c, \dots, P'_k \cap P'_c)$ to the subsets of $P'_c \subseteq W'$, it follows that for $d = 2^q$:

$$[\![\mathbf{P}^c]\!] =_d [\![\mathbf{P}'^c]\!] \text{ for each } c \in C \quad \Rightarrow \quad (W, \mathbf{P}) \equiv_q^{\mathsf{MSO}} (W', \mathbf{P}').$$

We can use this decomposition argument to deal with possibly infinite colour sets C in the treatment of the local a-structures of infinite $S5$ structures.

The following corollary transfers levels of MSO-equivalence in single-sorted structures with monadic predicates to levels of FO-equivalence in their two-sorted relational encodings. It also introduces the treatment of a coarsening of levels of bisimulation equivalence as occurs in the necessary passage from $\rho_\sim$ to some $\rho_{\sim^n}$ during the intended upgrading. We shall there need to replace the relational encodings of the local a-structures $\mathbb{M}{\restriction}[w]_a$, which are obtained along the vertical axes in Figure 5 that reflect full bisimulation equivalence, by the relational encodings of the $\mathfrak{M}_\ell \restriction [w]_a$ based on $\sim^n$, which are available across the horizontal axes that link $\mathfrak{M}$ and $\mathfrak{M}'$. We may w.l.o.g. assume that the relational structures $\mathfrak{M}$ and $\mathfrak{M}'$ at hand are the locally full encodings of models $\mathbb{M}$ and $\mathbb{M}'$ respectively (replacing $\mathfrak{M}$ by $\mathfrak{M}^{\mathsf{lf}}(\mathfrak{M}^*)$, the locally full relational encoding of the underlying inquisitive model encoded by $\mathfrak{M}$, if necessary).

The coarsening in question can be formalised through a projection map on the colour set C, which serves to identify colours that become equivalent under the coarser view. More specifically, consider the natural projection η associated with an equivalence relation $\approx$ on the colour set C,

$$\begin{array}{rcl} \eta\colon C & \longrightarrow & C/\!\approx \\ c & \longmapsto & \eta(c) = [c]_\approx = \{c' \in C \colon c' \approx c\}. \end{array}$$

For $\mathbb{M} = (W, \Sigma, \rho)$ with a colouring $\rho\colon W \to C$ we let $\mathbb{M}_\eta$ be the structure with the coarser colouring

$$\begin{array}{rcl} \rho_\eta = \eta \circ \rho\colon W & \longrightarrow & C/\!\approx \\ w & \longmapsto & \eta(\rho(w)). \end{array}$$

The application we have in mind is for $\rho = \rho_\sim$ as in $\mathbb{M} \restriction [w]_a$ and $\rho_\eta = \rho_{\sim^n}$ as in $\mathbb{M}_n \restriction [w]_a$, and their respective locally full relational encodings.

Corollary 9.12. *Consider any two local a-structures $\mathbb{M} \restriction [w]_a$ and $\mathbb{M}' \restriction [w']_a$ and their locally full two-sorted relational encodings $\mathfrak{M}^{\mathsf{lf}}(\mathbb{M} \restriction [w]_a)$ and $\mathfrak{M}^{\mathsf{lf}}(\mathbb{M}' \restriction [w']_a)$. If $\mathbb{M}$ and $\mathbb{M}'$ are both simple and sufficiently rich (in relation to $r \in \mathbb{N}$), then the following holds for any two worlds $v \in [w]_a$ and $v' \in [w']_a$:*

$$\mathbb{M} \restriction [w]_a, v \sim^1 \mathbb{M}' \restriction [w']_a, v' \quad \Rightarrow \quad \mathfrak{M}^{\mathsf{lf}}(\mathbb{M} \restriction [w]_a), v \equiv_r \mathfrak{M}^{\mathsf{lf}}(\mathbb{M}' \restriction [w']_a), v'.$$

Moreover, if again $\mathbb{M}$ and $\mathbb{M}'$ are both simple and sufficiently rich, then the following holds for any two worlds $w \in W$ and $w' \in W'$ and any coarsening $\eta\colon C \to C/\!\approx$ with induced $\mathbb{M}_\eta$ and $\mathbb{M}'_\eta$ and their locally full two-sorted relational encodings $\mathfrak{M}^{\mathsf{lf}}(\mathbb{M}_\eta \restriction [w]_a)$ and $\mathfrak{M}^{\mathsf{lf}}(\mathbb{M}'_\eta \restriction [w']_a)$ and $v \in [w]_a$ and $v' \in [w']_a$:

$$\mathbb{M}_\eta \restriction [w]_a, v \sim^1 \mathbb{M}'_\eta \restriction [w']_a, v' \quad \Rightarrow \quad \mathfrak{M}^{\mathsf{lf}}(\mathbb{M}_\eta \restriction [w]_a) \equiv_r \mathfrak{M}^{\mathsf{lf}}(\mathbb{M}'_\eta \restriction [w']_a).$$

By Observation 9.10, therefore, any two inquisitive $S5$ structures $\mathbb{M} = (W, (\Sigma_a)_{a\in\mathcal{A}}, V)$ and $\mathbb{M}' = (W', (\Sigma'_a)_{a\in\mathcal{A}}, V')$ admit, for and any $r \in \mathbb{N}$, globally bisimilar companions $\hat{\mathbb{M}} \sim \mathbb{M}$ and $\hat{\mathbb{M}}' \sim \mathbb{M}'$, finite if $\mathbb{M}$ and $\mathbb{M}'$ are finite, whose locally full relational encodings satisfy the above transfer within every one of their local a-structures (of prescribed granularity).

Proof of the corollary. Since the claims are entirely local to local a-structures involved, we may w.l.o.g. assume that the structures $\mathbb{M}$ and $\mathbb{M}'$ themselves are mono-modal inquisitive epistemic models — albeit with (propositional or relational encodings of) the colourings ρ and ρ_η that in our intended applications are imported as $\rho_\sim$ and $\rho_{\sim^n}$ from larger surrounding models. So let $\mathbb{M} = (W, \Sigma_a, V)$ and $\mathbb{M}' = (W', \Sigma'_a, V')$ be two local a-structures with propositional assignments corresponding to disjoint C-colourings such that $W = \dot{\bigcup}_{c \in C} P_c$ and $W' = \dot{\bigcup}_{c \in C} P'_c$. Assume also that $\mathbb{M}$ and $\mathbb{M}'$ are simple w.r.t. these colourings and K-rich for $K := 2^{r+2}$.

We first look at the case of finite structures $\mathbb{M}$ and $\mathbb{M}'$, and consequently finite colour set C across both structures combined. The existence of any pair of worlds w and w' such that $\mathbb{M}, w \sim^1 \mathbb{M}', w'$ implies not only that $\mathbb{M}$ and $\mathbb{M}'$ instantiate the same colours $c \in C$, but also that the sets of information states in $\Sigma_a(w)$ and $\Sigma'_a(w')$ are generated by matching finite families of colour-saturated sets.

As pointed out in connection with Lemma 9.9 above, these families generate boolean sub-algebras of the full power set algebras whose atoms are precisely the non-empty colour sets. It follows that for any fixed enumeration of these matching families $(s_j)_{j \in J}$ and $(s'_j)_{j \in J}$ of colour-saturated information states and the families of unary predicates $(P_c)_{c \in C}$ and $(P'_c)_{c \in C}$ that encode the colourings by the richness assumption:

$$[\![(P_c)_{c \in C}, (s_j)_{j \in J}]\!] =_d [\![(P'_c)_{c \in C}, (s'_j)_{j \in J}]\!] \text{ for } d = 2^{r+2},$$

so that the lemma implies

$$(W, (P_c)_{c \in C}, (s_j)_{j \in J}) \equiv^{\mathsf{MSO}}_{r+2} (W', (P'_c)_{c \in C}, (s'_j)_{j \in J}).$$

Since by assumption w and w' are such that they belong to matching colour classes and hence also matching colour-saturated information states, it further follows that

$$(W, (P_c)_{c \in C}, (s_j)_{j \in J}, \{w\}) \equiv^{\mathsf{MSO}}_{r+1} (W', (P'_c)_{c \in C}, (s'_j)_{j \in J}, \{w'\}),$$

which also guarantees that the two-sorted relational encodings are first-order equivalent up to quantifier rank r: the second player has a winning strategy for r rounds in the ordinary first-order game played on these two-sorted relational encodings, since the latter can be simulated in the MSO-game on the first sort. A move involving the second sort (a pebble on an information state) naturally corresponds to a set move in the MSO-game, while a move in the first sort (a pebble on a world) is mimicked by a singleton set move. W.r.t. membership of information states in the image of Σ_a or Σ'_a (i.e., the local restrictions of the relation E in the relational encodings) we observe that, due to simplicity, this issue is reduced to emptiness questions about intersections of these information states with the lists of colour-saturated information states $(s_j)_{j \in J}$ and $(s'_j)_{j \in J}$, cf. Observation 9.7.[14]

For the transition to a coarser assignment on the basis of an identification of colours from C according to a projection $\eta \colon C \to C/{\approx}$, it is clear that richness

[14]The one round to spare in the MSO levels of equivalence serves to ensure that responses respect set inclusion relationships: this concerns set inclusions between singleton sets introduced to mimic pebbles on worlds, i.e., $\in$-relationships, but also set inclusions between pebbled information states and their respective parent sets s or s' and checking for empty or non-empty intersections with the colour sets, which is necessary also to check membership in $\Sigma_a(w)$ and $\Sigma'_a(w')$.

and $=_d$-equivalence as well as simplicity are preserved under the merging of colour classes according to η.

The corresponding arguments for the setting of infinite structures has to deal with infinite sets C of colours, and correspondingly with infinite families $(P_c)_{c\in C}$ and $(P'_c)_{c\in C}$ as well with infinite matching families of colour-saturated information states $(s_j)_{j\in J}$ and $(s'_j)_{j\in J}$. But these families of sets can be decomposed into their disjoint restrictions to the individual colour sets P_c and P'_c that partition W and W', respectively. K-richness for $K = 2^{r+1}$ implies componentwise equivalence in the sense of $=_d$ for each pair of restrictions. In fact the restrictions of each colour-saturated s and matching s' to any individual P_c and P'_c are either both full or both empty, so both of size 0 or both of size $\geqslant K$. So $\equiv^{\mathsf{MSO}}_{r+2}$ follows as before, based on its compatibility with disjoint unions, and the argument for $\equiv_r$ for the pointed variants of the relational encodings remains the same. The transition to a coarsening via η again preserves richness, simplicity and $=_d$-equivalence. $\square$

FO-equivalence and Gaifman locality. We now turn to the global pattern of overlapping local a-classes and structures, again with a view to a partial bisimilar unfolding in a covering that simplifies the overall structure while removing obstacles to levels of first-order equivalence that are not governed by any level of bisimulation equivalence. While details towards this aspect of the upgrading argument will be given in Section 9.5, we here review the relevant technical background on Gaifman locality and its use towards establishing levels $\equiv_q$ of first-order equivalence between relational structures $\hat{\mathfrak{M}}$ and $\hat{\mathfrak{M}}'$ in the situation of Figure 5.

In our case of relational structures with just unary and binary relations, *Gaifman distance* is just ordinary graph distance w.r.t. the symmetrisations of all the binary relations in the vocabulary. It establishes a natural distance measure between elements (in our case across both sorts, worlds and information states).[15] For an element b of the relational structure $\mathfrak{B}$ and $\ell \in \mathbb{N}$, the set of elements at distance up to ℓ from b is the *ℓ-neighbourhood* of b, denoted as

$$N^\ell(b) = \{b' \in B \colon d(b,b') \leqslant \ell\}.$$

A first-order formula $\varphi(x)$ in a single free variable x is *ℓ-local* if its semantics at $b \in \mathfrak{B}$ depends just on the induced substructure $\mathfrak{B} \restriction N^\ell(b)$ on the ℓ-neighbourhood $N^\ell(b)$. In other words, $\varphi(x)$ is ℓ-local if, and only if, it is logically equivalent to its ℓ-local *relativisation* to the ℓ-neighbourhood of x, for which we write $\varphi^\ell(x)$ (which is itself expressible in FO). We refer to $\varphi^\ell(x) \in \mathsf{FO}$ as the *ℓ-localisation* of $\varphi(x)$. A finite set of elements of the relational structure $\mathfrak{B}$ is said to be *ℓ-scattered* if the ℓ-neighbourhoods of any two distinct members

[15] In general, 'infinite distances' occur between elements that are not linked by finite chains of binary relational edges, i.e. that are in distinct connected components. Due to downward closure and the rôle of the empty information state, all relational inquisitive epistemic models are actually connected; we shall technically deal with this artefact below, similar to the treatment in Section 8.1.

are disjoint. A *basic local sentence* is a sentence that asserts, for some formula $\varphi(x)$ and some $m \geqslant 1$, the existence of an ℓ-scattered set of m elements that each satisfy the ℓ-localisation $\varphi^\ell(x)$. Gaifman's theorem, in the case that we are interested in, asserts that any first-order formula in a single free variable x is logically equivalent to a boolean combination of ℓ-local formulae $\varphi^\ell(x)$ for some ℓ and some basic local sentences. We shall focus on the following definition of levels of Gaifman equivalence.

Definition 9.13. Two pointed relational structures $\mathfrak{B}, b$ and $\mathfrak{B}', b'$ are (ℓ, r, m) *Gaifman equivalent*, denoted $\mathfrak{B}, b \equiv^{(\ell)}_{r,m} \mathfrak{B}', b'$ if b and b' satisfy exactly the same ℓ-localisations of formulae $\varphi(x)$ of quantifier rank r in $\mathfrak{B}$ and $\mathfrak{B}'$ and if the structures $\mathfrak{B}$ and $\mathfrak{B}'$ satisfy exactly the same basic local sentences concerning ℓ'-localisations for formulae $\varphi(x)$ of quantifier rank r and ℓ'-scattered sets of size m', for any $\ell' \leqslant \ell$, $m' \leqslant m$.

$\mathfrak{B}, b$ and $\mathfrak{B}', b'$ are *ℓ-locally r-equivalent*, denoted $\mathfrak{B}, b \equiv^{(\ell)}_{r} \mathfrak{B}', b'$ if b and b' satisfy exactly the same ℓ-localisations of formulae $\varphi(x)$ of quantifier rank r.

The following is an immediate corollary of Gaifman's classical theorem [13, 12].

Proposition 9.14. *The semantics of any first-order formula $\psi(x)$ in a purely relational vocabulary is preserved under Gaifman equivalence $\equiv^{(\ell)}_{r,m}$ for sufficiently large values of the parameters $\ell, r, m \in \mathbb{N}$.*

We can therefore replace the target level $\equiv_q$ of FO-equivalence in the upgrading task of Proposition 9.1 as in Figure 5 by a suitable level $\equiv^{(\ell)}_{r,m}$ of Gaifman equivalence that is sufficient to preserve the semantics of the given bisimulation-invariant FO formula φ.

For formulae $\psi(x) \in \mathsf{FO}$ whose semantics is also preserved under disjoint unions of relational structures, preservation under $\equiv^{(\ell)}_{r,m}$ implies preservation under $\equiv^{\ell}_{r}$ in the following sense.

Observation 9.15. *Let $\psi(x) \in \mathsf{FO}$ be invariant under disjoint unions in the sense that for pointed $\mathfrak{B}, b$ and disjoint unions $\mathfrak{B} \oplus \mathfrak{C}$ always*

$$\mathfrak{B}, b \oplus \mathfrak{C} \models \psi \;\Leftrightarrow\; \mathfrak{B}, b \models \psi.$$

Then preservation of ψ under $\equiv^{(\ell)}_{r,m}$ implies preservation under $\equiv^{(\ell)}_{r}$.

Proof. The following simple argument from [22] encapsulates elements of the ad-hoc proof given for Theorem 1.2 in Section 8.1. We show that $\mathfrak{B}, b \equiv^{(\ell)}_{r} \mathfrak{B}', b'$ implies that $\mathfrak{B}, b \models \psi$ iff $\mathfrak{B}', b' \models \psi$. Since ψ is assumed to be invariant under $\equiv^{(\ell)}_{r,m}$, it suffices to observe that $\mathfrak{B}, b \equiv^{(\ell)}_{r} \mathfrak{B}', b'$ implies

$$m \otimes \mathfrak{B} \;\oplus\; \mathfrak{B}, b \;\oplus\; m \otimes \mathfrak{B}' \equiv^{(\ell)}_{r,m} m \otimes \mathfrak{B} \;\oplus\; \mathfrak{B}', b' \;\oplus\; m \otimes \mathfrak{B}',$$

where, e.g., $m \otimes \mathfrak{B}$ stands for the m-fold disjoint union of isomorphic copies of $\mathfrak{B}$, and to use robustness of ψ w.r.t. disjoint unions. $\square$

$$
\begin{tikzcd}
\mathfrak{M}, w \arrow[r, dash, "\sim^n"] \arrow[d, dash, "\sim"] & \mathfrak{M}', w' \arrow[d, dash, "\sim"] \\
\hat{\mathfrak{M}}, \hat{w} \arrow[r, dash, "\sim^n"] \arrow[d, dotted, dash] & \hat{\mathfrak{M}}', \hat{w}' \arrow[d, dotted, dash] \\
\hat{\mathfrak{M}}^\circ, \hat{w} \arrow[r, dash, "\equiv_r^{(\ell)}"] & \hat{\mathfrak{M}}'^\circ, \hat{w}'
\end{tikzcd}
$$

Figure 6: The upgrading in the epistemic setting, refined.

It follows that we may even replace the target level $\equiv_q$ for the upgrading for Proposition 9.1 according to Figure 5 by a suitable level of equivalence $\equiv_r^{(\ell)}$ — which, remarkably, is a level of local FO-equivalence, viz. of $\equiv_r$ in restriction to the ℓ-neighbourhoods of the distinguished worlds. But as the relevant classes $\mathcal{C}$ of inquisitive epistemic relational models do not allow for disjoint unions, we need to adapt the more generic argument from Observation 9.15 as follows. Recall the discussion of essentially disjoint unions and the operation $\mathfrak{M} \mapsto \mathfrak{M}^\circ$ that removes the empty information state from the second sort, as discussed in Section 8.1. For the following compare Figure 6.

Remark 9.16. *In order to show that $\varphi(x) \in \mathsf{FO}$ is preserved under $\sim^n$ for some suitable finite level of n in restriction to any one of the classes $\mathcal{C}$ of relational models from Theorem 1.3, according to Proposition 9.1, it suffices to upgrade $\mathfrak{M}, w \sim^n \mathfrak{M}', w'$ to a suitable level $\equiv_r^{(\ell)}$ of ℓ-local r-equivalence between the variants $\hat{\mathfrak{M}}^\circ, \hat{w}$ and $\hat{\mathfrak{M}}'^\circ, \hat{w}'$ of bisimilar companions $\hat{\mathfrak{M}}, \hat{w} \sim \mathfrak{M}, w$ and $\hat{\mathfrak{M}}', \hat{w}' \sim \mathfrak{M}', w'$ within $\mathcal{C}$. The analogous claim holds true for state-pointed models and $\varphi(x)$ preserved under $\sim^n$ as a state-property over $\mathcal{C}$.*

Proof. Let $\mathcal{C}^\circ := \{\mathfrak{M}^\circ : \mathfrak{M} \in \mathcal{C}\}$ and let $\varphi^\circ(x) \in \mathsf{FO}$ be such that, e.g. for the world-pointed case, for all $\mathfrak{M} \in \mathcal{C}$:

$$\mathfrak{M}, w \models \varphi \text{ iff } \mathfrak{M}^\circ, w \models \varphi^\circ.$$

Let ℓ, r, m be such that (ℓ, r, m)-Gaifman equivalence preserves the semantics of φ°. We note that, due to $\sim$-invariance over the class $\mathcal{C}$, its variant φ° is invariant under actual disjoint unions of models in $\mathcal{C}^\circ$. An application of the argument from Observation 9.15 to $\psi := \varphi^\circ$ over $\mathcal{C}^\circ$ yields the desired result. $\square$

9.5 Partial unfoldings in finite coverings

It remains to provide bisimilar companions for inquisitive epistemic relational models, for which a suitable finite level of $\sim^n$ can be upgraded to given a target level $\equiv_r^{(\ell)}$ of local FO-equivalence according to Remark 9.16. The required preprocessing needs to unclutter and smooth out the local pattern of overlapping local a-structures in such a manner that the ℓ-bisimulation type determines the first-order behaviour of ℓ-neighbourhoods up to quantifier rank r. This requires a local unfolding of the connectivity in the underlying Kripke structures and uses ideas developed in [22] to eliminate incidental cycles and overlaps in finite bisimilar coverings — ideas which have also been applied in the context of plain $S5$ Kripke structures in [11].

Consider the single-sorted, relational multi-modal $S5$ structure $\mathfrak{K} = \mathfrak{K}(\mathbb{M}) = (W, (R_a)_{a\in\mathcal{A}}, V)$ with equivalence relations R_a derived from the inquisitive $S5$ frame $\mathbb{M}$, with equivalence classes $[w]_a$. Generic constructions from [22], which are based on products with suitable Cayley groups, yield finite bisimilar coverings $\pi \colon \hat{\mathfrak{K}} \to \mathfrak{K}$ by another $S5$ structure $\hat{\mathfrak{K}}$ that is N-acyclic for some desired threshold $N \in \mathbb{N}$ in the following sense:

- a_i-classes and a_j-classes for agents $a_i \neq a_j$ can overlap in at most a single world and there is no non-trivial cyclic pattern of length up to N of overlapping a_i-classes.
- every a-class in the finite bisimilar covering is bijectively related to an a-class in the original structure by the projection map of the covering which furthermore induces the bisimulation.

Let $\pi \colon \hat{\mathfrak{K}} \to \mathfrak{K}$ be an N-acyclic finite bisimilar covering in this sense, which in particular means that the graph of the covering projection π is a bisimulation relation. Since a-classes in $\hat{\mathfrak{K}}$ are π-related isomorphic pre-images of corresponding a-classes in $\mathfrak{K}$, we may consistently endow them with an inquisitive Σ_a-assignment as pulled back from $\mathbb{M}$, to obtain a natural derived finite inquisitive bisimilar covering

$$\pi \colon \hat{\mathbb{M}} \longrightarrow \mathbb{M} = (W, (\Sigma_a), V)$$

in the sense of Definition 9.2, where $\hat{\mathbb{M}} = (\hat{W}, (\hat{\Sigma}_a), \hat{V})$ is based on $\hat{W}$ and compatible with the underlying $S5$ Kripke structure $\hat{\mathfrak{K}} = \mathfrak{K}(\hat{\mathbb{M}})$. More specifically, for $\hat{w} \in \hat{W}$ consider its $\hat{R}_a$ equivalence class $[\hat{w}]_a \subseteq \hat{W}$, which is bijectively related by π to the R_a equivalence class of $w = \pi(\hat{w})$ in $\mathfrak{K}$, $[w]_a \subseteq W$. In restriction to $[\hat{w}]_a$ we put

$$\begin{array}{rcl} \hat{\Sigma}_a \restriction [\hat{w}]_a \colon [\hat{w}]_a & \longrightarrow & \wp\wp([\hat{w}]_a) \\ \hat{u} & \longmapsto & \{\pi^{-1}(s) \colon s \in \Sigma_a(w) \subseteq \wp([w]_a)\}, \end{array}$$

where π^{-1} refers to the inverse of the local bijection between $[\hat{w}]_a$ and $[w]_a$ induced by π. In other words, we make the following diagram commute and note that the relevant restriction of π in the right-hand part of the diagram is bijective and part of a bisimulation at the level of the underlying $S5$ Kripke structures, which means in particular that it is compatible with the given propositional assignments:

$$
\begin{tikzcd}
\hat{\mathbb{M}}, \hat{w} \arrow[d, "\pi"'] & {[\hat{w}]_a} \arrow[r, "\hat{\Sigma}_a"] \arrow[d, "\pi"'] & \wp\wp([\hat{w}]_a) \arrow[d, "\pi"] \\
\mathbb{M}, w & {[w]_a} \arrow[r, "\Sigma_a"'] & \wp\wp([w]_a)
\end{tikzcd}
$$

We merely need to check that the resulting model $\hat{\mathbb{M}}$ is again an inquisitive epistemic model, and that π induces an inquisitive bisimulation, not just a bisimulation at the level of the underlying single-sorted, multi-modal $S5$-structures $\hat{\mathfrak{K}} = \mathfrak{K}(\hat{\mathbb{M}})$ and $\mathfrak{K} = \mathfrak{K}(\mathbb{M})$. That $\hat{\mathbb{M}}$ is an inquisitive $S5$ model is straightforward from its construction: $\hat{\Sigma}_a$-values are constant on the a-classes induced by $\hat{R}_a$ just as Σ_a-values are constant across a-classes induced by R_a; and $\hat{w} \in \hat{\sigma}_a(\hat{w})$ follows from the fact that $w \in \sigma_a(w) = [w]_a$. The back&forth conditions for an inquisitive bisimulation relation can be verified for

$$
\begin{aligned}
Z = \{(\hat{u}, \pi(\hat{u})) \in \hat{W} \times W \colon \hat{u} \in \hat{W}\} \\
\cup \{(\hat{s}, \pi(\hat{s})) \in \wp(\hat{W}) \times \wp(W) \colon a \in \mathcal{A}, \hat{w} \in \hat{W}, \hat{s} \in \hat{\Sigma}_a(\hat{w})\},
\end{aligned}
$$

which is the natural lifting of π to information states.

We summarise these findings as follows.

Lemma 9.17. *Any inquisitive $S5$ structure $\mathbb{M} = (W, (\Sigma_a), V)$ admits, for every $N \in \mathbb{N}$, a finite bisimilar covering of the form*

$$\pi \colon \hat{\mathbb{M}} \longrightarrow \mathbb{M}$$

by an inquisitive $S5$ structure $\hat{\mathbb{M}} = (\hat{W}, (\hat{\Sigma}_a), \hat{V})$ such that

- *(i) the global bisimulation induced by π is an isomorphism in restriction to each local a-structure $\hat{\mathbb{M}} \restriction [\hat{w}]_a$ of $\hat{\mathbb{M}}$, which is isomorphically mapped by π onto the local a-structure $\mathbb{M} \restriction [\pi(\hat{w})]_a$ of $\mathbb{M}$,*

and $\hat{\mathbb{M}}$ is N-acyclic in the sense that

- *(ii) no two distinct a-classes overlap in more than a single world, and*
- *(iii) there are no non-trivial cyclic patterns of length up to N formed by overlapping a-classes.*

It is noteworthy that condition (i) guarantees that simplicity and K-richness are preserved in the covering, simply because they are properties of the local a-structures, which up to isomorphism are the same in $\mathbb{M}$ and in $\hat{\mathbb{M}}$. This implies that the conditions imposed on structures $\mathbb{M}$ and $\mathbb{M}'$ in the following lemma can always be achieved simultaneously in finite bisimilar companions of arbitrary inquisitive epistemic models, by using finite N-acyclic coverings according to Lemma 9.17 after pre-processing according to Observation 9.10 for simplicity and richness.

Lemma 9.18. *Let $\mathbb{M}$ and $\mathbb{M}'$ be N-acyclic for $N > 2\ell$, simple and sufficiently rich to support the claim of Corollary 9.12. Then their locally full relational encodings $\mathfrak{M} = \mathfrak{M}^{\mathsf{lf}}(\mathbb{M})$ and $\mathfrak{M}' = \mathfrak{M}^{\mathsf{lf}}(\mathbb{M}')$ satisfy the following for any pair of worlds such that $\mathbb{M}, w \sim^n \mathbb{M}', w'$ for $n \geqslant \ell + 1$:*

$$\mathfrak{M}^\circ \restriction N^\ell(w), w \equiv_r \mathfrak{M}'^\circ \restriction N^\ell(w'), w'.$$

Proof. In the given situation, a strategy for the second second player in the r-round game on the pointed ℓ-neighbourhoods of w and w' in $\mathfrak{M}^\circ$ and $\mathfrak{M}'^\circ$ can be based on a combination of strategies in

(a) the r-round first-order game on the tree-like underlying $S5$ Kripke structures on the first sort (worlds) with a-classes for accessibility and a colouring of worlds that reflect sufficient finite approximations to their bisimulation types in $\mathbb{M}$ that which in particular suffice to guarantee 1-bisimulation equivalence and hence also $\equiv_r$ of local a-structures of appropriate granularities,[16] and

(b) local strategies in the r-round games on the relational encodings of the local a-structures (with colourings at the appropriate level of granularity) induced by pebbled pairs of worlds. These strategies are ultimately based on local MSO-games over the first sort, as discussed in connection with Corollary 9.12.

Consider any position $\mathbf{u}; \mathbf{u}'$ in the r-round first-order game on $\mathfrak{M}^\circ \restriction N^\ell(w)$ and $\mathfrak{M}'^\circ \restriction N^\ell(w')$, starting with pebbles on w and w' which are equivalent w.r.t. $\sim^{\ell+1}$ in $\mathbb{M}$ and $\mathbb{M}'$. It will be of the form $\mathbf{u} = (u_0, \ldots, u_k)$ and $\mathbf{u}' = (u'_0, \ldots, u'_k)$ for some $k \leqslant r$ with $u_0 = w$, $u'_o = w'$, and, for $1 \leqslant i \leqslant k$, either

(i) $u_i = w_i$ and $u'_i = w'_i$ are elements of the first sort (worlds) that are linked to w and w', respectively, by unique paths of overlapping a-classes, or

(ii) $u_i = s_i$ and $u'_i = s'_i$ are of the second sort (information states) in some $\Sigma_a(w_i)$, respectively $\Sigma'_a(w'_i)$, for a pair of uniquely determined worlds w_i and w'_i as in (i).

We describe such positions after k rounds, for $k \leqslant r$, in terms of the associated $\mathbf{w} = (w_0, \ldots, w_k)$ and $\mathbf{w}' = (w'_0, \ldots, w'_k)$ and (for some i) $s_i \in \Sigma_{a_i}(w_i)$ and $s'_i \in \Sigma_{a'_i}(w'_i)$, which we further augment as follows.

With the tuple $\mathbf{w}$ of worlds in $\mathfrak{M}$ we associate a tree-like hypergraph structure whose hyperedges represent overlapping a-classes in the underlying basic modal $S5$-frames over the first sort, and which corresponds to a minimal spanning sub-tree containing the worlds in $\mathbf{w}$. We write tree($\mathbf{w}$) for this tree structure which has as its vertices for every w_i in $\mathbf{w}$ the unique sequence of worlds in which the a-classes that make up the shortest connecting path from the root w to w_i intersect, as well as its end points, $w = w_0$ and w_i; its hyperedges are the non-trivial subsets of vertices that fall within the same a-class, labelled by the corresponding a, i.e., essentially abstract representations of the overlapping

[16] It is important here that just certain finite levels of bisimilarity, not the full $\sim$-type, can be controlled.

a-classes along these shortest connecting paths to the root. We label each vertex u of this tree structure, which is a world in $\mathbb{M}$, by its $\sim^{\ell-d+1}$-type in $\mathbb{M}$ for $d = d(w,u)$ (distance from the root $w = w_0$, or depth in the tree structure), and label its hyperedges by the appropriate $a \in \mathcal{A}$. The tree $\text{tree}(\mathbf{w}')$ on the side of $\mathfrak{M}'$ is similarly defined. After k rounds, each of these trees can have at most $1 + k\ell$ vertices and each individual hyperedge can have at most $k+1$ vertices. Also, the a-labelling uniquely distinguishes all hyperedges that are incident with any individual world u that is a vertex in the tree. For worlds u of $\mathbb{M}$, agent $a \in \mathcal{A}$ and $n \in \mathbb{N}$, we now denote as $\mathfrak{M}_n \restriction [w]_a$ the locally full (i.e. full) relational encoding of the local a-structure $\mathbb{M}_n \restriction [w]_a$ at granularity n, which has colours for $\sim^n$-classes, i.e. $\sim^n$-types in $\mathbb{M}$ (cf. Definition 9.3); and analogously for $\mathbb{M}'$.

We argue that the second player can maintain the following conditions in terms of these tree structures through r rounds, thus forcing a win:

(1) the tree structures $\text{tree}(\mathbf{w})$ and $\text{tree}(\mathbf{w}')$ spanned by $\mathbf{w}$ and $\mathbf{w}'$ are isomorphic (as hyperedge- and vertex-labelled hypergraphs) via an isomorphism ζ that maps $\mathbf{w}$ to $\mathbf{w}'$:

$$\zeta \colon \text{tree}(\mathbf{w}), \mathbf{w} \simeq \text{tree}(\mathbf{w}'), \mathbf{w}'$$

(2) for each local a-structure on $[u]_a$ for any pair of vertices $u \in \text{tree}(\mathbf{w})$ and $u' \in \text{tree}(\mathbf{w}')$ at depth $d = d(w,u) = d(w',u')$ that are related by the isomorphism ζ from (1):

$$\mathfrak{M}_{\ell-d} \restriction [u]_a, \mathbf{s} \equiv_{r-z+1} \mathfrak{M}'_{\ell-d} \restriction [u']_a, \mathbf{s}'$$

where $\mathbf{s}$ and $\mathbf{s}'$ are tuples of size z that coherently list any singleton information states corresponding to the tree vertices incident with that a-edge and any information states $s_i \in \Sigma_a(u)$ and $s'_i \in \Sigma_a(u')$ that may have been chosen in the second sort during the first k rounds of the game.

Clearly these conditions are satisfied at the start of the game: (2) in this case does not add anything beyond (1), which in turn is a consequence of the assumption that $\mathbb{M}, w_0 \sim^{\ell+1} \mathbb{M}', w'_0$.

We show how to maintain conditions (1) and (2) through round k, in which the first player may either choose an information state (s_k or s'_k) or a world (w_k or w'_k). We refer to the position before this round as described by parameters $\mathbf{w}, \mathbf{w}', \text{tree}(\mathbf{w}), \ldots$ as above but at level $k-1$, and assume conditions (1) and (2) for those. The following shows how the second player can find responses so as to maintain conditions (1) and (2).

Case 1. Suppose the first player chooses a non-empty, non-singleton information state, say s_k (a choice on the side of $\mathfrak{M}'$ can be is treated symmetrically), so that $s_k \in \Sigma_a(w_k)$ for some uniquely determined w_k and a.

Case 1.1: if $w_k = u$ for some vertex u in $\text{tree}(\mathbf{w})$, u at depth d say, then we look at one round in the game for

$$\mathfrak{M}_{\ell-d} \restriction [u]_a, \mathbf{s} \ \equiv_{r-z+1} \mathfrak{M}'_{\ell-d} \restriction [u']_a, \mathbf{s}'$$

with a move by the first player on s_k, which extends $\mathbf{s}$ to $\mathbf{s}s_k$; this has an adequate response s'_k for the second player, which extends $\mathbf{s}'$ to $\mathbf{s}s'_k$ and guarantees $\equiv_{r-z}$ (this is the appropriate level since the tuples $\mathbf{s}$ and $\mathbf{s}'$ have been extended by one component).

Case 1.2: w_k is "new" and the appropriate w'_k that satisfies conditions (1) and (2) has to be located in a first step that simulates a move on w_k (treated as Case 2), after which we may proceed as in Case 1.1.

Case 2. Suppose the first player chooses a world in the k-th round, say w_k. The choice of an appropriate match w'_k is treated by induction on the distance that the newly chosen world w_k has from tree($\mathbf{w}$). In the base case, distance 0 from tree($\mathbf{w}$), w_k is a vertex of tree($\mathbf{w}$) and nothing needs to be updated: the response is dictated by the existing isomorphism ζ according to (1).

In all other cases, tree($\mathbf{w}$) and ζ need to be extended to encompass the new w_k. The new world w_k can be joined to tree($\mathbf{w}$) by a unique shortest path of overlapping a-classes of length greater than 0 that connects it to tree($\mathbf{w}$). The new branch in tree($\mathbf{w}$) will be joined to tree($\mathbf{w}$) either through a new a-hyperedge emanating from an existing vertex u of tree($\mathbf{w}$) (treated in Case 2.1) or through a new vertex u in a local a-structure corresponding to an existing hyperedge (treated in Case 2.2).

Case 2.1: it is instructive to look at the special case of distance 1 from tree($\mathbf{w}$) and then argue how to iterate for larger distance. So let w_k be at distance 1 from tree($\mathbf{w}$) in the sense that $w_k \in [u]_a$ for some u in tree($\mathbf{w}$) at depth d that is not incident with an a-hyperedge in tree($\mathbf{w}$). Since $\mathbb{M}, u \sim^{\ell-d+1} \mathbb{M}', u'$ for $u' = \zeta(u)$ and

$$\mathfrak{M}_{\ell-d} \restriction [u]_a, u \equiv_r \mathfrak{M}'_{\ell-d} \restriction [u']_a u',$$

a suitable response to the move that pebbles the world w_k (or the singleton information state $\{w_k\}$) in that game yields a world $w'_k \in [u']_a$ such that $\mathbb{M}, w_k \sim^{\ell-d} \mathbb{M}', w'_k$ and

$$\mathfrak{M}_{\ell-d} \restriction [u]_a, \{u\}, \{w_k\} \ \equiv_{r-1} \mathfrak{M}'_{\ell-d} \restriction [u']_a, \{u'\}, \{w'_k\}.$$

These levels of equivalence and granularity are appropriate since the depth of w_k and w'_k is $d+1$, and since one new vertex contributes to the new hyperedge. So we may extend the isomorphism ζ to map w_k to w'_k in keeping with conditions (1) and (2). If w_k is at greater distance from its nearest neighbour u in tree($\mathbf{w}$) we can iterate this process of introducing new hyperedges with one new element at a time, degrading the level of inquisitive bisimulation equivalence by 1 in every step that takes us one step further away from the root, but maintaining equivalences $\equiv_{r-1}$ between the newly added local a-structures.

Case 2.2: it remains to argue for the case of $w_k \in [u]_a$ for some u in tree($\mathbf{w}$) that is already incident with an a-edge of tree($\mathbf{w}$). By (2) we have

$$\mathfrak{M}_{\ell-d} \restriction [u]_a, \mathbf{s} \ \equiv_{r-z} \mathfrak{M}'_{\ell-d} \restriction [u']_a, \mathbf{s}',$$

where z is the size of the tuples $\mathbf{s}$ and $\mathbf{s}'$ already incident with these local a-structures and d is the depth of u and u'. So we can find a response to a move on $\{w_k\}$ in this game that yields $w'_k \in [u']_a$ such that $\mathbb{M}, w_k \sim^{\ell-d} \mathbb{M}', w'_k$ and

$$\mathfrak{M}_{\ell-d} \restriction [u]_a, \mathbf{s}\{w_k\} \ \equiv_{r-z-1} \mathfrak{M}'_{\ell-d} \restriction [u']_a, \mathbf{s}'\{w'_k\}.$$

The levels of bisimulation equivalence and granularity are appropriate as the depth of w_k and w'_k is one greater than that of u and u', and as the length of the tuples $\mathbf{s}$ and $\mathbf{s}'$ has been increased by 1. So we may extend ζ by matching w_k with w'_k and extending tree($\mathbf{w}$) and tree($\mathbf{w}$)$'$ by these new vertices and stay consistent with conditions (1) and (2). □

We are now in a position to prove Theorem 1.3, which, as stated in Section 9.2, is a consequence of Proposition 9.1.

Proof of Proposition 9.1. Given a formula $\varphi(x) \in \mathsf{FO}$ whose semantics is invariant under $\sim$ over one of the relevant classes $\mathcal{C}$ of world-pointed inquisitive epistemic relational models, we look at two world-pointed models $\mathfrak{M}, w \sim^n \mathfrak{M}', w'$ from $\mathcal{C}$ that are n-bisimulation equivalent for some $n \geqslant \ell + 1$, where ℓ is the locality parameter in a level (ℓ, r, m) of Gaifman equivalence that preserves the semantics of the first-order formula φ° over arbitrary relational structures, as given in Proposition 9.14. Here $\varphi^\circ(x)$ is the FO-formula that transcribes φ for the variant structures $\mathfrak{M}^\circ$ without explicit representation of the trivial information state $\emptyset$ in their second sort. By Remark 9.16 (as illustrated in Figure 6), it suffices to exhibit bisimilar companions $\hat{\mathfrak{M}}, \hat{w} \sim \mathfrak{M}, w$ and $\hat{\mathfrak{M}}', \hat{w}' \sim \mathfrak{M}', w'$ within the same class $\mathcal{C}$ such that

$$\hat{\mathfrak{M}}^\circ, \hat{w} \ \equiv_r^{(\ell)} \ \hat{\mathfrak{M}}'^\circ, \hat{w}'.$$

To this end we may pass from $\mathfrak{M}$ and $\mathfrak{M}'$ to the associated inquisitive epistemic models $\mathbb{M}, w := \mathfrak{M}^*, w$ and $\mathbb{M}', w' := \mathfrak{M}^*, w$, and replace those by (finite) bisimilar companions $\hat{\mathbb{M}}, \hat{w} \sim \mathbb{M}, w$ and $\hat{\mathbb{M}}', \hat{w}' \sim \mathbb{M}', w'$ that are N-acyclic for $N > 2\ell$, simple and K-rich for a sufficiently high value of K as in Lemma 9.18. We observe that the locally full relational encodings of $\hat{\mathfrak{M}} := \mathfrak{M}^{\mathsf{lf}}(\hat{\mathbb{M}})$ and $\hat{\mathfrak{M}}' := \mathfrak{M}^{\mathsf{lf}}(\hat{\mathbb{M}}')$ are again in the relevant class $\mathcal{C}$ and, by Lemma 9.18, satisfy $\hat{\mathfrak{M}}^\circ, \hat{w} \equiv_r^{(\ell)} \hat{\mathfrak{M}}'^\circ, \hat{w}'$, as required in Remark 9.16.

The variation of this upgrading argument for state-pointed $\mathfrak{M}, s \sim^n \mathfrak{M}', s'$ proceeds as follows. Remark 9.16 is available also in this case to reduce the claim to an upgrading that achieves a suitable target level $\equiv_r^{(\ell)}$ of local FO equivalence in the ℓ-neighbourhoods of the distinguished states. To this end we may apply structural transformations as above, individually to each one of the world-pointed models $(\mathfrak{M}, w)_{w \in s}$ and $(\mathfrak{M}', w')_{w' \in s'}$ to obtain corresponding families $(\hat{\mathfrak{M}}_w, \hat{w})_{w \in s}$ and $(\hat{\mathfrak{M}}'_{w'}, \hat{w}')_{w' \in s'}$ where always $\hat{\mathfrak{M}}_w, \hat{w} \sim \mathfrak{M}, w$ and $\hat{\mathfrak{M}}'_{w'}, \hat{w}' \sim \mathfrak{M}', w'$.

We now use 2^r copies of the members in each of these families to obtain models

$$\hat{\mathfrak{M}} := \bigoplus_{w \in s} \hat{\mathfrak{M}}_w \times [2^r]$$

$$\hat{\mathfrak{M}}' := \bigoplus_{w' \in s'} \hat{\mathfrak{M}}'_{w'} \times [2^r]$$

over which the associated inquisitive $S5$ models $\mathbb{M} := \mathfrak{M}^*$, $\mathbb{M}' := \mathfrak{M}'^*$, $\hat{\mathbb{M}} := \hat{\mathfrak{M}}^*$ and $\hat{\mathbb{M}}' := \hat{\mathfrak{M}}'^*$ with distinguished information states $\hat{s} := \{\hat{w} \colon w \in s\} \times [2^r]$ and $\hat{s}' := \{\hat{w}' \colon w' \in s'\} \times [2^r]$ satisfy

$$\mathbb{M}, s \sim \hat{\mathbb{M}}, \hat{s} \ \sim^n \ \hat{\mathbb{M}}', \hat{s}' \sim \mathbb{M}', s'.$$

Note that the new distinguished states $\hat{s}$ and $\hat{s}'$ are totally scattered as sets of worlds in the underlying Kripke structures. In order to make these distinguished states available in the relational encodings of these models, we can adjoin $\wp(\hat{s})$ and $\wp(\hat{s}')$ to the second sorts of $\hat{\mathfrak{M}}$ and $\hat{\mathfrak{M}}'$ to obtain relational models

$$\hat{\mathfrak{M}} + \wp(\hat{s}) \ \text{ and } \ \hat{\mathfrak{M}}' + \wp(\hat{s}'),$$

in which $\hat{s}$ and $\hat{s}'$ together with all their subsets are represented as elements in the second sort.

Here, e.g. the shorthand "$+\wp(\hat{s})$" denotes the effect of taking the union of the second sort of the relational model $\hat{\mathfrak{M}}$ with the power set of $\hat{s}$ and adding corresponding ε-links of these information states to their elements in the first sort. Crucially, $\wp(\hat{s})$ and the second sort of $\hat{\mathfrak{M}}$ itself share just the empty information state and the singleton information states for worlds in $\hat{s}$, and the associated Kripke structure of $\hat{\mathfrak{M}} + \wp(\hat{s})$ is the same as for $\hat{\mathfrak{M}}$.

We can now extend the game argument of Lemma 9.18 to these models to show the local equivalence

$$(*) \quad (\hat{\mathfrak{M}} + \wp(\hat{s}))^\circ, \hat{s} \ \equiv_r^{(\ell)} \ (\hat{\mathfrak{M}}' + \wp(\hat{s}'))^\circ, \hat{s}',$$

which by Remark 9.16 provides the desired upgrading to clinch the argument for Proposition 9.1 in the state-pointed case.

Due to the richness level of 2^r implemented in $\hat{s}$ and $\hat{s}'$ we can use a game argument based on Lemma 9.11 to reduce $(*)$ to the equivalences

$$\hat{\mathfrak{M}}^\circ_w, \hat{w} \ \equiv_r^{(\ell)} \ \hat{\mathfrak{M}}'^\circ_{w'}, \hat{w}'$$

for all the constituent pairs arising from matching $\mathfrak{M}, w \sim^n \mathfrak{M}', w'$ over s and s', as guaranteed by Lemma 9.18. □

10 Conclusion

In this paper we have seen the beginnings of a model theory for inquisitive modal logic. Our contribution started in Section 3, where we described the natural notion of bisimulation for inquisitive modal structures. From a game-theoretic perspective, bisimilarity and its approximations can be characterised in terms of a game which interleaves two kinds of moves: world-to-state moves (from w to some $s \in \Sigma(w)$) and state-to-world moves (from s to some $w \in s$).

In Section 4 we saw that bisimilarity relates to modal equivalence in the familiar way: two pointed models are distinguishable in the n-round bisimulation game if and only if they are distinguished by a formula of modal depth n.

In Section 5 we compared inquisitive modal logic to neighbourhood semantics for modal logic, showing that, although these two logics are interpreted over similar structures, they are very different in terms of their expressive power, and are invariant under different notions of bisimulation equivalence.

In Section 6 we discussed how inquisitive modal models can be encoded as two-sorted relational structures on which we can interpret first-order formulae of a suitable relational signature. This enabled us to define a standard translation from InqML to first-order logic, and to view InqML as a syntactic fragment of first-order logic. We then asked whether (and if so, over what classes of structures) this syntactic fragment coincides, up to logical equivalence, with the fragment determined by the semantic property of bisimulation invariance.

Using an inquisitive analogue of partial tree unfoldings in Section 8, we established a positive answer to this question for the class of all relational inquisitive models, both in the general case and in restriction to finite models.

Finally, in Section 9 we turned to the case of inquisitive epistemic models—the inquisitive version of multi-modal S5 models. Technically, this case is much more challenging, since the partial unfolding procedure used in the previous section is incompatible with the S5 frame conditions. Nevertheless, we saw that the characterisation result still holds in this setting—again, both in general and in restriction to finite models.

The results obtained in this paper provide us with a better understanding of inquisitive modal logic in at least two ways. From a more concrete perspective, we have given a characterisation of the expressive power of InqML which is very helpful in order to tell what properties of pointed models can and cannot be expressed in the language: for instance, it is easy to see that properties like $\mathbb{P}(w) :=$ "$W \in \Sigma(w)$" or $\mathbb{P}(w) :=$ "$\{w\} \in \Sigma(w)$" are not bisimulation invariant, and thus not expressible in InqML. From a more abstract perspective, we have looked at a natural notion of behavioural equivalence for inquisitive modal structures, whose main constituent is a map $\Sigma : W \to \wp\wp(W)$, rather than $\sigma : W \to \wp(W)$ as in Kripke structures. We saw that, in terms of expressive power, InqML is a natural choice for a language designed to talk about properties which are invariant under this notion of equivalence: among first-order properties (and over various natural classes of models) InqML expresses all and only those properties that are invariant in this sense.

References

[1] S. Abramsky and J. Väänänen. From IF to BI. *Synthese*, 167(2):207–230, 2009.

[2] I. Ciardelli. Modalities in the realm of questions: axiomatizing inquisitive epistemic logic. In R. Goré, B. Kooi, and A. Kurucz, editors, *Advances in Modal Logic (AIML)*, pages 94–113, London, 2014. College Publications.

[3] I. Ciardelli. *Questions in Logic*, volume DS-2016-01 of *ILLC Dissertation Series*. Institute for Logic, Language and Computation, Amsterdam, 2015.

[4] I. Ciardelli. Dependency as question entailment. In S. Abramsky, J. Kontinen, J. Väänänen, and H. Vollmer, editors, *Dependence Logic: theory and applications*, pages 129–181. Springer International Publishing Switzerland, 2016.

[5] I. Ciardelli. Questions as information types. *Synthese*, 195(1):321–365, 2018.

[6] I. Ciardelli, J. Groenendijk, and F. Roelofsen. Inquisitive semantics: A new notion of meaning. *Language and Linguistics Compass*, 7(9):459–476, 2013.

[7] I. Ciardelli, J. Groenendijk, and F. Roelofsen. On the semantics and logic of declaratives and interrogatives. *Synthese*, 192(6):1689–1728, 2015.

[8] I. Ciardelli and M. Otto. Bisimulation in inquisitive modal logic. In Jerôme Lang, editor, *Theoretical Aspects of Rationality and Knowledge (TARK) 16*. EPTCS, 2017.

[9] I. Ciardelli and F. Roelofsen. Inquisitive logic. *Journal of Philosophical Logic*, 40:55–94, 2011.

[10] I. Ciardelli and F. Roelofsen. Inquisitive dynamic epistemic logic. *Synthese*, 192(6):1643–1687, 2015.

[11] A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. *Annals of Pure and Applied Logic*, 161:1–42, 2009.

[12] H.-D. Ebbinhaus and J. Flum. *Finite Model Theory*. Perspectives in Mathematical Logic. Springer, 2 edition, 1999.

[13] H. Gaifman. On local and nonlocal properties. In J. Stern, editor, *Logic Colloquium 81*, pages 105–135. North Holland, 1982.

[14] P. Galliani. Inclusion and exclusion dependencies in team semantics – on some logics of imperfect information. *Annals of Pure and Applied Logic*, 163(1):68–84, 2012.

[15] V. Goranko and M. Otto. Model theory of modal logic. In P. Blackburn, J. van Benthem, and F. Wolter, editors, *Handbook of Modal Logic*, pages 249–329. Elsevier, 2007.

[16] H. H. Hansen. Monotone modal logics. MSc Thesis, University of Amsterdam, 2003.

[17] H. H. Hansen, C. Kupke, and E. Pacuit. Neighbourhood structures: Bisimilarity and basic model theory. *Logical Methods in Computer Science*, 2, 2009.

[18] J. Hintikka. *Knowledge and Belief: an Introduction to the Logic of the Two Notions*. Cornell University Press, 1962.

[19] W. Hodges. Compositional semantics for a language of imperfect information. *Logic Journal of IGPL*, 5(4):539–563, 1997.

[20] W. Hodges. Some strange quantifiers. In *Structures in logic and computer science. Lecture notes in computer science*, pages 51–65. Springer, New York, 1997.

[21] M. Otto. Elementary proof of the van Benthem–Rosen characterisation theorem. Technical Report 2342, Fachbereich Mathematik, Technische Universität Darmstadt, 2004.

[22] M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130:173–205, 2004.

[23] M. Otto. Highly acyclic groups, hypergraph covers and the guarded fragment. *Journal of the ACM*, 59 (1), 2012.

[24] E. Pacuit. *Neighborhood Semantics for Modal Logic*. Springer, 2017. forthcoming.

[25] F. Roelofsen. Algebraic foundations for the semantic treatment of inquisitive content. *Synthese*, 190(1):79–102, 2013.

[26] E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:427–439, 1997.

[27] J. Väänänen. *Dependence Logic: A New Approach to Independence Friendly Logic*. Cambridge University Press, 2007.

[28] J. Väänänen. Modal dependence logic. In K. R. Apt and R. van Rooij, editors, *New Perspectives on Games and Interaction*, pages 237–254. Amsterdam University Press, 2008.

[29] J. van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Napoli, 1983.

[30] F. Yang. *On extensions and variants of dependence logic: A study of intuitionistic connectives in the team semantics setting*. PhD thesis, University of Helsinki, 2014.

[31] F. Yang and J. Väänänen. Propositional logics of dependence. *Annals of Pure and Applied Logic*, 167:557–589, 2016.