

INQUISITIVE BISIMULATION

IVANO CIARDELLI AND MARTIN OTTO

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, and characterise inquisitive modal logic as the bisimulation invariant fragment of first-order logic over various natural classes of two-sorted structures.

§1. Introduction. The recently developed framework of *inquisitive logic* [10, 7, 3, 5] 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* [32]. 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 [35, 4, 3, 36]). 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 [11] 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 present paper extends and subsumes the proceedings paper [8]. It is a revised version of the first part of a larger manuscript ([9]), accessible via arXiv, the second part of which we intend to develop into a separate paper. We are grateful to the reviewers for their precious comments and suggestions, which pushed us to clarify some important points. 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.

the formula $\Box p$) but also that she knows *whether* p ($\Box?p$) or that she *wonders whether* p (by the formula $\Box?p$)—a statement that cannot be expressed without the use of embedded questions. As shown in [11], 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. For instance, 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 $\Box?p$ 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. In standard modal logic, 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 similar: we still have a set $\Sigma(w)$ associated with each possible world w , and a set $[\varphi]_{\mathbb{M}}$ associated with a formula φ . Now, however, both $\Sigma(w)$ and $[\varphi]_{\mathbb{M}}$ are no longer sets of worlds, but sets of 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 $\Box\varphi$ 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. 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 [17, 34, 31, 26]. 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 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.*

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 generalise a number of notions and techniques developed in the model-theoretic analysis of modal logics over non-elementary classes of frames (cf. [17, 26, 13, 27], among others). In the present paper we technically focus on the general case of inquisitive modal models. This also sets the stage for the model-theoretic treatment of inquisitive epistemic models. That case, which is of particular interest from the point of view of logical modelling, also requires some further extensions of the technical apparatus presented here. We aim to present corresponding results from [9] in a sequel to the present paper.

§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 [20], information states are modelled extensionally as sets of worlds, namely, the set of those worlds which are compatible with the given information.¹

DEFINITION 2.1. [information states]

An *information state* over a set of worlds W is a subset $s \subseteq W$.

The empty set represents a state of inconsistent information, which is not compatible with any world. We refer to it as the *inconsistent state*.

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*.

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. As described in the next section, an inquisitive state can be seen as a combined representation of information and issues. For more discussion on the significance of this structure, see [6, 11, 12].

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 $\sigma(w)$, i.e., an information state: the set of worlds that are *accessible* from w .

Similarly, an inquisitive modal frame consists of a set W of worlds together with an *inquisitive assignment*, a map $\Sigma : W \rightarrow \wp(W)$ that assigns to each world a corresponding inquisitive state $\Sigma(w)$, i.e. a downward closed set of information states. A *model* is a frame enriched by a propositional assignment.

DEFINITION 2.3. [inquisitive modal models]

An inquisitive modal frame is a pair $\mathbb{F} = (W, \Sigma)$, where W is a set, whose

¹An 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 [21, 22] and dependence logic [32, 33, 1, 16, 35, 36], where sets of worlds are referred to as *teams*. Although they originated independently and for different purposes, inquisitive logic and dependence logic are tightly related. For detailed discussion of this connection, see [3, 4].

elements are referred to as *worlds*, and $\Sigma: W \rightarrow \wp(W)$ assigns an inquisitive state $\Sigma(w)$ to each world $w \in W$.

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 $V: \mathcal{P} \rightarrow \wp(W)$ 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 $\sigma: W \rightarrow \wp(W)$ induced by the inquisitive map Σ according to

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

A natural interpretation for inquisitive modal models is the epistemic one, developed in [11, 2]. In that 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*.² 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 $\Sigma_a: W \rightarrow \wp(W)$ 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).

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, $w_{pq}, w_{p\bar{q}}, w_{\bar{p}q}, w_{\bar{p}\bar{q}}$, 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.

$$\begin{aligned} \Sigma(w_{pq}) &= \Sigma(w_{p\bar{q}}) = \{\{w_{pq}\}, \{w_{q\bar{q}}\}\}^\downarrow \\ \Sigma(w_{\bar{p}q}) &= \Sigma(w_{\bar{p}\bar{q}}) = \{\{w_{\bar{p}q}, w_{\bar{p}\bar{q}}\}\}^\downarrow \end{aligned}$$

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

²In inquisitive semantics, the term *issue* is used to refer to the content of a question. For instance, the issues that a detective entertains might be those expressed by the questions *who committed the murder*, *whether they had an accomplice*, and *what the motive is*.




FIGURE 1. A single-agent inquisitive epistemic model

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 $w_{p\bar{q}}$, the agent’s knowledge state is $\{w_{pq}, w_{p\bar{q}}\}$: 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 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 $w_{\bar{p}q}$ and $w_{\bar{p}\bar{q}}$, the agent’s knowledge state is $\{w_{\bar{p}q}, w_{\bar{p}\bar{q}}\}$: 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:

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

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.³ We treat negation and disjunction as defined connectives (syntactic shorthands) according to

$$\neg \varphi := \varphi \rightarrow \perp \quad \text{and} \quad \varphi \vee \psi := \neg(\neg \varphi \wedge \neg \psi)$$

so that the above syntax emulates standard propositional formulae in terms of atoms and connectives \wedge and \rightarrow together with the defined \neg and \vee . The semantics of these 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. For instance, we read the formula $p \vee \neg p$ as the question *whether or not p*, and we abbreviate this formula as $?p$. Our language also 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 $\boxplus ?p$ intuitively says, roughly, that the agent is interested in the issue whether p .

³In [11, 2] the modalities \Box_a and \boxplus_a are denoted K_a and E_a , and read as “know” and “entertain” respectively.

As mentioned above, the semantics of INQML is given in terms of support in an information state, rather than truth at a possible world.⁴

DEFINITION 2.6. [semantics of INQML]

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

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

If a state s can be extended consistently to a state that supports a formula φ , we say that s is *compatible* with φ :

$$s \text{ is compatible with } \varphi \quad \text{iff} \quad \exists t \subseteq s : t \neq \emptyset \text{ and } \mathbb{M}, t \models \varphi,$$

The derived clauses for the defined connectives \neg and \vee then read as follows:

- $\mathbb{M}, s \models \neg \varphi \iff s \text{ is not compatible with } \varphi$
- $\mathbb{M}, s \models \varphi \vee \psi \iff \forall t \subseteq s, t \neq \emptyset : t \text{ is compatible with } \varphi \text{ or with } \psi$

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 by those information states that settle whether or not p is true.

PROPOSITION 2.7. *The following 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 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.

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:

⁴This means that INQML fits within the quickly growing family of logics based on a *team semantics*. See footnote 1 on page 4 and the references therein.

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

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*.⁵

DEFINITION 2.9. [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 [3], we view truth-conditional formulae as *statements*, and non-truth-conditional formulae as *questions*. The next proposition identifies a large class of formulae that are truth-conditional.

PROPOSITION 2.10. *Atomic formulae (including \perp) and all formulae of the form $\Box\varphi$ and $\boxplus\varphi$ 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 \boxplus , receive exactly the same truth-conditions as in standard modal logic.

PROPOSITION 2.11. *If φ is a formula not containing \vee or \boxplus , then we have $\mathbb{M}, w \models \varphi$ if and only if $\mathfrak{R}(\mathbb{M}), w \models \varphi$ holds 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.

EXAMPLE 2.12. Consider again the model of Example 2.5. The agent's knowledge state at world w_{pq} is $\sigma(w_{pq}) = \{w_{pq}, w_{p\bar{q}}\}$. Since $\{w_{pq}, w_{p\bar{q}}\}$ does not support $?q$ we have $\mathbb{M}, w \models \neg\Box?q$. On the other hand, since the agent's inquisitive state is $\Sigma(w_{pq}) = \{\{w_{pq}\}, \{w_{p\bar{q}}\}\}^\downarrow$, 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 ($\neg\Box?q$), but she's interested in finding out ($\boxplus?q$). By contrast, one can check that at world $w_{\bar{p}q}$ we have $\mathbb{M}, w_{\bar{p}q} \models \neg\Box?q \wedge \neg\boxplus?q$, witnessing that at this world, the agent is neither informed about whether q , nor interested in finding out.

⁵In team semantic terminology (e.g., [32, 36]), truth-conditional formulae are called *flat*.

2.4. Defining properties of worlds and states. Inquisitive modal formulae can be interpreted both relative to information states and (derivatively) relative to worlds. They can thus be seen both as expressing properties of state-pointed models, and as expressing properties of world-pointed models. We can identify these properties with the corresponding classes:

- $\mathcal{K}_\varphi^w = \{(M, w) : M, w \models \varphi\}$
- $\mathcal{K}_\varphi^s = \{(M, s) : M, s \models \varphi\}$

More generally, by a *property* of world- or state-pointed models we mean a class of such objects. We say that a property \mathcal{K} of world-pointed models is definable in INQML if $\mathcal{K} = \mathcal{K}_\varphi^w$ for some formula φ of INQML. Similarly, a property \mathcal{K} of state-pointed models is definable in INQML if $\mathcal{K} = \mathcal{K}_\varphi^s$ for some φ .

We can now formulate the main question that we will adress in this paper: which properties of world- or state-pointed models are definable in INQML?

For the case of state-pointed models, persistency and the semantic ex-falso condition (Proposition 2.7) impose an immediate constraint: in order for a property \mathcal{K} of state-pointed models to be definable in INQML, \mathcal{K} must be an *inquisitive property*, in the following sense.

DEFINITION 2.13. [inquisitive properties]

A property \mathcal{K} of state-pointed models is an *inquisitive* property if the following two conditions hold:

- (i) if $(\mathbb{M}, s) \in \mathcal{K}$ and $t \subseteq s$, then $(\mathbb{M}, t) \in \mathcal{K}$;
- (ii) $(\mathbb{M}, \emptyset) \in \mathcal{K}$ for any model \mathbb{M} .

In the rest of the paper, when dealing with properties of state-pointed models, we can restrict our attention to inquisitive properties.

What features must a world-property have in order to be INQML definable? Similarly, what features must an inquisitive state-property have? The following two sections provide a precise answer to this question.

§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 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, and 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.⁶

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

⁶This definition of bisimilarity between states is reminiscent of the corresponding definition given in [23] for modal team logic. Like inquisitive modal logic, modal team logic interprets formulae with respect to sets of possible worlds, and thus can be seen as expressing properties of state-pointed models. However, there are some major differences with the present setting. Most importantly, the structures for modal team logic are standard Kripke models. By contrast, INQML is interpreted on models having a richer structure; information states enter the picture not just as evaluation points, but also in determining the structure of the model itself, since each world is associated with a set $\Sigma(w)$ of “successors” which are not worlds, but information states. This difference is reflected in the respective notions of bisimulation. In modal team logic, bisimilarity between worlds is the standard notion, and bisimilarity between states is a simple derivative of it: two states are bisimilar if each world in the one is bisimilar to some world in the other. By contrast, in our setting, world-bisimilarity and state-bisimilarity are inextricably intertwined. It is helpful to view this in terms of the bisimulation game. The game for modal team logic starts with a pair of information states; player **I** selects a world from either state, and player **II** responds with a world in the other; after that, the standard bisimulation game for modal logic is played. Thus, information states play a very limited role: they only matter for the initial move, and moreover, there is no move where players have to pick an information state. By contrast, in the case of inquisitive modal logic, the game alternates indefinitely between world positions and state positions, and moves in which players pick information states are a crucial part of the game.

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 $s' \in \Sigma'(w')$ s.t. $s\overline{Y}s'$
 - for all $s' \in \Sigma'(w')$ there is $s \in \Sigma(w)$ 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, $Z^w := Z \cap (W \times W')$, 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 $Z^w \cup \overline{Z^w}$, but not necessarily identical to it. Thus, a bisimulation is not uniquely determined by its restriction to worlds. Rather, given a world-bisimulation Y , the bisimulation $Y \cup \overline{Y}$ is the largest among the bisimulations Z with $Z^w = Y$.

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

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. $\mathbb{M}, x \sim^n \mathbb{M}', x'$ iff there exists a back&forth system $(Z_i)_{i \leq n}$ such that xZ_nx' .*

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, INQML_n denotes 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 atomic propositions, n -bisimilarity coincides with logical indistinguishability in INQML_n , which we denote as \equiv_{INQML}^n :

$$\mathbb{M}, s \equiv_{\text{INQML}}^n \mathbb{M}', s' :\iff \left\{ \begin{array}{l} \mathbb{M}, s \models \varphi \iff \mathbb{M}', s' \models \varphi \\ \text{for all } \varphi \in \text{INQML}_n. \end{array} \right.$$

THEOREM 4.1 (Ehrenfeucht–Fraïssé 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) $\mathbb{M}, s \sim^n \mathbb{M}', s' \iff \mathbb{M}, s \equiv_{\text{INQML}}^n \mathbb{M}', s'$
- (ii) $\mathbb{M}, w \sim^n \mathbb{M}', w' \iff \mathbb{M}, w \equiv_{\text{INQML}}^n \mathbb{M}', w'$

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 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 $\chi_{\mathbb{M}, w}^n$ and $\chi_{\mathbb{M}, s}^n$ of modal depth n such that:

- (i) $\mathbb{M}', w' \models \chi_{\mathbb{M}, w}^n \iff \mathbb{M}', w' \sim^n \mathbb{M}, w$
- (ii) $\mathbb{M}', s' \models \chi_{\mathbb{M}, s}^n \iff \mathbb{M}', s' \sim^n \mathbb{M}, t \text{ for some } t \subseteq s$

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 $\chi_{\mathbb{M}, w}^n$ and $\chi_{\mathbb{M}, s}^n$ together with auxiliary formulae $\chi_{\mathbb{M}, \Pi}^n$ 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:

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

As a special case, we have $\chi_\emptyset^n = \perp$ (as $\bigvee \emptyset \equiv \perp$). 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 (indeed, it is easy to check that, for finite \mathcal{P} , INQML_n is finite up to logical equivalence). Note that, by Proposition 2.10, the formulae χ_w^n and χ_s^n are truth-conditional for all n .

We show the following:

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

We first show that, for each individual n , (1) \Rightarrow (2) \Rightarrow (3). The three claims are then established together by induction on n .

For (1) \Rightarrow (2), suppose $\mathbb{M}', s' \models \chi_{\mathbb{M}, s}^n$. By persistency (Proposition 2.7), $\chi_{\mathbb{M}, s}^n$ is true at each $w' \in s'$; that is, for all $w' \in s'$ we have $\mathbb{M}', w' \models \bigvee \{\chi_{\mathbb{M}, w}^n : w \in s\}$. Since connectives have the standard behaviour in terms of truth-conditions, this means 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 2.10), it follows that $\mathbb{M}', s' \models \chi_{\mathbb{M},s}^n$.

For (2) \Rightarrow (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 $\mathbb{M}', s' \models \chi_{\mathbb{M},\Pi}^n$.

We can now show (1) (and thus (2) and (3)) for all $n \in \mathbb{N}$ by induction. The claim $\mathbb{M}', w' \models \chi_{\mathbb{M},w}^0 \Leftrightarrow \mathbb{M}', w' \sim_0 \mathbb{M}, w$ 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 $\Pi \not\sim_n \Sigma(w)$.

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

For (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 $\mathbb{M}', w' \models \boxplus \chi_{\mathbb{M},\Sigma(w)}^n$.

For (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 claim (1). 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 $s \in \Sigma(w)$;
- (iii) any $s \in \Sigma(w)$ is n -bisimilar to some $s' \in \Sigma'(w')$.

For (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' satisfy the same atomic formulae.

For (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 induction hypothesis, this means that any $s' \in \Sigma'(w')$ is n -bisimilar to some $s \in \Sigma(w)$.

In preparation for (iii), consider the set Π of states in $\Sigma(w)$ that are n -bisimilar to some $s' \in \Sigma'(w')$. We have already seen that any 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 $\mathbb{M}', s' \models \chi_{\mathbb{M}, \Pi}^n$. As this holds for every $s' \in \Sigma'(w')$, we have $\mathbb{M}', w' \models \boxplus \chi_{\mathbb{M}, \Pi}^n$.

Now suppose towards a contradiction that, contrary to (iii), some $s \in \Sigma(w)$ were not n -bisimilar to any state in $\Sigma'(w')$. Then s could not be n -bisimilar to any state in Π either. This implies that $\Pi \not\sim_n \Sigma(w)$ so that $\neg \boxplus \chi_{\mathbb{M}, \Pi}^n$ would be a conjunct of $\chi_{\mathbb{M}, w}^{n+1}$. 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. \dashv

It is now easy to prove the non-trivial direction of Theorem 4.1.

PROOF OF THEOREM 4.1. We focus on the left-to-right direction in claim (i) of Theorem 4.1: the converse follows from the observation that INQML-formulae of depth up to n are invariant under n -bisimilarity, and claim (ii) follows from (i) by specialisation to singleton states. So 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, suppose 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 $\mathbb{M}, s \not\equiv_{\text{INQML}}^n \mathbb{M}', s'$. \dashv

As a corollary of Theorem 4.1, we have the following characterisation of properties definable in INQML.

COROLLARY 4.3. *A property of world-pointed models (resp., an inquisitive property of state-pointed models) over a finite set \mathcal{P} of atomic propositions is definable in INQML if and only if it is closed under \sim^n for some $n \in \mathbb{N}$.*

PROOF. If a property \mathcal{K} of pointed models is defined by a formula φ of depth n , then, since φ is invariant under n -bisimilarity, \mathcal{K} is closed under \sim^n .

Conversely, suppose that \mathcal{K} is a property of world-pointed models closed under \sim^n . Using Proposition 4.2 it is easy to show that \mathcal{K} is defined by the formula $\chi_n^{\mathcal{K}} := \bigvee \{ \chi_{\mathbb{M}, w}^n : (\mathbb{M}, w) \in \mathcal{K} \}$. Notice that the disjunction is well defined, since for a given n there are only finitely many distinct formulae of the form $\chi_{\mathbb{M}, w}^n$.

Similarly, if \mathcal{K} is an inquisitive property of state-pointed models closed under \sim^n , it follows from Proposition 4.2 that \mathcal{K} is defined by the inquisitive disjunction $\chi_n^{\mathcal{K}} = \bigvee \{ \chi_{\mathbb{M}, s}^n : (\mathbb{M}, s) \in \mathcal{K} \}$. \dashv

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 [29] 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, $V : \mathcal{P} \rightarrow \wp(W)$ is a propositional valuation, and $\Sigma : W \rightarrow \wp \wp(W)$, 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 [18], 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 big 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.⁷

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:

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

This property is clearly not invariant under inquisitive bisimulations (indeed, it is not preserved under disjoint unions!). Thus, it is not expressible in INQML.

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.

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

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

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

⁷In 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.

Given any $w \in W$, we have $\mathbb{M}_1, w \models \boxplus?p$ but $\mathbb{M}_2, w \not\models \boxplus?p$. However, the set $\{v, u\}$ is not the truth-set of any formula in neighbourhood semantics: the reason is that u and u' are indistinguishable, and so a truth-set always contains either both of them, or neither. Using this fact, we can show by induction that $\mathbb{M}_1, w \models_{\text{nhd}} \varphi \iff \mathbb{M}_2, w \models_{\text{nhd}} \varphi$ for all formulae φ . Hence, the property expressed by $\boxplus?p$ 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. For instance, consider again the above models \mathbb{M}_1 and \mathbb{M}_2 : according to the notion of bisimilarity \sim^N appropriate for neighbourhood semantics [19], the relation $R = \{(v, v), (u, u), (u, u'), (u', u), (u', u')\}$ is a bisimulation. This implies that $\mathbb{M}_1, v \sim^N \mathbb{M}_2, v$. 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 [18]. 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 [18], **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 $s' = \emptyset$.

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 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 we compare the expressive power of inquisitive modal logic with that of first-order logic. This 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 for the interpretation of each atomic proposition $p \in \mathcal{P}$. 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. We thus turn to two-sorted structures. In order to capture the fact that the second sort contains sets of elements of the first sort, our relational structures include a relation ϵ between these sorts, which simulates set-theoretic membership. This leads to the following notion.

6.1. Relational inquisitive modal models.

DEFINITION 6.1. [relational models]

A *relational inquisitive modal model* over atomic propositions $\mathcal{P} = \{p_i : i \in I\}$ 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 $P_i \subseteq W$ for $i \in I$.

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

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

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

By a *world (resp. state)-pointed relational model* we mean a pair (\mathfrak{M}, x) where x is an element in the first (resp. second) sort of \mathfrak{M} .

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, $\{\underline{s} : s \in S\} \subseteq \wp(W)$. 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 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:

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

so that $R[w] := \{w' : wRw'\} = \bigcup E[w]$ as the natural relational encoding of the map $\sigma : w \mapsto \bigcup \Sigma(w)$.

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 $\wp(R[w]) \subseteq S$ for all $w \in W$.

Note that, by local powerset, the condition of local fullness is equivalent to the condition that the information states $R[w] = \bigcup E[w]$ are represented in S for all $w \in W$.

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

$$\mathfrak{M}^* = (W, \Sigma, V) \quad \text{where } \Sigma(w) = E[w], V(p_i) = P_i.$$

Notice that the non-emptiness and downward closure conditions on E guarantee that \mathfrak{M}^* is indeed an inquisitive modal model. Since the passage from \mathfrak{M} to \mathfrak{M}^* 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 W , E and the P_i . But this leaves quite some choice with respect to the richness of the second sort. The following isolates some immediate choices.

DEFINITION 6.4. [relational encodings]

Given an inquisitive modal model $\mathbb{M} = (W, \Sigma, V)$, we define three relational encodings $\mathfrak{M}^{[\dots]}(\mathbb{M})$ 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 $\mathfrak{M}^{\text{rel}}(\mathbb{M})$, the minimal encoding of \mathbb{M} :
 $S := \text{image}(\Sigma)$;
- for $\mathfrak{M}^{\text{lf}}(\mathbb{M})$, the minimal locally full encoding of \mathbb{M} :
 $S := \{s \subseteq \sigma(w) : w \in W\}$;
- for $\mathfrak{M}^{\text{full}}(\mathbb{M})$, 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 InqML. The notions of $(n-)$ bisimilarity defined in Section 3 naturally lifts to relational models as follows.

DEFINITION 6.5. [(n-)bisimilarity for relational models]

Two state- or world-pointed relational models \mathfrak{M}_1, x_1 and \mathfrak{M}_2, x_2 are bisimilar, $\mathfrak{M}_1, x_1 \sim \mathfrak{M}_2, x_2$, if they encode bisimilar pointed models, i.e. if $\mathfrak{M}_1^*, x_1 \sim \mathfrak{M}_2^*, x_2$. Similarly for n -bisimilarity: $\mathfrak{M}_1, x_1 \sim_n \mathfrak{M}_2, x_2$ if $\mathfrak{M}_1^*, x_1 \sim_n \mathfrak{M}_2^*, x_2$.

In particular, any two encodings of the same pointed inquisitive model are bisimilar. We similarly lift the interpretation of InqML to relational models.

DEFINITION 6.6. Let \mathfrak{M}, x be a (world- or state-)pointed relational model and φ a formula of InqML. We define $\mathfrak{M}, x \models \varphi$ to mean $\mathfrak{M}^*, x \models \varphi$.

By a *property* of world-pointed relational models we mean a class of world-pointed relational models. Similarly, by an *inquisitive property* of state-pointed relational models we mean a class of state-pointed relational models which satisfies the analogues of the conditions in Definition 2.13. By a world- or state-property over a class of relational models \mathcal{C} we mean a class of pointed models (\mathfrak{M}, x) with $\mathfrak{M} \in \mathcal{C}$. We can then translate Corollary 4.3 into a characterisation of the properties of pointed relational models that are definable in InqML.

COROLLARY 6.7. Let \mathcal{K} be a property of world-pointed relational models, or an inquisitive property of state-pointed relational models, over a finite set \mathcal{P} of

atomic propositions. Then \mathcal{K} is definable in INQML if and only if it is closed under \sim^n for some $n \in \mathbb{N}$. More generally, if \mathcal{K} is a world property or an inquisitive state property over a class \mathcal{C} of relational models, \mathcal{K} is definable in INQML over \mathcal{C} iff it is closed under \sim^n within \mathcal{C} for some $n \in \mathbb{N}$.

PROOF. It follows from Theorem 4.1 that if a class \mathcal{K} is defined by an INQML formula φ , then \mathcal{K} is closed under \sim_n where n is the modal depth of φ .

Conversely, suppose \mathcal{K} is a property of world-pointed relational models closed under \sim^n for some $n \in \mathbb{N}$. Then $\mathcal{K}^* := \{(\mathfrak{M}^*, w) : (\mathfrak{M}, w) \in \mathcal{K}\}$ is also closed under \sim^n , so by Corollary 4.3 it is definable by a formula φ of INQML. It is easy to check that φ defines \mathcal{K} (note that \sim^n -closure of \mathcal{K} implies that $(\mathfrak{M}, w) \in \mathcal{K} \iff (\mathfrak{M}^*, w) \in \mathcal{K}^*$). If $\mathcal{K} \subseteq \mathcal{C}$ is closed under \sim^n in restriction to \mathcal{C} , we may similarly work with $\mathcal{K}^* := \{(\mathfrak{M}^*, w) : \mathfrak{M}, w \sim^n \mathfrak{M}', w' \text{ for some } (\mathfrak{M}', w') \in \mathcal{K}\}$. The reasoning for inquisitive state-pointed properties is exactly analogous, using the second part of Corollary 4.3. \dashv

6.4. Relational models and first-order logic. A relational inquisitive model supports a two-sorted first-order language having two relation symbols E and ϵ corresponding to the relations E and ε , respectively, and predicate symbols P_i for $i \in I$. We use $\mathbf{w}, \mathbf{v}, \mathbf{u}$ as variables for the first sort, and \mathbf{s}, \mathbf{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 :

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

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

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

In terms of this language we can define a pair of standard translations $\text{ST}_{\mathbf{w}}(\varphi)$ of $\text{ST}_{\mathbf{s}}(\varphi)$ of a formula, which capture its truth conditions in a world and its support conditions in an information state, respectively. Correspondingly, $\text{ST}_{\mathbf{w}}(\varphi)$ has a single free variable \mathbf{w} of the first sort while $\text{ST}_{\mathbf{s}}(\varphi)$ has the free variable \mathbf{s} of the second sort. Of $\text{ST}_{\mathbf{s}}(\varphi)$ we also use a substitution variant $\text{ST}_{\mathbf{t}}(\varphi)$ which is just like $\text{ST}_{\mathbf{s}}(\varphi)$ except that the roles of variables \mathbf{s} and \mathbf{t} are exchanged. The following define these standard translations by simultaneous induction:

- $\text{ST}_{\mathbf{w}}(p_i) = P_i(\mathbf{w})$
 $\text{ST}_{\mathbf{s}}(p_i) = \forall \mathbf{w}(\epsilon(\mathbf{w}, \mathbf{s}) \rightarrow \text{ST}_{\mathbf{w}}(p_i))$
- $\text{ST}_{\mathbf{w}}(\perp) = \perp$
 $\text{ST}_{\mathbf{s}}(\perp) = \forall \mathbf{w}(\epsilon(\mathbf{w}, \mathbf{s}) \rightarrow \text{ST}_{\mathbf{w}}(\perp)) \quad (\equiv \neg \exists \mathbf{w} \epsilon(\mathbf{w}, \mathbf{s}))$
- $\text{ST}_{\mathbf{w}}(\varphi \wedge \psi) = \text{ST}_{\mathbf{w}}(\varphi) \wedge \text{ST}_{\mathbf{w}}(\psi)$
 $\text{ST}_{\mathbf{s}}(\varphi \wedge \psi) = \text{ST}_{\mathbf{s}}(\varphi) \wedge \text{ST}_{\mathbf{s}}(\psi)$
- $\text{ST}_{\mathbf{w}}(\varphi \vee \psi) = \text{ST}_{\mathbf{w}}(\varphi) \vee \text{ST}_{\mathbf{w}}(\psi)$
 $\text{ST}_{\mathbf{s}}(\varphi \vee \psi) = \text{ST}_{\mathbf{s}}(\varphi) \vee \text{ST}_{\mathbf{s}}(\psi)$
- $\text{ST}_{\mathbf{w}}(\varphi \rightarrow \psi) = \text{ST}_{\mathbf{w}}(\varphi) \rightarrow \text{ST}_{\mathbf{w}}(\psi)$
 $\text{ST}_{\mathbf{s}}(\varphi \rightarrow \psi) = \forall \mathbf{t}(\mathbf{t} \subseteq \mathbf{s} \rightarrow (\text{ST}_{\mathbf{t}}(\varphi) \rightarrow \text{ST}_{\mathbf{t}}(\psi)))$

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

- $\text{ST}_w(\boxplus\varphi) = \forall s(\mathbf{E}(w, s) \rightarrow \text{ST}_s(\varphi))$
 $\text{ST}_s(\boxplus\varphi) = \forall w(\epsilon(w, s) \rightarrow \text{ST}_w(\boxplus\varphi))$
- $\text{ST}_w(\Box\varphi) = \forall s(\mathbf{e}(w, s) \rightarrow \text{ST}_s(\varphi))$
 $\text{ST}_s(\Box\varphi) = \forall w(\epsilon(w, s) \rightarrow \text{ST}_w(\Box\varphi))$

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 $\text{ST}_w(\varphi)$ and $\text{ST}_s(\varphi)$ in any locally full relational encoding of \mathbb{M} .

PROPOSITION 6.8. *Let \mathbb{M} be a locally full relational inquisitive model, $\varphi \in \text{INQML}$. For all worlds $w \in W$ and all states $s \in S$:*

- (i) $\mathbb{M}, w \models \varphi \iff \mathbb{M}, w \models \text{ST}_w(\varphi)$
- (ii) $\mathbb{M}, s \models \varphi \iff \mathbb{M}, s \models \text{ST}_s(\varphi)$

The assumption that \mathbb{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 \mathbb{M} . If so, there will be no state $s \in S$ satisfying $\epsilon(w, s)$, which means that $\text{ST}_w(\Box\varphi)$ will come out as vacuously true at w , regardless of whether or not $\mathbb{M}^*, w \models \Box\varphi$. However, even when \mathbb{M} is not locally full, preservation still holds for all \Box -free formulae, as one can easily verify.

PROPOSITION 6.9. *Let \mathbb{M} be a relational inquisitive model, $\varphi \in \text{INQML}$ a \Box -free formula. Then for all worlds $w \in W$ and all states $s \in S$:*

- (i) $\mathbb{M}, w \models \varphi \iff \mathbb{M}, w \models \text{ST}_w(\varphi)$
- (ii) $\mathbb{M}, s \models \varphi \iff \mathbb{M}, s \models \text{ST}_s(\varphi)$

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 6.10. *For any $\varphi \in \text{INQML}$ there exist first-order formulae $\varphi_w^* := \text{ST}_w(\varphi^*)$ and $\varphi_s^* := \text{ST}_s(\varphi^*)$ such that for any relational inquisitive model \mathbb{M} , world $w \in W$ and $s \in S$:*

- (i) $\mathbb{M}, w \models \varphi \iff \mathbb{M}, w \models \varphi_w^*$
- (ii) $\mathbb{M}, s \models \varphi \iff \mathbb{M}, s \models \varphi_s^*$

The corollary allows us to view INQML as a syntactic fragment of first-order logic, $\text{INQML} \subseteq \text{FO}$, over the class of all relational inquisitive models, just as standard modal logic ML may be regarded as a fragment $\text{ML} \subseteq \text{FO}$ 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. In fact, first-order logic is not compact over this class, as the following example shows.

EXAMPLE 6.11. There is a first-order formula $\varphi(s)$ in a single free variable s of the second sort (information state) which over any relational inquisitive model says of an element s that there are no infinite R -paths inside s . Combining this, for instance, with a formula that says that R in restriction to s defines a

discrete linear ordering with a minimal element, and formulae $\psi_n(\mathbf{s})$ saying that s comprises at least n distinct worlds, we get a violation of compactness.

PROOF. The induced modal accessibility relation R is definable according to

$$R(\mathbf{u}, \mathbf{v}) \leftrightarrow \exists s(E(\mathbf{u}, s) \wedge \epsilon(\mathbf{v}, \mathbf{s})).$$

The local power set condition implies that the entire power set $\wp(s)$ of the designated state s is represented in the second sort of the relational model. So the following formula faithfully emulates the standard monadic second-order formalisation of the relevant property:

$$\varphi(\mathbf{s}) := \neg \exists \mathbf{t} (\mathbf{t} \subseteq \mathbf{s} \wedge \exists \mathbf{u} \epsilon(\mathbf{u}, \mathbf{t}) \wedge \forall \mathbf{u} (\epsilon(\mathbf{u}, \mathbf{t}) \rightarrow \exists \mathbf{v} (\epsilon(\mathbf{v}, \mathbf{t}) \wedge R(\mathbf{u}, \mathbf{v}))))).$$

where again $\mathbf{t} \subseteq \mathbf{s}$ abbreviates $\forall \mathbf{v} (\epsilon(\mathbf{v}, \mathbf{t}) \rightarrow \epsilon(\mathbf{v}, \mathbf{s}))$. \dashv

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.

In an extension and variation of the above, Silke Meißner [24] has proposed an alternative standard translation, which in some way is more uniform as it allows for a direct treatment of \Box . As outlined in [25], it also relaxes the constraints on relational encodings so as to extend the scope of the standard translation to an elementary class of relational structures, which in turns gives rise to a model-theoretic compactness proof for INQML. While our translation could in principle be replaced by the more recent one from [25], adherence to our narrower classes of natural relational encodings of the *intended* inquisitive models can be seen as a strength of our characterisation theorems.

§7. Bisimulation invariance.

7.1. Bisimulation invariance as a semantic constraint. As discussed above, INQML can be thought of as a fragment of first-order logic when interpreted over relational models. We may think of \sim -invariance as a characteristic semantic feature of this fragment. The question we are interested in is: with respect to what classes \mathcal{C} of relational models can INQML be characterised as being *exactly* the \sim -invariant fragment of first-order logic? Let us first make precise what this means.

DEFINITION 7.1. We say that INQML is the \sim -invariant fragment of FO for world-properties with respect to a class \mathcal{C} of relational models, in symbols

$$\text{INQML} \equiv_{\mathcal{C}}^w \text{FO}/\sim,$$

if for every property \mathcal{K} of state-pointed models over \mathcal{C} , \mathcal{K} is definable in INQML if and only if it is both definable in FO and \sim -invariant.

Similarly, we say that INQML is the \sim -invariant fragment of FO for inquisitive state-properties with respect to \mathcal{C} , in symbols

$$\text{INQML} \equiv_{\mathcal{C}}^s \text{FO}/\sim,$$

if for any inquisitive property \mathcal{K} of state-pointed models over \mathcal{C} , \mathcal{K} is definable in INQML if and only if it is definable in FO and \sim -invariant.

REMARK 7.2. For any class \mathcal{C} , $\text{INQML} \equiv_{\mathcal{C}}^s \text{FO}/\sim$ implies $\text{INQML} \equiv_{\mathcal{C}}^w \text{FO}/\sim$.

PROOF. With the \sim -invariant world property defined by $\varphi(\mathbf{w}) \in \mathbf{FO}$ associate the property defined by $\varphi'(\mathbf{s}) = \forall \mathbf{w}(\mathbf{w} \in \mathbf{s} \rightarrow \varphi(\mathbf{w}))$. This property is inquisitive and \sim -invariant. From $\text{INQML} \equiv_{\mathcal{C}}^s \mathbf{FO}/\sim$ we obtain a formula $\psi \in \text{INQML}$ expressing this property. Then the same formula ψ , on the level of worlds, defines the world-property defined by $\varphi(\mathbf{w})$. \dashv

Our question then can be formulated succinctly as follows: over which classes \mathcal{C} do we have $\text{INQML} \equiv_{\mathcal{C}}^s \mathbf{FO}/\sim$ (and thus also $\text{INQML} \equiv_{\mathcal{C}}^w \mathbf{FO}/\sim$)? Equivalently, the question is: over which classes \mathcal{C} is INQML sufficiently expressive to capture *all* first-order definable properties of world- or state-pointed relational models that are invariant under inquisitive bisimulation?

Section 8 will establish that INQML is expressively complete for \sim -invariant first-order properties in this sense over each of the following classes of relational models: the class of all relational models, the class of all finite relational models, the class of all locally full models, and the class of all finite locally full models. Before delving into the proof, however, we discuss some underlying model-theoretic concerns and limitations. In particular we 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 dealing with first-order logic over non-elementary classes of models.

7.2. Bisimulation invariance and compactness. Let \mathcal{K} be a property of world-pointed relational models over a class \mathcal{C} . Suppose \mathcal{K} is INQML -definable by a formula φ : then \mathcal{K} is both \mathbf{FO} -definable (by the standard translation φ^*) and \sim -invariant (as it is \sim_n invariant, with n the modal depth of φ). Thus, one direction of the equivalence $\text{INQML} \equiv_{\mathcal{C}}^w \mathbf{FO}/\sim$ holds for any class \mathcal{C} . By Corollary 6.7, the converse direction amounts to the claim that if \mathcal{K} is \mathbf{FO} -definable and \sim -invariant, then it is in fact \sim_n -invariant for some n .⁹ That is, it amounts to the claim that, for every \mathbf{FO} -definable property \mathcal{K} over \mathcal{C} , \sim -invariance implies \sim_n -invariance for some n .

Analogous reasoning establishes the same connection w.r.t. inquisitive state-properties. We summarise these in the following, where a first-order formula $\varphi(\mathbf{s})$ in a free variable of the second sort is called *inquisitive* over the class \mathcal{C} if the state-property expressed by $\varphi(\mathbf{s})$ over \mathcal{C} is an inquisitive property, i.e. is downward closed and always holds of the empty state.

OBSERVATION 7.3. *For any class \mathcal{C} of relational inquisitive models, the following are equivalent:*

- (i) $\text{INQML} \equiv_{\mathcal{C}}^w \mathbf{FO}/\sim$,
- (ii) *for any formula $\varphi(\mathbf{w}) \in \mathbf{FO}$ in a single free variable of the first sort, \sim -invariance over \mathcal{C} implies \sim^n -invariance over \mathcal{C} for some finite n .*

Similarly, the following are equivalent:

- (i) $\text{INQML} \equiv_{\mathcal{C}}^s \mathbf{FO}/\sim$,
- (ii) *for any formula $\varphi(\mathbf{s}) \in \mathbf{FO}$ in a single free variable of the second sort that is inquisitive over the class \mathcal{C} , \sim -invariance over \mathcal{C} implies \sim^n -invariance over \mathcal{C} for some finite n .*

⁹Notice that if \mathcal{K} is \mathbf{FO} -definable, the defining formula contains only finitely many atoms. Thus, the property \mathcal{K} depends only on the restriction of a model to a finite set \mathcal{P} of atoms, and we can use Corollary 6.7 to conclude that if \mathcal{K} is \sim_n -invariant for some n , it is INQML -definable.

In both contexts, condition (ii) 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. Interestingly, this compactness principle for \sim -invariance of FO-properties of worlds fails relative to the class of *full* relational models.

PROPOSITION 7.4. *There is a first-order formula $\varphi(\mathbf{w})$ in a single free variable \mathbf{w} of the first sort that, relative to the class of full relational models, is \sim -invariant but not \sim^n -invariant for any n .*

PROOF. Compare Example 6.11 for the following well-foundedness property:

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

On one hand \mathcal{P} clearly is \sim invariant but not \sim^n -invariant for any n . On the other hand \mathcal{P} is first-order definable over the class of full relational inquisitive models since, over these models, first-order logic affords the full expressive power of monadic second-order quantification over the first sort W : first-order quantification over the second sort $S = \wp(W)$ is quantification over subsets of the first sort. The formula

$$\varphi(\mathbf{w}) := \neg \exists \mathbf{s} (\epsilon(\mathbf{w}, \mathbf{s}) \wedge \forall \mathbf{u} (\epsilon(\mathbf{u}, \mathbf{s}) \rightarrow \exists \mathbf{v} (\epsilon(\mathbf{v}, \mathbf{s}) \wedge R(\mathbf{u}, \mathbf{v}))))).$$

defines the world-property \mathcal{P} over any full relational model. \dashv

A similar well-foundedness property can 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 distinct agents $a, b \in \mathcal{A}$. This shows that $\text{INQML} \not\equiv_{\mathcal{C}}^w \text{FO}/\sim$ when \mathcal{C} is the class of full relational models or full relational epistemic models. Over these classes, there are FO-definable, \sim -invariant world properties that are not INQML-definable. Although this is in sharp contrast with our Theorem 1.2, the fact that the analogue of the theorem fails 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 section, we establish semantic correspondences:

$$\text{INQML} \equiv_{\mathcal{C}}^w \text{FO}/\sim \qquad \text{INQML} \equiv_{\mathcal{C}}^s \text{FO}/\sim$$

These are assertions about equal expressive power between two systems presented in very different style: while INQML is based on concrete syntax with clearly defined semantics, FO/\sim is defined in terms of the semantic constraint of \sim -invariance. In fact, \sim -invariance is easily seen to be undecidable as a property of first-order formulae, so that FO/\sim itself cannot be regarded as a syntactic fragment. As discussed in the previous subsection, proving one of these equivalences boils down to establishing a compactness principle relating \sim -invariance to \sim^n -invariance for some finite level n , in the non-classical context of non-elementary classes of relational inquisitive models.




FIGURE 2. Generic upgrading patterns.

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 [26] of van Benthem’s characterisation theorem [34] and its finite model theory version due to Rosen [31] (for ramifications of this method, see also [27, 13] and [28]). 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 providing, 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 on the left 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 chases the diagram through its lower rung to check that, for φ that is preserved under both \sim and \equiv_q , we have $\mathfrak{M}, w \models \varphi$ iff $\mathfrak{M}', w' \models \varphi$.

The reasoning for inquisitive properties of information states is analogous, using a corresponding upgrading for state-pointed models (cf. the right hand side in Figure 2). 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.

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 \mathcal{C} !)—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).

We show how to achieve the required upgrades for various classes \mathcal{C} of models in the next section, and thus establish our characterisation theorems.

We use a variation of the upgrading technique from [26] to instantiate the above idea. In effect we shall deviate slightly from the generic picture in Figure 2 by interleaving \sim -preserving pre-processing steps and \equiv_q -preserving steps as shown in Figure 4. The upgrading itself is based on an inquisitive analogue of partial tree unfoldings, combined with locality arguments for first-order Ehrenfeucht–Fraïssé games. We start with two technical remarks.

Essentially disjoint unions and essential parts. 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, the disjoint union of two relational inquisitive models would fail to satisfy extensionality (and thus fail to be a relational inquisitive model) unless we take care of identifying the respective empty states. In the following, if \mathfrak{M} and \mathfrak{M}' are relational models, we denote by $\mathfrak{M} \oplus \mathfrak{M}'$ their *essentially disjoint union*, i.e., the model obtained from the disjoint union by identifying the empty information states of \mathfrak{M} and \mathfrak{M}' . Indeed, the empty information state plays a special, albeit somewhat trivial rôle for various purposes. To isolate the structurally distinctive part of a relational inquisitive model $\mathfrak{M} = (W, S, E, \in, (P_i))$ we may consider its *essential part* as obtained by removal of the empty information state:

$$\mathfrak{M}^\circ := \mathfrak{M} \upharpoonright (W \cup S^\circ) = (W, S^\circ, E^\circ, \in^\circ, (P_i))$$

where $S^\circ := S \setminus \{\emptyset\}$ and $E^\circ, \in^\circ \subseteq W \times S^\circ$ are corresponding restrictions. While \mathfrak{M}° is not itself a relational inquisitive model, it uniquely determines the original model. Writing $\mathfrak{M}^\circ * \{\emptyset\}$ for the unique extension by re-insertion of \emptyset , which reproduces \mathfrak{M} , we have the one-to-one correspondence

$$(\dagger) \quad \mathfrak{M}^\circ = \mathfrak{M} \upharpoonright (W \cup S^\circ) \quad \longleftrightarrow \quad \mathfrak{M} = \mathfrak{M}^\circ * \{\emptyset\}$$

Moreover, essentially disjoint unions of models (or of subsets of their domains) are disjoint unions at the level of the essential parts.

Locality and truncation of models. 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*. 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 \in ; 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. 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 \mathfrak{M} .¹⁰ Passage to the essential part \mathfrak{M}° , however, overcomes this obstacle. The empty state plays a trivial rôle not just w.r.t. bisimulation, where it only occurs as a dead end, but also w.r.t. FO expressiveness: the relational model \mathfrak{M} is uniformly quantifier-free FO-interpretable in its essential part \mathfrak{M}° . It follows that \equiv_q between the essential parts of (pointed) relational inquisitive models implies \equiv_q between the actual models.¹¹ So the correspondence in (\dagger) is compatible with individual levels of FO-equivalence. Meaningful locality arguments can therefore be based on ℓ -neighbourhoods w.r.t. essential parts, where Gaifman distance is not trivialised by \emptyset . If \mathfrak{M} is a relational inquisitive model, w a world in \mathfrak{M} , and ℓ an even

¹⁰In 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 .

¹¹Uniform syntactic rewriting provides for every $\varphi(x) \in \text{FO}$ a translation $\varphi^\circ(x) \in \text{FO}$ of the same quantifier rank such that $\mathfrak{M}, w \models \varphi$ iff $\mathfrak{M}^\circ, w \models \varphi^\circ$.

number, we define *truncation* of \mathfrak{M} to depth ℓ as

$$\mathfrak{M}_w^\ell := (\mathfrak{M}^\circ \upharpoonright N^\ell(w)) * \{\emptyset\},$$

where $N^\ell(w)$ consists of those worlds or states at Gaifman distance at most ℓ from w in \mathfrak{M}° (!). It is easy to see that \mathfrak{M}_w^ℓ is a relational inquisitive model. Similarly, for a state $s \neq \emptyset$ and even ℓ , we define $\mathfrak{M}_s^\ell := (\mathfrak{M}^\circ \upharpoonright N^{\ell+1}(s)) * \{\emptyset\}$, which again is a relational inquisitive model.

§8. Characterisation theorem for InqML. Our aim is to show the following main characterisation theorem.

THEOREM 1.2. *Let \mathcal{C} be any of the following classes of relational models: the class of all models; of finite models; of locally full models; of finite locally full models. Then $\text{InqML} \equiv_{\mathcal{C}}^v \text{FO}/\sim$ and $\text{InqML} \equiv_{\mathcal{C}}^s \text{FO}/\sim$.*

By Observation 7.3 it suffices to establish the following.

PROPOSITION 8.1. *Let \mathcal{C} be any of the above classes. Any first-order formula which is \sim -invariant over \mathcal{C} is \sim^n -invariant over \mathcal{C} for some finite level $n \in \mathbb{N}$.*

8.1. Partial unfolding and stratification. To establish the compactness principle for \sim -invariance expressed in Proposition 8.1 for the relevant classes of relational models we make use of a process of *stratification*. This is similar to tree-like unfoldings in standard modal logic.

DEFINITION 8.2. A relational inquisitive model \mathfrak{M} is *stratified* if its two domains W and S consist of essentially disjoint strata, i.e., of two families $(W_i)_{i \in \mathbb{N}}$ and $(S_i)_{i \in \mathbb{N}}$ such that:

- (i) $W = \bigcup W_i$ and $S = \bigcup S_i$;
- (ii) for each $i, j \in \mathbb{N}$: $W_i \cap W_j = \emptyset$ and $S_i \cap S_j = \{\emptyset\}$;
- (iii) $E[w] \subseteq S_i$ for all $w \in W_i$ and $S_i \subseteq \wp(W_{i+1})$.

For an even $\ell \neq 0$ and a world w , we say that \mathfrak{M} is *stratified to depth ℓ from w* if its truncation \mathfrak{M}_w^ℓ (see the definition at the end of Section 7.3) is stratified with $W_0 = \{w\}$. \mathfrak{M} is *stratified to depth ℓ from a non-empty state $s \in S$* , if its truncation \mathfrak{M}_s^ℓ 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 . An analogous result holds 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.*

PROOF SKETCH. 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. Starting from a model \mathfrak{M} , we define a stratified model \mathfrak{M}' , whose essential part \mathfrak{M}'° consists of \mathbb{N} -tagged copies of worlds and non-empty 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 choose a downward closed set $S_n \supseteq \bigcup_{(u,n) \in W'_n} E[u]$ and put

$$\begin{aligned} S'_n{}^\circ &:= S_n^\circ \times \{n\}, \\ W'_{n+1} &:= \bigcup_{s \in S_n^\circ} s \times \{n+1\}, \\ \text{and define } E'^\circ &:= \{((u,n), (s,n)) : (u,s) \in E\}, \\ \varepsilon'^\circ &:= \{((u,n+1), (s,n)) : u \in s\}, \\ P'_i &:= \{(u,n) : u \in P_i\}. \end{aligned}$$

This uniquely determines $\mathfrak{M}' = \mathfrak{M}'^\circ * \{\emptyset\}$. It is easy to check 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 $S'_{n+1}{}^\circ$ and augment E'° by all of E° . The resulting \mathfrak{M}', w' 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 $S'_n{}^\circ$, viz. $S'_n{}^\circ := S^\circ \times \{n\}$, 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 proceed inductively as above. \dashv

OBSERVATION 8.4. *Let \mathfrak{M}, w and \mathfrak{M}', w' be world-pointed relational models that are stratified to depth ℓ for some even ℓ from their respective worlds. Let \mathfrak{M}_w^ℓ, w and $\mathfrak{M}'_{w'}, w'$ be their ℓ -truncations. Then, for $n \geq \ell/2$:*

$$\mathfrak{M}_w^\ell, w \sim^n \mathfrak{M}'_{w'}^\ell, w' \implies \mathfrak{M}_w^\ell, w \sim \mathfrak{M}'_{w'}^\ell, w'.$$

Analogously for state-pointed models that are stratified to depth ℓ from their distinguished states.


This is because, due to stratification and cut-off, the n -round game exhausts all possibilities in the unbounded game. After m rounds of the bisimulation game, which starts from the pairing of the roots w and w' of the stratified models, the position is a pairing of worlds from strata W_m and W'_m . So player **II** wins the unbounded game if she does not lose within n rounds.

PROOF OF THEOREM 1.2. We first 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 $\varphi(x) \in \text{FO}_q$ 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 4. 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 $\ell = n$ from w . Let \mathfrak{M}_w^ℓ be its ℓ -truncation, which is then fully stratified. 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}_w^ℓ . In addition, \mathfrak{M}_0 contains a copy of \mathfrak{M}_w^ℓ with the distinguished world w , while \mathfrak{M}_1 contains a copy of \mathfrak{M} with the distinguished world w .

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

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

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

FIGURE 3. The structures \mathfrak{M}_0^o, w and \mathfrak{M}_1^o, w in the game argument.

As argued in connection with the correspondence (\dagger) at the end of Section 7, due to quantifier-free interpretability of \mathfrak{M}_i, w in \mathfrak{M}_i^o, w , ($*$) is equivalent to

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

The diagram in Figure 3 suggests the arrangement, with open cones for copies of \mathfrak{M}^o and truncated cones for $\mathfrak{M}^o \upharpoonright N^\ell(w)$ (the essential part of \mathfrak{M}_w^ℓ), and with filled circles for the distinguished worlds.

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, player **II** 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^o and $\mathbf{u}' = (u'_0, u'_1, \dots, u'_m)$ with $u'_0 = w$ in \mathfrak{M}_1^o 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^o and \mathfrak{M}_1^o or in isomorphic configurations in $\mathfrak{M}_0^o \upharpoonright N^\ell(w)$ and $\mathfrak{M}_1^o \upharpoonright N^\ell(w)$.

This condition is satisfied at the start of the game, for $m = 0$ ($\ell_0 = 2^q = n$). 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^o or $u' = u'_{m+1}$ in \mathfrak{M}_1^o 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^o or $u' = u'_{m+1}$ in \mathfrak{M}_1^o 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.




FIGURE 4. Upgrading pattern for Theorem 1.2/Proposition 8.1.

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 4 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 6.7: indeed, φ is logically equivalent over \mathcal{C} to the disjunction over the characteristic formulae $\chi_{\mathfrak{M}, w}^n$ for all $\mathfrak{M}, w \in \mathcal{C}$ that satisfy φ . \dashv

The case of state properties. To show Proposition 8.1 for state properties, we can similarly upgrade the situation $\mathfrak{M}, s \sim^n \mathfrak{M}', s'$ for non-empty s, s' in companion structures through passage to truncations of fully bisimilar models that are stratified to depth ℓ from their distinguished states.¹² Assuming w.l.o.g. that \mathfrak{M}, s is itself stratified to depth $\ell = 2^q$ with stratified restriction \mathfrak{M}_s^ℓ we define as before the following essentially disjoint unions

$$\begin{aligned}\mathfrak{M}_0, s &:= q \otimes \mathfrak{M} \oplus \mathfrak{M}_s^\ell, s \oplus q \otimes \mathfrak{M}_s^\ell \\ \mathfrak{M}_1, s &:= q \otimes \mathfrak{M} \oplus \mathfrak{M}, s \oplus q \otimes \mathfrak{M}_s^\ell\end{aligned}$$

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 4 for the relevant state-pointed models.

§9. Conclusion. We have seen the beginnings of a model theory for inquisitive modal logic. Our contribution started in Section 3, where we described a natural notion of bisimulation for inquisitive modal structures. From a game-theoretic perspective, bisimilarity and its approximations can be characterised

¹²Note that the upgrading argument trivialises for $\varphi(\mathbf{s}) \in \mathbf{FO}$ in the case of state-pointed models \mathfrak{M}, s with $s = \emptyset$: For any n , $\mathfrak{M}, \emptyset \sim^n \mathfrak{M}', s'$ iff $\mathfrak{M}, \emptyset \sim \mathfrak{M}', s'$ iff $s' = \emptyset$. If $\varphi(\mathbf{s})$ is \sim -invariant over \mathcal{C} , it must be satisfied by $s = \emptyset$ across all of \mathcal{C} or nowhere; and being inquisitive over \mathcal{C} , the former must in fact be true.

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 usual way: two pointed models over a finite vocabulary are distinguishable in the n -round bisimulation game iff 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 naturally 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 with respect to those relational structures.

We then asked 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 several natural classes, including the class of all relational inquisitive models, and the class of all 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 $\mathcal{P}(w) := "W \in \Sigma(w)"$ or $\mathcal{P}(w) := "\{w\} \in \Sigma(w)"$ 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 $\Sigma : W \rightarrow \wp\wp(W)$, rather than $\sigma : W \rightarrow \wp(W)$ 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: over various classes of structures, INQML expresses all and only the first-order properties that are invariant in this sense.

In a separate paper, we will tackle the case of inquisitive epistemic models—the inquisitive version of multi-modal *S5* models. Conceptually, this class of models is interesting in light of the natural interpretation of inquisitive modalities in the epistemic setting, as described in Section 2. Technically, it presents interesting challenges as the stratifications used in Section 8 are incompatible with the *S5* frame conditions. Nevertheless, it can be shown that the counterpart of our characterisation result still holds in this setting—again, both in general and in restriction to finite models. A preliminary account is given in [9].

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 Jérôme Lang, editor, *Theoretical Aspects of Rationality and Knowledge (TARK) 16*. EPTCS, 2017.
- [9] I. Ciardelli and M. Otto. Inquisitive bisimulation. arXiv preprint, 57 pages, arXiv:1803.03483, 2018.
- [10] I. Ciardelli and F. Roelofsen. Inquisitive logic. *Journal of Philosophical Logic*, 40:55–94, 2011.
- [11] I. Ciardelli and F. Roelofsen. Inquisitive dynamic epistemic logic. *Synthese*, 192(6):1643–1687, 2015.
- [12] I. Ciardelli, J. Groenendijk, and F. Roelofsen. *Inquisitive Semantics*. Oxford University Press, 2018.
- [13] A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. *Annals of Pure and Applied Logic*, 161:1–42, 2009.
- [14] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Perspectives in Mathematical Logic. Springer, 2 edition, 1999.
- [15] H. Gaifman. On local and nonlocal properties. In J. Stern, editor, *Logic Colloquium 81*, pages 105–135. North Holland, 1982.
- [16] 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.
- [17] 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.
- [18] H. H. Hansen. Monotone modal logics. MSc Thesis, University of Amsterdam, 2003.
- [19] H. H. Hansen, C. Kupke, and E. Pacuit. Neighbourhood structures: Bisimilarity and basic model theory. *Logical Methods in Computer Science*, 2, 2009.
- [20] J. Hintikka. *Knowledge and Belief: an Introduction to the Logic of the Two Notions*. Cornell University Press, 1962.
- [21] W. Hodges. Compositional semantics for a language of imperfect information. *Logic Journal of IGPL*, 5(4):539–563, 1997.
- [22] W. Hodges. Some strange quantifiers. In *Structures in logic and computer science. Lecture notes in computer science*, pages 51–65. Springer, New York, 1997.
- [23] J. Kontinen, J. Steffen Müller, H. Schnoor, and H. Vollmer. A van Benthem Theorem for Modal Team Semantics. In Stephan Kreutzer, editor, *Proceedings of CSL 24*, 277–291. Leibniz International Proceedings in Informatics, 2015.
- [24] S. Meißner, On the Model Theory of Inquisitive Modal Logic, Bachelor Thesis. Department of Mathematics, TU Darmstadt, 2018.
- [25] S. Meißner and M. Otto. A first-order framework for inquisitive modal logic. arXiv preprint, 16 pages, arXiv:1906.04981, 2019.
- [26] M. Otto. Elementary proof of the van Benthem–Rosen characterisation theorem. Technical Report 2342, Fachbereich Mathematik, Technische Universität Darmstadt, 2004.
- [27] M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130:173–205, 2004.
- [28] M. Otto. Highly acyclic groups, hypergraph covers and the guarded fragment. *Journal of the ACM*, 59 (1), 2012.
- [29] E. Pacuit. *Neighborhood Semantics for Modal Logic*. Springer, 2017. forthcoming.

- [30] F. Roelofsen. Algebraic foundations for the semantic treatment of inquisitive content. *Synthese*, 190(1):79–102, 2013.
- [31] E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:427–439, 1997.
- [32] J. Väänänen. *Dependence Logic: A New Approach to Independence Friendly Logic*. Cambridge University Press, 2007.
- [33] 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.
- [34] J. van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Napoli, 1983.
- [35] 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.
- [36] F. Yang and J. Väänänen. Propositional logics of dependence. *Annals of Pure and Applied Logic*, 167:557–589, 2016.