

Bisimulation in Inquisitive Modal Logic*

Ivano Ciardelli

Institute for Logic, Language, and Computation
University of Amsterdam
i.a.ciardelli@uva.nl

Martin Otto

Department of Mathematics, Logic Group
Technische Universität Darmstadt
otto@mathematik.tu-darmstadt.de

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, and characterise inquisitive modal logic as the bisimulation invariant fragments of first-order logic over various classes of two-sorted relational structures. These results crucially require non-classical methods in studying bisimulations and first-order expressiveness over non-elementary classes.

1 Introduction

The recently developed framework of *inquisitive logic* [8, 2, 6, 4] 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* [20]. 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 [3]. This connection explains the deep similarities existing between systems of inquisitive logic and systems of dependence logic (see [22, 3, 2, 23]). 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 [9] in the setting of epistemic logic. The resulting *inquisitive epistemic logic* (IEL) 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 IEL can express not only that an agent knows that p ($\Box p$) but also that she knows *whether* p ($\Box?p$) or that she *wonders whether* p ($\boxplus?p$)—a statement that cannot be expressed without the use of embedded questions. As shown in [9], 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, IEL 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 and their raising in a communication scenario.

*Ivano Ciardelli’s research was financially supported by the European Research Council (ERC) 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*.

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., [2] suggests to interpret INQML as a logic of action. On this interpretation, a modal formula $\Box?p$ expresses that whether a certain fact p will come about is determined independently of the agent’s choices, while $\boxplus?p$ expresses that whether p will come about is fully determined by her 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 [11, 21, 19, 14]. 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 for INQML of the classical Ehrenfeucht–Fraïssé correspondence. This result is non-trivial in the inquisitive setting, because of some subtle issues stemming from the interleaving of first- and second-order features in inquisitive modal logic.

Theorem 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 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.

Theorem 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.*

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 ([11, 14, 10, 15], among others).

2 Inquisitive modal logic

In this section we provide an essential introduction to inquisitive modal logic, INQML [2]. For details and proofs, see §7 of [2].

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 (worlds), but relative to states of information, modelled extensionally as sets of worlds (viz., those worlds compatible with the given information).

Definition 3 (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 : 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 [5, 18, 9]).

Definition 4 (inquisitive states)

An inquisitive state over W is a non-empty set of information states $\Pi \subseteq \wp(W)$ satisfying

- $s \in \Pi$ and $t \subseteq s$ implies $t \in \Pi$ (*downward closure*).

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*, i.e., a map Σ that assigns to each world an inquisitive state. An inquisitive modal model is an inquisitive frame equipped with a valuation function.

Definition 5 (inquisitive modal models)

An inquisitive modal frame is a pair $\mathbb{F} = \langle W, \Sigma \rangle$, where $\Sigma: W \rightarrow \wp\wp(W)$ associates to each world $w \in W$

an inquisitive state $\Sigma(w)$. An inquisitive modal model is a pair $\mathbb{M} = \langle \mathbb{F}, V \rangle$ where \mathbb{F} is an inquisitive modal frame, and $V: \mathcal{P} \rightarrow \wp(W)$ is a propositional valuation function. 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 $\sigma: W \rightarrow \wp(W)$ induced by the inquisitive map Σ according to $\sigma(w) := \bigcup \Sigma(w)$.

Under an epistemic interpretation [9, 1], Σ is taken to describe not only an agent's *knowledge*, as in epistemic logic, but also her *issues*, i.e., the questions she is interested in. The agent's knowledge state at w , $\sigma(w) = \bigcup \Sigma(w)$, consists of those worlds that are compatible with what the agent knows. The agent's inquisitive state at w , $\Sigma(w)$, consists of those information states where her issues are settled.

2.3 Inquisitive Modal Logic

The syntax of inquisitive modal logic INQML is given by:

$$\varphi ::= p \mid \perp \mid (\varphi \wedge \psi) \mid (\varphi \rightarrow \psi) \mid (\varphi \vee \psi) \mid \Box \varphi \mid \boxplus \varphi$$

We treat negation and disjunction as defined connectives (syntactic shorthands) according to $\neg \varphi := \varphi \rightarrow \perp$, and $\varphi \vee \psi := \neg(\neg \varphi \wedge \neg \psi)$. In this sense, the above syntax includes standard propositional formulae in terms of atoms and connectives \wedge and \rightarrow 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, \vee , called *inquisitive disjunction*. We may read formulae built up by means of this connective as propositional questions. E.g., we read the formula $p \vee \neg p$ 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 when applied to statements, but crucially differ when applied to questions. Under an epistemic interpretation, $\Box ?p$ expresses the fact that the agent knows whether p , while $\boxplus ?p$ expresses (roughly) the fact that she wants to find out whether p .

While models for INQML are formally a class of neighbourhood models, the semantics of INQML is very different from neighbourhood semantics for modal logic.¹ As mentioned above, the semantics of INQML is given in terms of support relative to an information state, rather than truth at a possible world.

Definition 6 (semantics of INQML)

Let $\mathbb{M} = \langle W, \Sigma, V \rangle$ be an inquisitive modal model, $s \subseteq W$:

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

¹This different perspective on the models leads us to a notion of bisimulation which is different from the one that has been considered for neighbourhood models [12, 17, 13]. Due to space limitations we cannot discuss the difference in detail here.

As an illustration, consider the support conditions for the formula $?p := p \vee \neg p$: 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:

- 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 $\varphi \in \text{INQML}$.

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 state—vacuously supports everything. 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 obtained as a defined notion.

Definition 7 (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 7 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 8 (truth conditions for modal formulae)

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

Notice that truth in INQML cannot be given a direct recursive definition, as the truth conditions for modal formulae $\Box\varphi$ and $\Box\Box\varphi$ 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*.²

Definition 9 (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 [2], we view truth-conditional formulae as statements, and non-truth-conditional formulae as questions. The next proposition identifies a large class of truth-conditional formulae.

Proposition 10 Atomic formulae, \perp , and all formulae of the form $\Box\varphi$ and $\Box\Box\varphi$ are truth-conditional. The class of truth-conditional formulae is closed under all connectives except for \vee .

Using this fact, it is easy to see that all formulae of standard modal logic, i.e., formulae which do not contain \vee or $\Box\Box$, receive exactly the same truth conditions as in standard modal logic.

Proposition 11 If φ is a formula not containing \vee or $\Box\Box$, then $\mathbb{M}, w \models \varphi \iff \mathfrak{R}(\mathbb{M}), w \models \varphi$ in standard Kripke semantics.

²In dependence logic (e.g., [20, 23]) truth-conditional formulae are called *flat* formulae.

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, we have:

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

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 information 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.

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 chose the latter for its more immediate and intuitive appeal to the underlying dynamics of a ‘‘probing’’ of behavioural equivalence.

The 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 $\mathbb{M} = \langle W, \Sigma, V \rangle$ and $\mathbb{M}' = \langle W', \Sigma', V' \rangle$. We denote world-positions as $\langle w, w' \rangle$ and state-positions as $\langle s, s' \rangle$, 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 $\langle w, w' \rangle$, **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 $\langle s, s' \rangle$. Playing from a state-position $\langle s, s' \rangle$, **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 $\langle w, w' \rangle$. 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 $\langle w, w' \rangle$ 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, including infinite runs of the unbounded game, are won by **II**.

Definition 12 (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 $\langle w, w' \rangle$. \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 $\langle w, w' \rangle$.*

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.

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 INQML_n the class of INQML-formulae of depth up to n . It is easy to see that the semantics of any formula in 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 basic propositions, n -bisimilarity coincides with logical indistinguishability in INQML_n , which we denote as \equiv_{INQML}^n :

$$\mathbb{M}, s \equiv_{\text{INQML}}^n \mathbb{M}', s' \stackrel{\text{def}}{\iff} \begin{cases} \mathbb{M}, s \models \varphi \iff \mathbb{M}', s' \models \varphi \\ \text{for all } \varphi \in \text{INQML}_n. \end{cases}$$

Theorem 1. *Given a 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' :*

$$\mathbb{M}, s \sim^n \mathbb{M}', s' \iff \mathbb{M}, s \equiv_{\text{INQML}}^n \mathbb{M}', s'$$

Notice that, by taking s and s' to be singleton states, we obtain the corresponding connection for world-pointed models as a special case: $\mathbb{M}, w \sim^n \mathbb{M}', w' \iff \mathbb{M}, w \equiv_{\text{INQML}}^n \mathbb{M}', w'$. As customary, the crucial implication of the theorem, from right to left, follows from the existence of *characteristic formulae* that define \sim^n -classes of worlds, information states and inquisitive states over models—and it is here that the finiteness of \mathcal{P} is crucial.

Proposition 13 (characteristic formulae for \sim^n -classes)

For any world-pointed model \mathbb{M}, w over a finite set of atomic propositions \mathcal{P} , and for any $n \in \mathbb{N}$ there is a formula $\chi_{\mathbb{M}, w}^n \in \text{INQML}$ of modal depth n s.t. $\mathbb{M}', w' \models \chi_{\mathbb{M}, w}^n \iff \mathbb{M}', w' \sim^n \mathbb{M}, w$.

Proof. By simultaneous induction on n , we define formulae $\chi_{\mathbb{M}, w}^n$ together with auxiliary formulae $\chi_{\mathbb{M}, s}^n$ and $\chi_{\mathbb{M}, \Pi}^n$ 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:

$$\begin{aligned} \chi_w^0 &= \bigwedge \{p : w \in V(p)\} \wedge \bigwedge \{\neg p : w \notin V(p)\} \\ \chi_s^n &= \bigvee \{\chi_w^n : w \in s\} \\ \chi_\Pi^n &= \bigvee \{\chi_s^n : s \in \Pi\} \\ \chi_w^{n+1} &= \chi_w^n \wedge \boxplus \chi_{\Sigma(w)}^n \wedge \bigwedge \{\neg \boxplus \chi_\Pi^n : \Pi \subseteq \Sigma(w), \Pi \not\sim^n \Sigma(w)\} \end{aligned}$$

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_w^n , and analogously for χ_s^n or χ_Π^n . We can then prove by simultaneous induction on n that these formulae satisfy the following properties:

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

The details of the inductive proof are given in appendix A.1. □

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 (in which φ is supported).

Corollary 14 *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 $n \in \mathbb{N}$.*

5 Relational inquisitive models

In this paper, we want to compare the expressive power of inquisitive modal logic with that of first-order logic. However, this is not straightforward. 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 $\Sigma : W \rightarrow \wp \wp(W)$. 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.

5.1 Relational Inquisitive Models

Definition 15 (relational models) *A relational inquisitive modal model is a relational structure*

$$\mathfrak{M} = \langle W, S, E, \varepsilon, (P_i)_{p_i \in \mathcal{P}} \rangle$$

where $W, S \neq \emptyset$ are sets, $E, \varepsilon \subseteq W \times S$, and $P_i \subseteq W$. With $s \in S$ we associate the set $\underline{s} := \{w \in W \mid w \varepsilon s\} \subseteq W$ and require the following conditions, which enforce resemblance with inquisitive modal models:

- *Extensionality:* if $\underline{s} = \underline{s}'$, then $s = s'$.
- *Non-emptiness:* for every w , $E[w] \neq \emptyset$.
- *Downward closure:* if $s \in E[w]$ and $t \subseteq \underline{s}$, there is an $s' \in S$ such that $\underline{s}' = t$ and $s' \in E[w]$.

By extensionality, the second sort S can be identified with a domain $\{\underline{s} \mid s \in S\} \subseteq \wp(W)$ of sets over the first sort. We will always make this identification and view a relational model as a structure $\mathfrak{M} = \langle W, S, E, \varepsilon, (P_i) \rangle$ where $S \subseteq \wp(W)$ and ε is the actual membership relation. In the following we shall therefore also specify relational models by just $\mathfrak{M} = \langle W, S, E, (P_i) \rangle$, when the fact that $S \subseteq \wp(W)$ and the natural interpretation of ε are understood.

Notice that a relational model \mathfrak{M} induces a corresponding Kripke model $\mathfrak{K}(\mathfrak{M})$ on W . We simply let wRw' if for some $s \in S$ we have $w \varepsilon s$ and $w' \varepsilon s$, and we let $R[w] := \{w' \mid wRw'\}$.

5.2 Natural Classes of Relational Models

In addition to extensionality and downward closure, 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 perspective of each world w .

Definition 16 (classes of relational models)

A relational model $\mathfrak{M} = \langle W, S, E, (P_i) \rangle$ is called:

- full if $S = \wp(W)$;
- locally full if $S \supseteq \wp(R[w])$ for all $w \in W$.

These conditions suggest different ways of encoding a concrete inquisitive modal model $M = \langle W, \Sigma, V \rangle$ as a relational model.

Definition 17 (relational encodings) Let $\mathbb{M} = \langle W, \Sigma, V \rangle$ be an inquisitive modal model. We define three relational encodings $\mathfrak{M}^{\text{rel}}, \mathfrak{M}^{\text{lf}}, \mathfrak{M}^{\text{full}}$ 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 $\mathfrak{M}^{\text{rel}}(\mathbb{M})$: $S = \text{image}(\Sigma)$;
- for $\mathfrak{M}^{\text{lf}}(\mathbb{M})$: $S = \{s \subseteq \sigma(w) : w \in W\}$;
- for $\mathfrak{M}^{\text{full}}(\mathbb{M})$: $S = \wp(W)$.

Clearly, $\mathfrak{M}^{\text{rel}}(\mathbb{M})$ is the minimal relational counterpart of \mathbb{M} , $\mathfrak{M}^{\text{lf}}(\mathbb{M})$ its minimal counterpart that is locally full, and $\mathfrak{M}^{\text{full}}(\mathbb{M})$ its unique counterpart that is full.

5.3 Relational Models and First-Order Logic

A relational inquisitive model supports a two-sorted first-order language having two relation symbols E and ε , and a number of predicate symbols P_i for $i \in I$. It is easy to translate formulae $\varphi \in \text{INQML}$ to FO-formulae $\varphi^*(x)$ in a single free variable x of the second sort in such a way that, if \mathbb{M} is an inquisitive modal model and $\mathfrak{M}(\mathbb{M})$ is any of the above encodings, we have:

$$\mathbb{M}, s \models \varphi \iff \mathfrak{M}(\mathbb{M}) \models \varphi^*(s)$$

This translation can be seen as an analogue of the standard translation of modal logic to first-order logic. The framework of relational inquisitive models thus allows us to view INQML as a syntactic fragment of FO, $\text{INQML} \subseteq \text{FO}$, just as standard modal logic ML over Kripke structures may be regarded as a fragment $\text{ML} \subseteq \text{FO}$.

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

6 The \sim -invariant fragment of FO

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 for the rest of this paper is to which extent INQML may express all

properties of worlds that are FO-expressible.³ 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

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

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

6.1 Bisimulation Invariance and Compactness

The inquisitive Ehrenfeucht–Fraïssé theorem, Theorem 1, implies \sim -invariance for all of INQML. By Corollary 14 it further implies *expressive completeness* of INQML_n for any \sim^n -invariant property of world-pointed models. In order to prove (\dagger) in restriction to some particular class \mathcal{C} of relational inquisitive models, it is thus necessary and sufficient to show that, for any $\varphi(x) \in \text{FO}$, \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 18 *For any class \mathcal{C} of relational inquisitive models, the following are equivalent:*

- (i) $\text{INQML} \equiv \text{FO}/\sim$ 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 .

Interestingly, first-order logic does not satisfy compactness in restriction to the (non-elementary) class of relational inquisitive models (see Example 22 in the appendix). More importantly, over the class of *full* relational models, violations of compactness can even be exhibited for \sim -invariant formulae.

Observation 19 *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, first-order logic violates compactness over full relational models.*

6.2 The Characterisation Theorem

In light of Observation 18, Observation 19 means that (\dagger) fails over the class of full relational models. This is not too surprising: on full relational models, FO has access to full-fledged second-order quantification, while INQML can only quantify over subsets within the range of Σ . This is in sharp contrast with our main theorem:

THEOREM 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, $\text{INQML} \equiv \text{FO}/\sim$, 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} .*

Without recourse to compactness, the most useful tool from first-order model theory for our purposes is the *local nature* of first-order logic over relational structures, in terms of Gaifman distance. In the setting of a relational model, 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

³Our results can be extended easily to address the corresponding question for properties of information states, but due to space limitations we will not make the corresponding results explicit here.




Figure 1: Generic upgrading pattern.

$N^\ell(w)$ of a world w consists of all worlds or states at distance up to ℓ from w in this sense. 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 $\mathfrak{M} \upharpoonright N^\ell(w), w$, is also a world-pointed relational model.

In light of Observation 18, to show that (\dagger) holds over a class \mathcal{C} we need to show that a first-order formula $\varphi(x)$ whose semantics is invariant under \sim over the class \mathcal{C} , is in fact invariant under one of the much coarser finite approximations \sim^n over \mathcal{C} , for some value n 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 [14] of van Benthem’s classical characterisation of basic modal logic [21] and its finite model theory version due to Rosen [19] (for ramifications of this method, see also [15, 10] and [16]). This approach involves an *upgrading* of a sufficiently high finite level \sim^n of bisimulation equivalence (or \equiv_{INQML}^n) to a finite target level \equiv_q of elementary equivalence, where q is the quantifier rank of φ . Concretely, this amounts to finding, for any world-pointed relational model \mathfrak{M}, w , a fully bisimilar pointed model \mathfrak{M}^*, w^* with the property that, if $\mathfrak{M}, w \sim^n \mathfrak{M}', w'$, then $\mathfrak{M}^*, w^* \equiv_q \mathfrak{M}'^*, w'^*$. The diagram in Figure 1 shows how \sim -invariance of φ , together with its nature as a first-order formula of quantifier rank q , entails its \sim^n -invariance —simply by taking the detour via the lower rung. In the following section, we show how to achieve the required upgradings for various classes \mathcal{C} of relational models; we use a variation on an upgrading technique from [14], based on an inquisitive analogue of partial tree unfoldings.

6.3 Partial Unfolding and Stratification

Theorem 2 boils down to the compactness property expressed in Observation 18 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.

Definition 20 *We say that a relational inquisitive model \mathfrak{M} is stratified if its two domains W and S consist of essentially disjoint⁴ strata $(W_i)_{i \in \mathbb{N}}$ and $(S_i)_{i \in \mathbb{N}}$ s.t.*

- (i) $W = \dot{\bigcup} W_i$ and $S \setminus \{\emptyset\} = \dot{\bigcup} (S_i \setminus \{\emptyset\})$;⁴
- (ii) $S_i \subseteq \wp(W_{i+1})$, and $E[w] \subseteq S_i$ for all $w \in W_i$.

For an even number $\ell \neq 0$ and a world w , we say that \mathfrak{M} is stratified to depth ℓ from w if $\mathfrak{M} \upharpoonright N^\ell(w)$ is stratified.

It is not hard to see that any world-pointed relational inquisitive model is bisimilar to one that is stratified. Moreover, for any even number $\ell \neq 0$, a finite world-pointed relational model is bisimilar to one that is finite and stratified to depth ℓ from its distinguished world. If the original model is locally full, the process of partial unfolding leading to an (ℓ) -stratified model preserves local fullness.

⁴The S_i will share the trivial information state \emptyset , by extensionality and the downward closure requirement. This is unproblematic for our purposes.




Figure 2: Upgrading pattern for Theorem 2.

Observation 21 For relational models \mathfrak{M} and \mathfrak{M}' that are stratified to depth ℓ for some even $\ell \neq 0$, and for $n \geq \ell/2$:

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

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

Proof of Theorem 2. Let \mathcal{C} be any one of the classes in the theorem and let $\varphi(x) \in \text{FO}_q$ be \sim -invariant 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 2. Towards its ingredients, consider a world-pointed relational model \mathfrak{M}, w in \mathcal{C} . Since φ is \sim -invariant, we can 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. Both models contain q distinct isomorphic copies of \mathfrak{M} as well as of $\mathfrak{M} \upharpoonright N^\ell(w)$. In addition, \mathfrak{M}_0 contains a copy of $\mathfrak{M} \upharpoonright N^\ell(w)$ with the distinguished world w , while \mathfrak{M}_1 contains a copy of \mathfrak{M} with the distinguished world w :⁵

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

Using an Ehrenfeucht-Fraïssé game argument for FO it is possible to show (cf. Appendix A.3) that

$$\mathfrak{M}_0, w \equiv_q \mathfrak{M}_1, w.$$

Given any two pointed models $\mathfrak{M}, w \sim^n \mathfrak{M}', w'$ in \mathcal{C} , we can see that φ is preserved between them by chasing the diagram in Figure 2 along the path through the auxiliary models, which are all in \mathcal{C} . \square

7 Conclusion

We have seen the foundations of a model theory for inquisitive modal logic in two main aspects. Firstly, the notion of *inquisitive bisimulation equivalence* has been established as the appropriate notion of semantic invariance by an Ehrenfeucht-Fraïssé correspondence, which provides a precious tool for studying

⁵To obey extensionality, when taking disjoint unions we identify the various copies of the empty information state coming from the different components of the model.

the expressive power of inquisitive modal logic. Secondly, we have seen that INQML admits *model-theoretic characterisations* as the bisimulation-invariant fragment of classical first-order logic over certain classes of relational structures with two sorts for worlds and information states. Our result holds both in the general setting, and in restriction to finite models. The model-theoretic challenges arise in dealing with non-elementary classes of models, whose essentially two-sorted nature extends first-order expressiveness in the direction of monadic second-order logic. Unpublished work [7] indicates that this approach can be taken considerably further: characterisations analogous to those presented here for basic INQML can be obtained for *inquisitive epistemic logic*—the multi-agent, S5-like variant of INQML. In that setting, the model unfolding procedure that we used here to establish our Theorem 2 can no longer be used, because the resulting structures would no longer satisfy the inquisitive S5 constraints. Instead, new and more complex techniques are needed.

References

- [1] I. Ciardelli (2014): *Modalities in the realm of questions: axiomatizing inquisitive epistemic logic*. In R. Goré, B. Kooi & A. Kurucz, editors: *Advances in Modal Logic (AIML)*, College Publications, London, pp. 94–113.
- [2] I. Ciardelli (2015): *Questions in Logic*. *ILLC Dissertation Series DS-2016-01*, Institute for Logic, Language and Computation, Amsterdam.
- [3] I. Ciardelli (2016): *Dependency as question entailment*. In S. Abramsky, J. Kontinen, J. Väänänen & H. Vollmer, editors: *Dependence Logic: theory and applications*, Springer International Publishing Switzerland, pp. 129–181, doi:10.1007/978-3-319-31803-5_8.
- [4] I. Ciardelli (2016): *Questions as information types*. *Synthese*, doi:10.1007/s11229-016-1221-y.
- [5] I. Ciardelli, J. Groenendijk & F. Roelofsen (2013): *Inquisitive semantics: A new notion of meaning*. *Language and Linguistics Compass* 7(9), pp. 459–476, doi:10.1111/lnc3.12037.
- [6] I. Ciardelli, J. Groenendijk & F. Roelofsen (2015): *On the semantics and logic of declaratives and interrogatives*. *Synthese* 192(6), pp. 1689–1728, doi:10.1007/s11229-013-0352-7.
- [7] I. Ciardelli & M. Otto (2017): *Bisimulation in Inquisitive Modal Logic*. Unpublished manuscript.
- [8] I. Ciardelli & F. Roelofsen (2011): *Inquisitive logic*. *Journal of Philosophical Logic* 40, pp. 55–94, doi:10.1007/s10992-010-9142-6.
- [9] I. Ciardelli & F. Roelofsen (2015): *Inquisitive dynamic epistemic logic*. *Synthese* 192(6), pp. 1643–1687, doi:10.1007/s11229-014-0404-7.
- [10] A. Dawar & M. Otto (2009): *Modal characterisation theorems over special classes of frames*. *Annals of Pure and Applied Logic* 161, pp. 1–42, doi:10.1016/j.apal.2009.04.002.
- [11] V. Goranko & M. Otto (2007): *Model Theory of Modal Logic*. In P. Blackburn, J. van Benthem & F. Wolter, editors: *Handbook of Modal Logic*, Elsevier, pp. 249–329, doi:10.1016/S1570-2464(07)80008-5.
- [12] H. Hansen (2003): *Monotone Modal Logics*. MSc Thesis, University of Amsterdam.
- [13] H. Hansen, C. Kupke & E. Pacuit (2009): *Neighbourhood structures: Bisimilarity and basic model theory*. *Logical Methods in Computer Science* 2.
- [14] M. Otto (2004): *Elementary Proof of the van Benthem–Rosen Characterisation Theorem*. Technical Report 2342, Fachbereich Mathematik, Technische Universität Darmstadt.
- [15] M. Otto (2004): *Modal and guarded characterisation theorems over finite transition systems*. *Annals of Pure and Applied Logic* 130, pp. 173–205, doi:10.1016/j.apal.2004.04.003.
- [16] M. Otto (2012): *Highly acyclic groups, hypergraph covers and the guarded fragment*. *Journal of the ACM* 59 (1), doi:10.1145/2108242.2108247.

- [17] E. Pacuit (2007): *Neighborhood semantics for modal logic: An introduction*. Lecture notes for a course at ESSLLI.
- [18] F. Roelofsen (2013): *Algebraic foundations for the semantic treatment of inquisitive content*. *Synthese* 190(1), pp. 79–102, doi:10.1007/s11229-013-0282-4.
- [19] E. Rosen (1997): *Modal logic over finite structures*. *Journal of Logic, Language and Information* 6, pp. 427–439, doi:10.1023/A:1008275906015.
- [20] J. Väänänen (2007): *Dependence Logic: A New Approach to Independence Friendly Logic*. Cambridge University Press, doi:10.1017/CBO9780511611193.
- [21] J. van Benthem (1983): *Modal Logic and Classical Logic*. Bibliopolis, Napoli, doi:10.2307/2274406.
- [22] F. Yang (2014): *On extensions and variants of dependence logic: A study of intuitionistic connectives in the team semantics setting*. Ph.D. thesis, University of Helsinki.
- [23] F. Yang & J. Väänänen (2016): *Propositional logics of dependence*. *Annals of Pure and Applied Logic* 167, pp. 557–589, doi:10.1016/j.apal.2016.03.003.

A Appendix

A.1 Proof of Theorem 1

Proof of Theorem 1/Proposition 13. We fill in the missing part in the proof of Theorem 1, showing that the characteristic formulae, as defined in connection with Proposition 13, have the following properties:

1. $\mathbb{M}', w' \models \chi_{\mathbb{M}, w}^n \iff \mathbb{M}', w' \sim^n \mathbb{M}, w$
2. $\mathbb{M}', s' \models \chi_{\mathbb{M}, s}^n \iff \mathbb{M}', s' \sim^n \mathbb{M}, t$ for some $t \subseteq s$
3. $\mathbb{M}', s' \models \chi_{\mathbb{M}, \Pi}^n \iff \mathbb{M}', s' \sim^n \mathbb{M}, s$ 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 $\mathbb{M}', s' \models \chi_{\mathbb{M}, s}^n$, that is, suppose $\mathbb{M}', s' \models \bigvee \{\chi_{\mathbb{M}, w}^n \mid w \in s\}$. This requires that for any $w' \in s'$ we have $\mathbb{M}', w' \models \chi_{\mathbb{M}, w}^n$ 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 $\mathbb{M}', w' \models \chi_{\mathbb{M}, w}^n$, which implies $\mathbb{M}', w' \models \chi_{\mathbb{M}, s}^n$. Since this holds for any $w' \in s'$, and since $\chi_{\mathbb{M}, s}^n$ is a truth-conditional formula (by Proposition 10), it follows that $\mathbb{M}', s' \models \chi_{\mathbb{M}, s}^n$.

For claim (3), suppose $\mathbb{M}', s' \models \chi_{\mathbb{M}, \Pi}^n$. This implies $\mathbb{M}', s' \models \chi_{\mathbb{M}, s}^n$ 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), $\mathbb{M}', s' \models \chi_{\mathbb{M}, t}^n$, and since $t \in \Pi$, also $\mathbb{M}', s' \models \chi_{\mathbb{M}, \Pi}^n$.

Next, we use these facts to show that claim (1) holds for all $n \in \mathbb{N}$, by induction on n . The claim $\mathbb{M}', w' \models \chi_{\mathbb{M}, w}^0 \iff \mathbb{M}', w' \sim_0 \mathbb{M}, w$ follows immediately from the definition of $\chi_{\mathbb{M}, w}^0$. 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 $\mathbb{M}', w' \models \chi_{\mathbb{M}, w}^{n+1}$. This amounts to showing that: (i) $\mathbb{M}', w' \models \chi_{\mathbb{M}, w}^n$; (ii) $\mathbb{M}', w' \models \boxplus \chi_{\mathbb{M}, \Sigma(w)}^n$; (iii) $\mathbb{M}', w' \models \neg \boxplus \chi_{\mathbb{M}, \Pi}^n$ when $\Pi \subseteq \Sigma(w)$ and $\Pi \not\sim_n \Sigma(w)$. 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 $\mathbb{M}', w' \models \chi_{\mathbb{M}, w}^n$.
- (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, $\mathbb{M}', s' \models \chi_{\mathbb{M}, \Sigma(w)}^n$. This holds for all $s' \in \Sigma'(w')$, and so $\mathbb{M}', w' \models \boxplus \chi_{\mathbb{M}, \Sigma(w)}^n$.

(iii) Suppose for a contradiction that for some $\Pi \subseteq \Sigma(w)$, $\Pi \not\sim_n \Sigma(w)$ and $\mathbb{M}', w' \models \boxplus \chi_{\mathbb{M}, \Pi}^n$. This means that every $s' \in \Sigma'(w')$ supports $\chi_{\mathbb{M}, \Pi}^n$ and thus, by our induction hypothesis, is n -bisimilar to some $s \in \Pi$. Since $\Pi \subseteq \Sigma(w)$ and $\Pi \not\sim_n \Sigma(w)$, 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 $\mathbb{M}', w' \models \chi_{\mathbb{M}, w}^{n+1}$. 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 $\chi_{\mathbb{M}, w}^n$ is a conjunct of $\chi_{\mathbb{M}, w}^{n+1}$, 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 $\boxplus \chi_{\mathbb{M}, \Sigma(w)}^n$ is a conjunct of $\chi_{\mathbb{M}, w}^{n+1}$, $\mathbb{M}', w' \models \boxplus \chi_{\mathbb{M}, \Sigma(w)}^n$. This implies that any $s' \in \Sigma'(w')$ supports $\chi_{\mathbb{M}, \Sigma(w)}^n$. 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 $\mathbb{M}', s' \models \chi_{\mathbb{M}, \Pi}^n$. And since this is true for each $s' \in \Sigma'(w')$, we have $\mathbb{M}', w' \models \boxplus \chi_{\mathbb{M}, \Pi}^n$. 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 $\Pi \not\sim_n \Sigma(w)$, which means that $\neg \boxplus \chi_{\mathbb{M}, \Pi}^n$ is a conjunct of $\chi_{\mathbb{M}, w}^{n+1}$. But then, since $\mathbb{M}', w' \models \chi_{\mathbb{M}, w}^{n+1}$, we should have $\mathbb{M}', w' \models \neg \boxplus \chi_{\mathbb{M}, \Pi}^n$ contrary to what we found above.

This completes the proof of Proposition 13. We can then use the properties of our characteristic formulae to prove the non-trivial direction of Theorem 1. For suppose $\mathbb{M}, s \not\sim_n \mathbb{M}', s'$: then either of the states s and s' is not n -bisimilar to any subset of the other. Without loss of generality, say it is s' . By the property of the formula $\chi_{\mathbb{M}, s}^n$ we have $\mathbb{M}, s \models \chi_{\mathbb{M}, s}^n$ but $\mathbb{M}', s' \not\models \chi_{\mathbb{M}, s}^n$. Since the modal depth of $\chi_{\mathbb{M}, s}^n$ is n , this shows that $\mathbb{M}, s \not\equiv_{\text{INQML}}^n \mathbb{M}', s'$. \square

A.2 Failures of Compactness

Observation 19 refers to the following example of a first-order property of worlds w in *full* relational models that is \sim -invariant and (as a well-foundedness assertion) obviously incompatible with compactness. In terms of the accessibility relation $R = \{(u, v) : v \in s \text{ for some } s \text{ with } uEs\}$:

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

This property is not expressible in INQML: although it is \sim -invariant, it is clearly not invariant under any finite level of bisimulation equivalence. 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} over the underlying Kripke frame, can therefore be expressed in two-sorted first-order logic over full relational models:

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

This shows that the analogue of our Theorem 2 fails for the class of full relational models: over this class, there are properties that are FO-definable and \sim -invariant, but not definable in INQML.

It is also possible to show that compactness fails over the (non-elementary) class of all relational models. However, in this case, Theorem 2 together with the compactness of INQML implies that this cannot happen for FO-formulae that are \sim -invariant.

The counterexample to compactness is obtained by relativising the property \mathbb{P} given above to a non-empty information state s in the image of Σ : in restriction to s , any relational encoding of $\langle W, \Sigma \rangle$ provides the full power set $\wp(s)$ in the second sort.

Example 22 There is a first-order formula in a single free variable the second sort (information state) which over any relational inquisitive model says of an element s that (i) $s \in E[W]$ and (ii) there are no infinite R -paths included in s . Condition (i) is expressed simply by $\exists w.Ews$; for (ii), we relativise to s the above formula that defines property \mathbb{P} , and then universally quantify over $w \in s$. \square

A.3 Locality in Stratified Models

We present the game argument at the heart of the proof of Theorem 2 in Section 6, which needs to establish

$$(**) \quad \mathfrak{M}_{0,w} \equiv_q \mathfrak{M}_{1,w}$$

for structures

$$\begin{aligned} \mathfrak{M}_{0,w} &= q \otimes \mathfrak{M} \oplus \mathfrak{M} \upharpoonright N^\ell(w), w \oplus q \otimes \mathfrak{M} \upharpoonright N^\ell(w) \\ \mathfrak{M}_{1,w} &= q \otimes \mathfrak{M} \oplus \mathfrak{M}, w \oplus q \otimes \mathfrak{M} \upharpoonright N^\ell(w) \end{aligned}$$

obtained as essentially disjoint sums from a world-pointed relational inquisitive model \mathfrak{M}, w that is stratified to depth $\ell = 2^q$, so that its truncation $\mathfrak{M} \upharpoonright N^\ell(w)$ is stratified. ‘Essentially disjoint sums’ here refers to disjoint sums with identifications of the empty information states \emptyset across the disjoint parts. It is not hard to see that the q -equivalence claim in (**), however, is insensitive to whether the empty information states, which are uniformly present in the second sort of each component, are identified or not. So we may as well work with proper disjoint unions in the following proof.

Proof of ()** We argue that the second player has a winning strategy in the classical q -round Ehrenfeucht–Fraïssé game over these two structures 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 and $\mathbf{u}' = (u'_0, u'_1, \dots, u'_m)$ with $u'_0 = w$ in \mathfrak{M}_1 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 corresponding clusters are in isomorphic configurations of isomorphic component structures of \mathfrak{M}_0 and \mathfrak{M}_1 or in isomorphic configurations in $\mathfrak{M}_0 \upharpoonright N^\ell(w)$ and $\mathfrak{M}_1 \upharpoonright N^\ell(w)$.

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 or $u' = u'_{m+1}$ in \mathfrak{M}_1 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 or $u' = u'_{m+1}$ in \mathfrak{M}_1 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.

All steps of this proof restrict naturally to the scenarios of (finite or general) locally full relational inquisitive structures.