Describing neighborhoods in inquisitive modal logic

Ivano Ciardelli [1]

Munich Center for Mathematical Philosophy
LMU Munich

Abstract

We introduce and investigate an inquisitive modal logic interpreted on neighborhood models. The logic extends propositional inquisitive logic with a binary modality $\Rrightarrow$, where $(\varphi \Rrightarrow \psi)$ is a statement true at a world iff every neighborhood that supports φ also supports ψ. This logic provides a natural language to describe properties of neighborhood structures, has an interpretation as a logic of ability, and generalizes previous versions of inquisitive modal logic. We give an appropriate notion of bisimulation, establish a Hennessy-Milner theorem, and relate a fragment of the logic to the instantial neighborhood logic of [6] by truth-preserving translations. We also prove a completeness theorem, showing that our modal conditional is axiomatized by four simple principles familiar from the study of strict conditionals.

Keywords: Inquisitive logic, neighborhood structures, bisimulation, axiomatization, strict conditional, logic of ability, team semantics, instantial neighborhood logic.

1 Introduction

This paper introduces and investigates an inquisitive modal logic InqCM which extends propositional inquisitive logic with a conditional modality denoted $\Rrightarrow$. This logic can be motivated from three different directions.

One line of motivation starts from considering neighborhood structures. A standard version of neighborhood semantics for modal logic [2,7,8,21,24] interprets a formula $\Box\varphi$ as true at a world w if φ is true throughout some neighborhood s of w (in this case, we say that *s supports* φ). [2] However, the standard modal language with this interpretation does not allow us to express much about the configurations that arise in the neighborhoods of a given point. Consider three worlds w_1, w_2, w_3 associated with sets of neighborhoods $\Sigma(w_i)$ as depicted in Figure 1. Here are some respects in which the situation at these

[1] Gefördert durch die Deutsche Forschungsgemeinschaft (DFG) - Projektnummer 446711878. Funded by the German Research Foundation (DFG) - Project number 446711878.

[2] This is different from the Scott-Montague neighborhood semantics [23,27], which interprets $\Box\varphi$ as true at w if the set of worlds satisfying φ is a neighborhood of w, though the two are equivalent if the set of neighborhoods is closed under supersets. See [24] for discussion.

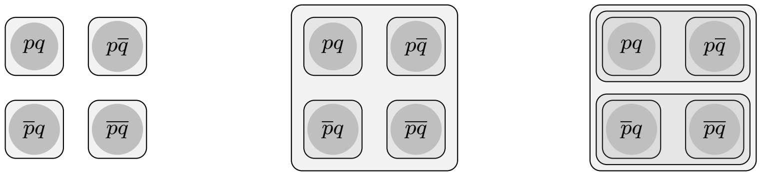

(a) $\Sigma(w_1)$ (b) $\Sigma(w_2)$ (c) $\Sigma(w_3)$

Fig. 1. Sets of neighborhoods associated with three worlds.

worlds is different: every neighborhood of w_1 settles whether p (i.e., the truth value of p is constant within each neighborhood) while this is not the case for w_2 and w_3; moreover, every neighborhood of w_2 that settles whether p also settles whether q, whereas this is not the case for w_3. These facts cannot be expressed in the standard modal language with the above semantics.

More generally, in a neighborhood structure it is natural to regard two worlds as equivalent only if any neighborhood of the one is 'matched' by some neighborhood of the other; in turn, it is natural to regard two neighborhoods as matching if every world in one is equivalent to some world in the other. This leads to a natural notion of bisimilarity for neighborhood structures. It seems worthwhile to pursue a modal language that, at least in restriction to finite models, can distinguish points that are not bisimilar in this sense. The *instantial neighborhood logic* recently developed in [6] is such a language. Modal formulas of this logic have the form $\Box(\psi_1, \ldots, \psi_n; \varphi)$ for $n \geq 0$ and express the existence of a neighborhood s that supports φ and is compatible with each ψ_i.

We will explore a different way to set up a language that can express the relevant properties. Our approach uses a modal conditional $\Rrightarrow$, where a modal formula $\varphi \Rrightarrow \psi$ is true at a world if all neighborhoods supporting φ also support ψ. Crucially, however, φ and ψ are allowed to be questions, and for this reason, the relevant modal conditional is added to an underlying inquisitive logic. In our logic, the notion of support at a set of worlds is taken as primitive, while still boiling down to global truth for a large class of formulas. Thus, e.g., our language will include a formula $?p$ which is supported by a set of worlds s just in case s settles whether p, i.e., if the truth value of p is constant in s. Then the fact that every neighborhood settles whether p can be expressed as $\top \Rrightarrow ?p$ (true at w_1 but not at w_2, w_3), whereas the fact that every neighborhood that settles whether p also settles whether q can be expressed as $?p \Rrightarrow ?q$ (true at w_1, w_2 but not at w_3). One attraction of this approach is that the modal implication $\Rrightarrow$ is very natural both conceptually and formally. Its logic is completely axiomatized by the following four properties, familiar from the study of strict conditionals: (i) if φ entails ψ, then $\varphi \Rrightarrow \psi$ is a logical truth; (ii) $\Rrightarrow$ is transitive; (iii) if φ implies two formulas, it implies their conjunction; (iv) if ψ is implied by two formulas, it is implied by their (inquisitive) disjunction.

A related but more concrete take on our logic arises from thinking about action and ability. Consider three scenarios in which an agent a is deliberating whether to make p and q true or false. In the first scenario, a must make a choice about each of p and q. In the second scenario, a can make both choices, or she can delegate both choices to another agent. In the third scenario, a can make both choices, or make the choice about p and delegate the one about q, or she can delegate both choices. The three scenarios correspond to the pictures in Fig. 1, where each neighborhood corresponds to an action available to a, and the worlds in a neighborhood represent the possible outcomes of that action.

Standard analyses of strategic ability [1,3,4,7,25] focus on the powers of agents to force certain outcomes. From the perspective of these analyses, these situations are the same: in all of them, a is in effect a dictator who can force any of the outcomes, while other agents cannot prevent any outcome. Yet, there is a clear sense in which these situations are very different. In the first scenario, a not only *can*, but also *must* decide on p and q, while in the other two scenarios, she can delegate (which may well be the optimal action—perhaps the other agent is in a better position to make the relevant decisions). Moreover, in the second scenario, a must decide on q if she wants to decide on p, whereas in the third scenario she can can choose on p while delegating the choice on q. Thus, while these situations are the same in terms of what the agent *can* force, they differ in terms of what she *must* force, or what she must force if she wants to force something else. The logic that we develop in this paper will allow us to express these significant facts in a natural way. So, for instance, the fact that the agent is forced to choose on p is expressed by $\top \Rrightarrow ?p$, while the fact that she is forced to choose on q if she wants to choose on p is expressed by $?p \Rrightarrow ?q$.

A third way to view the logic is as an extension of previous work on inquisitive modal logic [9,10,15,17,22,26]. In that work, formulas are interpreted over *downward-closed* neighborhood models, i.e., models where the set of neighborhoods is closed under subsets; the key modality, denoted $\boxplus$, is a universal quantifier over neighborhoods. The language with this modality was shown to be expressively adequate for the appropriate notion of bisimulation in [13]. But while the downward closure constraint is well-motivated under some interpretations of the model (as in inquisitive epistemic logic, [15]), it may not be under other interpretations. This paper thus generalizes inquisitive modal logic by dropping the downward closure requirement. As we show, in this more general setting the universal modality $\boxplus$ is no longer sufficient to have an expressively adequate language. Instead, we obtain an adequate language by taking the modal conditional $\Rrightarrow$ as primitive (one can then define $\boxplus\varphi$ as $(\top \Rrightarrow \varphi)$).

The paper is structured as follows: in Section 2 we introduce and illustrate the logic InqCM; in Section 3 we introduce the notion of bisimulation for which our language is invariant and show a Hennessy-Milner theorem; in Section 4 we describe a Hilbert-style proof system for InqCM; in Section 5 we adapt the standard notion of resolutions and state some useful technical lemmas; in Section 6 we relate a fragment of InqCM to instantial neighborhood logic by defining truth-preserving translations in both directions; in Section 7 we show the completeness of our proof system; finally, in Section 8 we discuss a number of directions in which the present work can be extended.

2 Inquisitive neighborhood logic

In this section we introduce our logic InqCM. Given a set $\mathcal{P}$ of atomic sentences, the language $\mathcal{L}$ of the logic is given by the following BNF definition:

$$\mathcal{L} \qquad \varphi ::= p \mid \bot \mid (\varphi \wedge \varphi) \mid (\varphi \to \varphi) \mid (\varphi \vee\!\!\!\vee \varphi) \mid (\varphi \Rrightarrow \varphi) \qquad p \in \mathcal{P}$$

As standard in inquisitive logic, we also use the following defined connectives: $\neg\varphi := (\varphi \to \bot)$, $\top := \neg\bot$, $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$, $?\varphi := \varphi \vee\!\!\!\vee \neg\varphi$.

The connective $\vee\!\!\!\vee$, called *inquisitive disjunction*, is thought of as a question-forming disjunction. Thus, e.g., whereas $p \vee \neg p$ stands for the tautological statement that p or $\neg p$ is the case, $p \vee\!\!\!\vee \neg p$ (denoted $?p$) stands for the polar question whether p or $\neg p$ is the case, i.e., whether p is true or false.

The propositional fragment of $\mathcal{L}$, consisting of formulas without the binary modality $\Rrightarrow$, is just inquisitive propositional logic [10,14]. Propositional formulas without $\vee\!\!\!\vee$ can be simply identified with standard propositional formulas.

As anticipated in the introduction, we interpret this language over neighborhood models, with the only constraint that neighborhoods be nonempty.

Definition 2.1 An *inhabited neighborhood model* (*in-model* for short) is a tuple $M = \langle W, \Sigma, V\rangle$ where W is a nonempty set, $\Sigma : W \to \wp\wp(W)$ is a map such that $\emptyset \notin \Sigma(w)$ for all $w \in W$, and $V : \mathcal{P} \to \wp(W)$ is a valuation function. We refer to elements $w \in W$ as *worlds* or *points*, to subsets $s \subseteq W$ as *information states* or simply *states*, and to the elements $s \in \Sigma(w)$ as the *neighborhoods* of w.

As mentioned in the introduction, one possible interpretation of these models is as follows. Worlds $w \in W$ are stages of a dynamic process unfolding over time. At each stage w, the relevant agent has at her disposal a set of actions, each of which may produce a range of different outcomes depending on factors such as randomness and the choices of other agents. Each neighborhood $s \in \Sigma(w)$ collects all the outcomes that may result from a given action. The constraint $\emptyset \notin \Sigma(w)$ captures the idea that an action must allow at least one outcome.

To interpret statements like p and questions like $?p$ uniformly, the semantics is not given by a recursive definition of truth at a world, but by a recursive definition of *support* at an information state, as customary in inquisitive logic.

Definition 2.2 [Support] Given an in-model M and a state $s \subseteq W$ we define:

- $M, s \models p \iff s \subseteq V(p)$
- $M, s \models \bot \iff s = \emptyset$
- $M, s \models \varphi \wedge \psi \iff M, s \models \varphi$ and $M, s \models \psi$
- $M, s \models \varphi \vee\!\!\!\vee \psi \iff M, s \models \varphi$ or $M, s \models \psi$
- $M, s \models \varphi \to \psi \iff \forall t \subseteq s : M, t \models \varphi$ implies $M, t \models \psi$
- $M, s \models \varphi \Rrightarrow \psi \iff \forall w \in s\, \forall t \in \Sigma(w) : M, t \models \varphi$ implies $M, t \models \psi$

When no confusion arises, we suppress reference to M and write simply $s \models \varphi$. We assume that the propositional clauses are familiar from previous work in inquisitive logic (see, e.g., [10,12]). The novelty is the presence of the modality

$\Rrightarrow$. Before discussing it in detail, it will be useful to recall some standard facts and notions. First, the following standard properties of inquisitive logics hold:

- Persistency: if $M, s \models \varphi$ and $t \subseteq s$ then $M, t \models \varphi$;
- Empty state property: $M, \emptyset \models \varphi$ for all $\varphi \in \mathcal{L}$.

Logical entailment and equivalence are defined in terms of support, as follows.

Definition 2.3 Let $\Phi \cup \{\psi\} \subseteq \mathcal{L}$. We say that Φ *entails* ψ, denoted $\Phi \models \psi$, if for all in-models M and states s, if $M, s \models \varphi$ for all $\varphi \in \Phi$, then $M, s \models \psi$. We say that ψ is *valid* if $\emptyset \models \psi$. We say that $\varphi, \psi \in \mathcal{L}$ are *logically equivalent*, denoted $\varphi \equiv \psi$, if φ and ψ are supported by the same states in every in-model.

Although the semantics is defined in terms of support, a notion of truth at a world is recovered by defining $M, w \models \varphi$ as $M, \{w\} \models \varphi$. One can then check that all standard connectives have their classical truth-functional behavior, for instance $M, w \models \neg\varphi \iff M, w \not\models \varphi$. In particular, this means that all standard (i.e., $\mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}}$-free) propositional formulas have their usual truth conditions. The set of worlds where φ is true is denoted $|\varphi|_M$, or simply $|\varphi|$ if M is clear.

For many formulas in the language, support at a state s simply boils down to truth at each world in s. We call such formulas *truth-conditional*.

Definition 2.4 A formula $\varphi \in \mathcal{L}$ is truth-conditional if for all in-models M and states $s \subseteq W$ we have $M, s \models \varphi \iff \forall w \in s : M, w \models \varphi$.

We now define a set $\mathcal{L}_!$ of declarative formulas (or *declaratives*) as follows:

$$\mathcal{L}_! \qquad \alpha \ ::= \ p \mid \bot \mid (\alpha \wedge \alpha) \mid (\alpha \to \alpha) \mid (\varphi \Rrightarrow \varphi) \qquad\qquad p \in \mathcal{P}, \varphi \in \mathcal{L}$$

In words, a formula α is a declarative if all occurrences of $\mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}}$ in α are within the scope of the modality $\Rrightarrow$ (either in the 'antecedent' or in the 'consequent'). In the following, we will use $\alpha, \beta, \gamma, \delta$ as meta-variables ranging over $\mathcal{L}_!$, while we use φ, ψ, χ for arbitrary formulas in $\mathcal{L}$.

Up to logical equivalence, $\mathcal{L}_!$ is exactly the truth-conditional fragment of $\mathcal{L}$.

Proposition 2.5 *Every $\alpha \in \mathcal{L}_!$ is truth-conditional. Conversely, if $\varphi \in \mathcal{L}$ is truth-conditional, $\varphi \equiv \alpha$ for some $\alpha \in \mathcal{L}_!$.*

Proof. The first part of the statement is proved by a plain induction on $\alpha \in \mathcal{L}_!$. For the second part, given any $\varphi \in \mathcal{L}$, let φ^c be the declarative obtained by replacing any occurrence of $\mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}}$ which is not in the scope of $\Rrightarrow$ by $\vee$. A simple induction shows that φ and φ^c are true at the same worlds in any model. If φ is truth-conditional, then since φ and φ^c are both truth-conditional with the same truth conditions, they are equivalent. $\square$

Note that all standard (i.e., $\mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}}$-free) propositional formulas are declaratives and thus truth-conditional. For these formulas, our support semantics is equivalent to the standard truth-conditional semantics. For an example of a formula which is *not* truth-conditional, take $?p$. This is supported if the truth value of p is the same is all worlds in s, that is, we have $M, s \models ?p \iff s \subseteq |p|$ or $s \cap |p| = \emptyset$.

With these preliminaries in place, let us now turn to the modality $\Rrightarrow$. This modality can be applied to $\varphi, \psi \in \mathcal{L}$ to produce the modal formula $\varphi \Rrightarrow \psi$. This formula is a declarative and so truth-conditional by Proposition 2.5. Hence, to understand its semantics it suffices to focus on its truth conditions, which are:

$$M, w \models \varphi \Rrightarrow \psi \iff \forall t \in \Sigma(w) : M, t \models \varphi \text{ implies } M, t \models \psi$$

In words, $\varphi \Rrightarrow \psi$ is true at a world w if every neighborhood of w that supports φ also supports ψ. To appreciate what can be expressed by modal formulas of this form, it is helpful to consider a number of special cases.

First, suppose α, β are declaratives. By Proposition 2.5, the formula $\alpha \Rrightarrow \beta$ expresses a kind of global consequence from the perspective of the world w: for any neighborhood s, if α is true everywhere in s, then β is true everywhere in s. More generally, if $\alpha, \beta_1, \ldots, \beta_n$ are declaratives, $\alpha \Rrightarrow (\beta_1 \vee\!\!\!\vee \cdots \vee\!\!\!\vee \beta_n)$ expresses a kind of multiple-conclusion global consequence: for any neighborhood s of w, if α is true everywhere in s then some β_i is true everywhere in s:

$$w \models (\alpha \Rrightarrow \bigvee\!\!\!\bigvee_{i \leq n} \beta_i) \iff \forall s \in \Sigma(w) : \text{if } s \subseteq |\alpha| \text{ then } (s \subseteq |\beta_i| \text{ for some } i)$$

Note that, as a consequence, the negation $\neg(\alpha \Rrightarrow (\neg\beta_1 \vee\!\!\!\vee \cdots \vee\!\!\!\vee \neg\beta_n))$ expresses the existence of a neighborhood s such that α is true everywhere in s and for each $i \leq n$, β_i is true somewhere in s. This is exactly the meaning of a modal formula of the form $\Box(\beta_1, \ldots, \beta_n; \alpha)$ in instantial neighborhood semantics [6]; we will come back to the connection in Section 6. As another example, consider the formula $(?p \Rrightarrow ?q)$. This expresses the fact that any neighborhood that settles whether p also settles whether q, in symbols:

$$w \models (?p \Rrightarrow ?q) \iff \forall s \in \Sigma(w) : \text{if } (s \subseteq |p| \text{ or } s \subseteq \overline{|p|}) \text{ then } (s \subseteq |q| \text{ or } s \subseteq \overline{|q|})$$

Looking at Figure 1, this formula is true at worlds w_1 and w_2, but not at w_3.

An interesting feature of the language is the interplay of the two conditionals $\Rrightarrow$ and $\to$, that can be used to specify global and local restriction respectively: while $\Rrightarrow$ allows us to restrict the class of neighborhoods under consideration, $\to$ allows us to restrict the worlds in each neighborhood. Consider, e.g., the formula $(p \Rrightarrow (q \to ?r))$. This expresses the following fact: if we restrict to those neighborhoods that support p and then we restrict each of these neighborhoods to the q-worlds, all the resulting states settle whether r:

$$w \models (p \Rrightarrow (q \to ?r)) \iff \forall s \in \Sigma(w) : s \subseteq |p| \text{ implies } s \cap |q| \models ?r$$

Similarly, the formula $(?p \Rrightarrow (?q \to ?r))$ expresses that if we restrict to neighborhoods that settle whether p, and then look at the parts of such neighborhoods where the truth value of q is settled, each of these parts settles whether r.

$$w \models (?p \Rrightarrow (?q \to ?r)) \iff \forall s \in \Sigma(w) : (s \subseteq |p| \text{ or } s \subseteq \overline{|p|}) \text{ implies } (s \cap |q| \models ?r \text{ and } s \cap \overline{|q|} \models ?r)$$

Another interesting observation is that from the binary modality $\Rrightarrow$, we can define two unary modalities as follows:

$$\boxplus\varphi := (\top \Rrightarrow \varphi) \qquad \diamondplus\varphi := \neg(\varphi \Rrightarrow \bot)$$

Since all neighborhoods of a point support $\top$ and none support $\bot$ (note: here we use the requirement that neighborhoods are nonempty!) we find that $\boxplus$ and $\diamondplus$ are, respectively, a universal and an existential modality over neighborhoods:

$$w \models \boxplus\varphi \iff \forall s \in \Sigma(w) : s \models \varphi \qquad w \models \diamondplus\varphi \iff \exists s \in \Sigma(w) : s \models \varphi$$

In particular, if $\alpha \in \mathcal{L}_!$ then $\diamondplus\alpha$ is true at a world w if there is some neighborhood s of w such that α is true everywhere in s; as discussed in the introduction, this is the sort of fact expressed by modal formulas in a standard version of neighborhood semantics. When $\alpha \in \mathcal{L}_!$, the formula $\boxplus\alpha$ is not especially interesting: it says that α is true throughout every neighborhood of the given point w, which simply means that α is true everywhere in the set $R_\Sigma(w) = \bigcup\Sigma(w)$. However, the modality $\boxplus$ becomes very interesting when applied to a question. For instance, consider the formula $\boxplus?p$: this is true if the truth value of p is settled, one way or the other, in every neighborhood of the given world. Looking at Figure 1, this formula is true at world w_1, but false at w_2 and w_3. Previous work on inquisitive modal logic [9,10,13,15,17,22,26] focused on a language where $\boxplus$ is the main primitive modality.[3] However, the semantics was based on models that are downward closed in the following sense.[4]

Definition 2.6 Let $M = \langle W, \Sigma, V\rangle$ be an in-model. Its downward-closure is $M^\downarrow = \langle W, \Sigma^\downarrow, V\rangle$, where $\Sigma^\downarrow(w) = \{t \subseteq W \mid t \neq \emptyset \text{ and } t \subseteq s \text{ for some } s \in \Sigma(w)\}$. We say that M is *downward closed* if $M = M^\downarrow$.

Now consider the language $\mathcal{L}_\boxplus$ given by the following BNF definition:

$$\mathcal{L}_\boxplus \qquad \eta ::= p \mid \bot \mid (\eta \wedge \eta) \mid (\eta \to \eta) \mid (\eta \vee\!\!\!\vee\, \eta) \mid \boxplus\eta \qquad\qquad p \in \mathcal{P}$$

In the context of downward-closed models, $\mathcal{L}_\boxplus$ is equi-expressive with $\mathcal{L}$. This is a consequence of the following proposition (the easy proof is left as an exercise).

Proposition 2.7 *If M is a downward-closed in-model then for any $s \subseteq W$, $M, s \models (\varphi \Rrightarrow \psi) \iff M, s \models \boxplus(\varphi \to \psi)$.*

In general, however, $\mathcal{L}_\boxplus$ is strictly less expressive than $\mathcal{L}$. Indeed, formulas of $\mathcal{L}_\boxplus$ cannot distinguish a model from its downward closure (cf. Prop. 2.6 in [22]).

Proposition 2.8 *Let $\eta \in \mathcal{L}_\boxplus$. Then for any M, s: $M, s \models \eta \iff M^\downarrow, s \models \eta$.*

Proof. By induction on η. The key case is the one for $\eta = \boxplus\theta$. Suppose $M, s \models \boxplus\theta$. Take any $w \in s$ and any state $t \in \Sigma^\downarrow(w)$. Then $t \subseteq t'$ for some $t' \in \Sigma(w)$. Since $M, s \models \boxplus\theta$ and $w \in s$ we have $M, t' \models \theta$ and so by persistency $M, t \models \theta$, which by induction hypothesis gives $M^\downarrow, t \models \theta$. This shows that $M^\downarrow, s \models \boxplus\theta$. The converse implication is immediate since $\Sigma(w) \subseteq \Sigma^\downarrow(w)$. $\square$

[3] Normally, a second primitive modality $\Box$ is considered as well, but at least in propositional modal logic, it can be removed from the language without loss of expressive power (see [10]).

[4] A subtlety: in previous work $\Sigma(w)$ is required to always contain $\emptyset$, whereas here we require $\Sigma(w)$ to never contain $\emptyset$. However, the presence or absence of $\emptyset$ in $\Sigma(w)$ is immaterial to the semantics of $\mathcal{L}$ and $\mathcal{L}_\boxplus$ by the empty state property. Thus, the models considered in previous work can be safely identified with their counterparts where $\emptyset$ is removed from all $\Sigma(w)$.

It is now easy to see that many $\varphi \in \mathcal{L}$ are not equivalent to any $\eta \in \mathcal{L}_{\boxplus}$. For instance, consider a model M with $W = \{v, v'\}$, $\Sigma(v) = \{W\}$ and $V(p) = \{v\}$. Then $M, v \not\models \oplus p$ but $M^{\downarrow}, v \models \oplus p$. Thus, $\oplus p$ is not equivalent to any $\eta \in \mathcal{L}_{\boxplus}$.[5]

3 Bisimulation and expressive power

In a neighborhood model, it is natural to regard two worlds as equivalent if they agree on atomic sentences and every neighborhood of the one is equivalent to some neighborhood of the other. In turn, it is natural to regard two neighborhoods as equivalent if every world in the one is equivalent to some world in the other. This leads naturally to the following notion of bisimulation.

Definition 3.1 [(cf. [6,13])] Given models $M = \langle W, \Sigma, V\rangle, M' = \langle W', \Sigma', V'\rangle$, a relation $Z \subseteq W \times W'$ is a *bisimulation* if whenever wZw' holds we have:

Atoms: $w \in V(p) \iff w' \in V'(p)$ for all $p \in \mathcal{P}$;

Forth: $\forall s \in \Sigma(w)\, \exists s' \in \Sigma'(w')$ with $s\overline{Z}s'$;

Back: $\forall s' \in \Sigma'(w')\, \exists s \in \Sigma(w)$ with $s\overline{Z}s'$.

Here, $\overline{Z}$ is the Egli-Milner lifting of Z, that is, $s\overline{Z}s'$ holds just in case we have:

$$(\forall w \in s\, \exists w' \in s' : wZw') \text{ and } (\forall w' \in s'\, \exists w \in s : wZw')$$

Two worlds w, w' are *bisimilar* ($w \sim w'$) if there is a bisimulation Z with wZw'; two states s, s' are bisimilar ($s \sim s'$) if there is a bisimulazion Z with $s\overline{Z}s'$.

Bisimilarity can be characterized in terms of a game with two players, S(poiler) and D(uplicator). The game alternates between world-positions $\langle w, w'\rangle \in W \times W'$ and state-positions $\langle s, s'\rangle \in \wp(W) \times \wp(W')$. Playing from a world position, S picks a neighborhood of either world and D responds with a neighborhood of the other world, leading to a state position. Playing from a state position, S picks a world in either state and D responds with a world in the other state. If D is unable to make a move, or if a world-position is reached where the worlds disagree on some atomic sentence, S wins; in all other cases, D wins. Then worlds w, w' (respectively, states s, s') are bisimilar iff D has a winning strategy in the game starting from position $\langle w, w'\rangle$ (respectively, $\langle s, s'\rangle$).

This notion of bisimulation is especially useful since it allows for standard model-theoretic constructions like disjoint unions and (partial) tree unfoldings (see [6] for the details), while preserving the satisfaction of InqCM formulas.

Definition 3.2 [Modal equivalence] Given models M, M', we say that two states $s \subseteq W, s' \subseteq W'$ are *modally equivalent* ($s \leftrightsquigarrow s'$) if they support the same formulas of InqCM. Similarly, two worlds $w \in W$ and $w' \in W'$ are modally equivalent ($w \leftrightsquigarrow w'$) if they make the same formulas true.

[5] There is also a tight relation between our operator $\Rrightarrow$ and the inquisitive strict conditional $\Rightarrow$ studied in [11]. With every Kripke model $M = \langle W, R, V\rangle$ we can associate a neighborhood model $M^n = \langle W, \Sigma_R, V\rangle$ where $\Sigma_R(w) = \{s \subseteq R[w] \mid s \neq \emptyset\}$. Then the semantics of a formula $\varphi \Rightarrow \psi$ in the setting of M coincides with the semantic of $\varphi \Rrightarrow \psi$ in M^n, and the logic $\mathsf{inq}^{\Rightarrow}$ of [11] can be viewed as the special case of InqCM over a restricted class of models.

Proposition 3.3 (Invariance under bisimulation) *For any worlds w, w', $w \sim w'$ implies $w \leftrightsquigarrow w'$; for any states s, s', $s \sim s'$ implies $s \leftrightsquigarrow s'$.*

The proof is straightforward; we include it in Appendix A for completeness.

Conversely, we can show a Hennessy-Milner theorem. For an in-model M, we say M is *image-finite* if the set $\bigcup \Sigma(w)$ is finite for every $w \in W$. In such models, modal equivalence implies bisimilarity for worlds and for finite states.

Theorem 3.4 *If M, M' are image-finite, then for any worlds w, w', $w \leftrightsquigarrow w'$ implies $w \sim w'$, and for all finite states s, s', $s \leftrightsquigarrow s'$ implies $s \sim s'$.*

Proof. We will show that the modal equivalence relation $\leftrightsquigarrow$ on worlds is a bisimulation. Suppose $w \leftrightsquigarrow w'$. Clearly, w and w' satisfy the same atoms. We show that the Forth condition is satisfied (the proof for Back is analogous). Take a state $s \in \Sigma(w)$. By image-finiteness we can write $s = \{w_1, \ldots, w_n\}$ and $\bigcup \Sigma'(w') = \{v_1, \ldots, v_m\}$. Note that $n \geq 1$ since neighborhoods are required to be nonempty, and $m \geq 1$ since if $m = 0$ we would have $\Sigma'(w') = \emptyset$, in which case w and w' are distinguished by the formula $\top \Rrightarrow \bot$. Now for each i, j, if $w_i \leftrightsquigarrow v_j$ we define $\delta_{ij} = \top$, while if $w_i \not\leftrightsquigarrow v_j$ we let δ_{ij} be a declarative such that $M, w_i \models \delta_{ij}$ and $M, v_j \models \neg\delta_{ij}$. Note that such a declarative exists: if $w_i \not\leftrightsquigarrow v_j$, then w_i and v_j disagree about the truth of some formula $\xi \in \mathcal{L}$. Now ξ has the same truth conditions as the declarative ξ^c defined as in the proof of Proposition 2.5, so w_i and v_i disagree on the truth of ξ^c. We can choose $\delta = \xi^c$ if $w_i \models \xi^c$ and $\delta = \neg\xi^c$ otherwise.

Let $\gamma_i := \bigwedge_{j \leq m} \delta_{ij}$. For all $j \leq m$ we have $M', v_j \models \gamma_i \iff w_i \leftrightsquigarrow v_j$. Now consider the formula:

$$\varphi := (\bigvee_{i=1}^{n} \gamma_i) \Rrightarrow (\mathbb{V}_{i=1}^{n} \neg\gamma_i)$$

It is easy to check that s supports the antecedent of φ but not the consequent, so $M, w \not\models \varphi$. Since $w \leftrightsquigarrow w'$, we have that $M', w' \not\models \varphi$. So there is a $s' \in \Sigma'(w')$ that supports the antecedent of φ but not the consequent. We are now going to show that every world in s is modally equivalent to a world in s' and vice versa, thus showing that the Forth condition holds for $\leftrightsquigarrow$. Since $s' \models \bigvee_{i \leq n} \gamma_i$, by persistency every world $v \in s'$ satisfies some formula γ_i and thus is modally equivalent to some world in s, namely w_i. For the converse, take a world $w_i \in s$. Since s' does not support the consequent of φ we have $s' \not\models \neg\gamma_i$. Since $\neg\gamma_i$ is a declarative and thus truth-conditional, there is a world $v \in s'$ such that $v \not\models \neg\gamma_i$ and so $v \models \gamma_i$, which means that v is modally equivalent to w_i.

This proves the claim for worlds. For the claim about states, let s, s' be finite states with $s \leftrightsquigarrow s'$. Since $\leftrightsquigarrow$ on worlds is a bisimulation, it suffices to show that every world in s is modally equivalent to a world in s' and vice versa. So, take $w \in s$. Proceeding as above, we can define a declarative γ_w such that $w \models \gamma_w$ and for all $v \in s'$ we have $v \models \gamma_w \iff w \leftrightsquigarrow v$. Now since $w \in s$, by persistency we have $s \not\models \neg\gamma_w$ and since $s \leftrightsquigarrow s'$ also $s' \not\models \neg\gamma_w$. Since $\neg\gamma_w$ is a declarative and thus truth-conditional, there is some $w' \in s'$ with $w' \not\models \neg\gamma_w$ and so $w' \models \gamma_w$, which implies $w \leftrightsquigarrow w'$. The converse is proved analogously. □

It is not hard to construct examples showing that the result fails if M, M' are not image-finite, or if s, s' are infinite states (even within image-finite models).

We can also introduce notions of n-step bisimilarity and show that, if $\mathcal{P}$ is finite, two worlds/states are n-step bisimilar iff they cannot be distinguished by a formula of modal depth up to n. We omit the details due to space constraints.

4 Axiomatization

In this section we describe a Hilbert-style proof system for InqCM. The propositional basis for the system consists of all instances of the axioms for intuitionistic propositional logic, with $\vee\!\!\!\vee$ in the role of intuitionistic disjunction, and all instances of the following two axioms, where $\alpha \in \mathcal{L}_!$ and $\varphi, \psi \in \mathcal{L}$:

(DDN) $\neg\neg\alpha \to \alpha$
(Split) $(\alpha \to \varphi \vee\!\!\!\vee \psi) \to (\alpha \to \varphi) \vee\!\!\!\vee (\alpha \to \psi)$

An explicit list of the propositional axiom schemata is included in Appendix B. We then have three axiom schemata for $\Rrightarrow$, namely:

(Tran) $(\varphi \Rrightarrow \psi) \wedge (\psi \Rrightarrow \chi) \to (\varphi \Rrightarrow \chi)$
(Conj) $(\varphi \Rrightarrow \psi) \wedge (\varphi \Rrightarrow \chi) \to (\varphi \Rrightarrow (\psi \wedge \chi))$
(Disj) $(\varphi \Rrightarrow \chi) \wedge (\psi \Rrightarrow \chi) \to ((\varphi \vee\!\!\!\vee \psi) \Rrightarrow \chi)$

The inference rules are modus ponens and a conditional version of necessitation:

$$\frac{\varphi \quad \varphi \to \psi}{\psi}\ (\texttt{MP}) \qquad\qquad \frac{\varphi \to \psi}{\varphi \Rrightarrow \psi}\ (\texttt{CN})$$

If $\Phi, \Psi \subseteq \mathcal{L}$, we write $\Phi \vdash \Psi$ if there are $\varphi_1, \dots, \varphi_n \in \Phi$ and $\psi_1, \dots, \psi_m \in \Psi$ such that the formula $\varphi_1 \wedge \dots \wedge \varphi_n \to \psi_1 \vee\!\!\!\vee \dots \vee\!\!\!\vee \psi_m$ is derivable in the system (we allow for $n = 0$, in which case the relevant antecedent is $\top$, and for $m = 0$, in which case the consequent is $\bot$). We write $\varphi_1, \dots, \varphi_n \vdash \psi_1, \dots, \psi_m$ instead of $\{\varphi_1, \dots, \varphi_n\} \vdash \{\psi_1, \dots, \psi_m\}$, and we write $\varphi \dashv\vdash \psi$ in case $\varphi \vdash \psi$ and $\psi \vdash \varphi$. As usual, soundness is proved by checking that the axioms are valid and the inference rules preserve validity. We omit the straightforward proof.

Proposition 4.1 (Soundness) *For all $\Phi \cup \{\psi\} \subseteq \mathcal{L}$, $\Phi \vdash \psi$ implies $\Phi \models \psi$.*

The completeness of the system will be shown in Section 7. For that, we need to introduce one more notion which is standard in inquisitive logic: resolutions.

5 Resolutions

Following a standard recipe, we will associate to each $\varphi \in \mathcal{L}$ a finite nonempty set of declaratives $\mathcal{R}(\varphi) \subseteq \mathcal{L}_!$ such that φ is equivalent to $\bigvee\!\!\!\bigvee \mathcal{R}(\varphi)$.

Definition 5.1 [Resolutions (cf. Definition 7.1.24 in [10])]

- $\mathcal{R}(\alpha) = \{\alpha\}$ if α is an atom, $\bot$, or a modal formula $(\varphi \Rrightarrow \psi)$
- $\mathcal{R}(\varphi \wedge \psi) = \{\alpha \wedge \beta \mid \alpha \in \mathcal{R}(\varphi), \beta \in \mathcal{R}(\psi)\}$
- $\mathcal{R}(\varphi \vee\!\!\!\vee \psi) = \mathcal{R}(\varphi) \cup \mathcal{R}(\psi)$

- $\mathcal{R}(\varphi \to \psi) = \{\bigwedge_{\alpha \in \mathcal{R}(\varphi)}(\alpha \to f(\alpha)) \mid f : \mathcal{R}(\varphi) \to \mathcal{R}(\psi)\}$

It is immediate to verify that for every declarative $\alpha \in \mathcal{L}_!$ we have $\mathcal{R}(\alpha) = \{\alpha\}$.

The following provable normal form result uses only the propositional component of the proof system, and can thus be proved in exactly the same way as for inquisitive propositional logic (see Lemma 3.3.4 in [10] for the details).

Lemma 5.2 *For all* $\varphi \in \mathcal{L}$, $\varphi \dashv\vdash \bigvee\!\!\!\!\bigvee \mathcal{R}(\varphi)$.

As an immediate corollary we have the following fact.

Lemma 5.3 *For all* $\varphi \in \mathcal{L}$ *and all* $\alpha \in \mathcal{R}(\varphi)$, $\alpha \vdash \varphi$.

The next lemma says that if a set of declaratives derives a formula, it derives a particular resolution of it (note that the converse holds by the previous lemma). The proof is again standard (see, e.g., Lemma 5.6 in [11]), but we include it in Appendix C for the sake of completeness.

Lemma 5.4 *Let* $\Gamma \subseteq \mathcal{L}_!$ *and* $\varphi \in \mathcal{L}$. *If* $\Gamma \vdash \varphi$, *then* $\Gamma \vdash \alpha$ *for some* $\alpha \in \mathcal{R}(\varphi)$.

We will also need a notion of resolutions for sets of formulas. A resolution of a set Φ is a set obtained by replacing each element $\varphi \in \Phi$ by a resolution of it.

Definition 5.5 Given $\Phi \subseteq \mathcal{L}$, a resolution function for Φ is a function f that associates to each $\varphi \in \Phi$ some $f(\varphi) \in \mathcal{R}(\varphi)$. A resolution of Φ is the image of Φ under a resolution function: $\mathcal{R}(\Phi) = \{f[\Phi] \mid f$ a resolution function for $\Phi\}$.

Note that if Γ is a set of declaratives then $\mathcal{R}(\Gamma) = \{\Gamma\}$.

The next lemma says that if Φ fails to derive Ψ, then Φ can be strengthened to a resolution $\Gamma \in \mathcal{R}(\Phi)$ that still fails to derive Ψ. The proof is again standard. Since it uses only Lemma 5.2 and the propositional axioms for $\bigvee\!\!\!\!\bigvee$, we omit it and refer, e.g., to pp. 87-88 in [10], where the argument is given in detail.

Lemma 5.6 *For all* $\Phi, \Psi \subseteq \mathcal{L}$, *if* $\Phi \nvdash \Psi$ *then there is* $\Delta \in \mathcal{R}(\Phi)$ *s.t.* $\Delta \nvdash \Psi$.

With these standard results at hand, we are now ready to prove the completeness of our system. Before turning to that, however, we will make use of resolutions to relate InqCM more precisely to instantial neighborhood logic.

6 Translating to and from instantial neighborhood logic

As discussed in the introduction, instantial neighborhood logic (INL, [6]) is a modal language interpreted on neighborhood structures which is invariant under the notion of bisimulation that we discussed in Section 3. The language $\mathcal{L}_{\mathsf{INL}}$ of INL is a modal language with primitive connectives $\neg$ and $\wedge$, and where modal formulas have the form $\Box(\rho_1, \ldots, \rho_n; \sigma)$ with $n \geq 0$. The semantics is given by a standard definition of truth at a world, where the modal clause is:

$$M, w \models \Box(\rho_1, \ldots, \rho_n; \sigma) \iff \exists s \in \Sigma(w) : (\forall v \in s : M, v \models \sigma) \text{ and } (\forall i \leq n \exists v \in s : M, v \models \rho_i)$$

We will show that over inhabited neighborhood models, INL has the same expressive power as the declarative fragment of InqCM. We prove this by defining two translations that preserve truth conditions. We define a translation

$(\cdot)^* : \mathcal{L}_{\mathsf{INL}} \to \mathcal{L}_!$ by letting $p^* = p$, $(\neg\sigma)^* = \neg\sigma^*$, $(\rho \wedge \sigma)^* = \rho^* \wedge \sigma^*$ and, crucially,

$$\Box(\rho_1, \dots, \rho_n; \sigma)^* = \neg(\sigma^* \Rrightarrow (\neg\rho_1^* \mathbin{\vee\!\!\!\vee} \cdots \mathbin{\vee\!\!\!\vee} \neg\rho_n^*))$$

As a straightforward induction shows, this map preserves truth conditions.

Proposition 6.1 *Let $\sigma \in \mathcal{L}_{\mathsf{INL}}$. For every in-model M and every world w, $M, w \models \sigma \iff M, w \models \sigma^*$.*

Translating declaratives of InqCM to INL is more tricky. Take a modal formula $(\varphi \Rrightarrow \psi)$: in general, φ and ψ are not declaratives, so the translation will not be defined on them. Instead, we first compute the resolutions of φ and ψ, and then assemble a translation from the translations of these. To make this precise, we need a non-standard notion of complexity. Given $\alpha, \beta \in \mathcal{L}_!$, we let $\alpha \prec \beta$ in case either α has lower modal depth than β, or α and β have the same modal depth and α is a subformula of β. Clearly, $\prec$ is well-founded and thus suitable for induction. Now we define $(\cdot)^\star : \mathcal{L}_! \to \mathcal{L}_{\mathsf{INL}}$ recursively on $\prec$ as follows: $p^\star = p$; $\bot^\star = (p \wedge \neg p)$ for an arbitrary $p \in \mathcal{P}$; $(\alpha \wedge \beta)^\star = \alpha^\star \wedge \beta^\star$; $(\alpha \to \beta)^\star = \neg(\alpha^\star \wedge \neg\beta^\star)$; and finally:

$$(\varphi \Rrightarrow \psi)^\star = \bigwedge_{i=1}^{n} \neg\Box(\neg\beta_1^\star, \dots, \neg\beta_m^\star; \alpha_i^\star)$$

where $\{\alpha_1, \dots, \alpha_n\} = \mathcal{R}(\varphi)$ and $\{\beta_1, \dots, \beta_m\} = \mathcal{R}(\psi)$. Note that $\alpha_i^\star$ is defined since $\alpha_i \prec (\varphi \Rrightarrow \psi)$: this is because α_i has the same modal depth as φ, which is lower that the modal depth of $(\varphi \Rrightarrow \psi)$. A similar argument goes for β_j.

Proposition 6.2 *Let $\alpha \in \mathcal{L}_!$ be any declarative in InqCM. For any in-model M and world w we have $M, w \models \alpha \iff M, w \models \alpha^\star$.*

Proof. By induction on $\prec$. We focus on the induction step for $\alpha = (\varphi \Rrightarrow \psi)$. Let $\mathcal{R}(\varphi) = \{\alpha_1, \dots, \alpha_n\}$ and $\mathcal{R}(\psi) = \{\beta_1, \dots, \beta_m\}$. Take any in-model M and any world w in M. We will show that $M, w \not\models (\varphi \Rrightarrow \psi)$ iff $M, w \not\models (\varphi \Rrightarrow \psi)^\star$. The second step uses Lemma 5.2, while the fourth step uses Proposition 2.5.

$$\begin{aligned}
& w \not\models (\varphi \Rrightarrow \psi) \\
\iff & \exists s \in \Sigma(w) : s \models \varphi \text{ and } s \not\models \psi \\
\iff & \exists s \in \Sigma(w) : s \models \textstyle\mathbin{\vee\!\!\!\vee}_{i \le n} \alpha_i \text{ and } s \not\models \textstyle\mathbin{\vee\!\!\!\vee}_{j \le m} \beta_j \\
\iff & \exists s \in \Sigma(w) : \exists i \le n (s \models \alpha_i) \text{ and } \forall j \le m (s \not\models \beta_j) \\
\iff & \exists s \in \Sigma(w) : \exists i \le n (\forall v \in s : v \models \alpha_i) \text{ and } \forall j \le m (\exists v \in s : v \not\models \beta_j) \\
\iff & \exists i \le n \exists s \in \Sigma(w) : (\forall v \in s : v \models \alpha_i) \text{ and } \forall j \le m (\exists v \in s : v \models \neg\beta_j) \\
\iff & \exists i \le n \exists s \in \Sigma(w) : (\forall v \in s : v \models \alpha_i^\star) \text{ and } \forall j \le m (\exists v \in s : v \models \neg\beta_j^\star) \\
\iff & \exists i \le n : w \models \Box(\neg\beta_1^\star, \dots, \neg\beta_m^\star; \alpha_i^\star) \\
\iff & w \not\models \textstyle\bigwedge_{i=1}^{n} \neg\Box(\neg\beta_1^\star, \dots, \neg\beta_m^\star; \alpha_i^\star) \\
\iff & w \not\models (\varphi \Rrightarrow \psi)^\star
\end{aligned}$$

Again, note that we can use the induction hypothesis on α_i since $\alpha_i \prec (\varphi \Rrightarrow \psi)$: indeed, α_i has the same modal depth as φ, which is lower than the one of $(\varphi \Rrightarrow \psi)$. A similar argument goes for β_j. □

A couple of remarks on this translation. First, note that given a formula $\sigma \in \mathcal{L}_{\mathsf{INL}}$, the size of σ^* grows linearly on the size of σ. By contrast, since the number of resolutions of a formula $\varphi \in \mathcal{L}$ grows exponentially in the length of φ due to the clause for implication, the size of the translation of a formula $\alpha \in \mathcal{L}_!$ may grow exponentially in the size of α. It seems natural to conjecture that this is inevitable for such a translation, and thus that InqCM is exponentially more succinct than INL, but we will not try to prove this here.

It is also worth noting that this strategy, that relies crucially on resolutions, would not be viable in the setting of inquisitive predicate logic, where no analogue of resolutions is available. It is natural to conjecture that a first-order version of InqCM would be strictly more expressive than a first-order version of INL (for readers familiar with inquisitive logic, a challenge would be to translate formulas like $\boxplus(\forall x?Px \to \forall x?Qx)$, which say that in every neighborhood s of w the extension of Q is functionally determined by the extension of P).

7 Completeness

We will now prove that our axiomatization is complete by constructing a canonical model. Call a set $\Gamma \subseteq \mathcal{L}_!$ of declaratives a *complete theory of declaratives* (CTD for short) if (i) Γ is deductively closed w.r.t. declaratives: if $\Gamma \vdash \alpha$ and $\alpha \in \mathcal{L}_!$ then $\alpha \in \Gamma$; (ii) Γ is complete: for all $\alpha \in \mathcal{L}_!$, exactly one of α and $\neg\alpha$ is in Γ (thus, $\bot \notin \Gamma$, since $\top \in \Gamma$ by (i)). Now if S is a set of CTDs, we let $\bigcap S = \{\alpha \in \mathcal{L}_! \mid \alpha \in \Gamma \text{ for all } \Gamma \in S\}$ (thus, in particular, $\bigcap \emptyset = \mathcal{L}_!$).

Remark 7.1 For any set S of CTDs, $\bigcap S \vdash \alpha$ and $\alpha \in \mathcal{L}_!$ implies $\alpha \in \bigcap S$.

With any set $\Delta \subseteq \mathcal{L}_!$ of declaratives we can associate a set of CTDs, namely, the set of its complete extensions: $S_\Delta = \{\Gamma \mid \Gamma \text{ is a CTD and } \Delta \subseteq \Gamma\}$. The following Lindenbaum-type lemma is proved by the usual saturation argument.

Lemma 7.2 *If $\Delta \subseteq \mathcal{L}_!$ and $\Delta \nvdash \bot$, then $S_\Delta \neq \emptyset$.*

For any $\Delta \subseteq \mathcal{L}_!$, the sets Δ and $\bigcap S_\Delta$ prove the same formulas.

Lemma 7.3 *For any $\Delta \subseteq \mathcal{L}_!$ and $\varphi \in \mathcal{L}$: $\Delta \vdash \varphi \iff \bigcap S_\Delta \vdash \varphi$.*

Proof. The direction $\Rightarrow$ is clear as $\Delta \subseteq \bigcap S_\Delta$. For the converse, suppose for a contradiction that for some φ we had $\bigcap S_\Delta \vdash \varphi$ but $\Delta \nvdash \varphi$. Since $\bigcap S_\Delta$ is a set of declaratives, by Lemma 5.4 we have $\bigcap S_\Delta \vdash \alpha$ for some $\alpha \in \mathcal{R}(\varphi)$. Since $\Delta \nvdash \varphi$, it follows by Lemma 5.3 that $\Delta \nvdash \alpha$. By the axiom $\neg\neg\alpha \to \alpha$, this implies $\Delta \nvdash \neg\neg\alpha$, and therefore $\Delta, \neg\alpha \nvdash \bot$. So by Lemma 7.2 there is a CTD Γ' such that $\Delta \cup \{\neg\alpha\} \subseteq \Gamma'$. Since $\Gamma' \in S_\Delta$ and $\alpha \notin \Gamma'$ we have $\alpha \notin \bigcap S_\Delta$, so by Remark 7.1 we have $\bigcap S_\Delta \nvdash \alpha$, which is a contradiction. □

We now define a canonical model based on complete theories of declaratives.

Definition 7.4 The canonical model for InqCM is $M^c = \langle W^c, \Sigma^c, V^c \rangle$ where:

- W^c is the set of complete theories of declaratives;
- $\Sigma^c(\Gamma) = \{S \neq \emptyset \mid \forall \varphi, \psi \in \mathcal{L} : (\varphi \Rrightarrow \psi) \in \Gamma \text{ and } \bigcap S \vdash \varphi \text{ implies } \bigcap S \vdash \psi\}$
- $V^c(p) = \{\Gamma \in W^c \mid p \in \Gamma\}$

The following analogue of the standard existence lemma is the key to the completeness proof.

Lemma 7.5 (Existence Lemma) *If* Γ *is a CTD and* $(\varphi \Rrightarrow \psi) \notin \Gamma$*, there exists a state* $S \in \Sigma^c(\Gamma)$ *such that* $\bigcap S \vdash \varphi$ *and* $\bigcap S \nvdash \psi$*.*

Proof. Here we provide the outline of the argument, referring to Appendix D for some technical details. Given two sets $\Phi, \Psi \subseteq \mathcal{L}$, we write $\Phi \Rrightarrow_\Gamma \Psi$ if there are $\varphi_1, \ldots, \varphi_n \in \Phi$ and $\psi_1, \ldots, \psi_m \in \Psi$ with $((\bigwedge_{i \leq n} \varphi_i) \Rrightarrow (\bigvee\!\!\!\!\bigvee_{i \leq m} \psi_i)) \in \Gamma$.

Step 1. We show that $\Phi \cup \{\chi\} \Rrightarrow_\Gamma \Psi$ and $\Phi \Rrightarrow_\Gamma \Psi \cup \{\chi\}$ implies $\Phi \Rrightarrow_\Gamma \Psi$.

This is where the modal part of the proof system is used in a crucial way. We refer to to Appendix D for the details of this step.

Step 2. We partition $\mathcal{L}$ into two sets L, R with $\varphi \in \mathsf{L}$, $\psi \in \mathsf{R}$, and $\mathsf{L} \not\Rrightarrow_\Gamma \mathsf{R}$.

For this, we use Step 1 within a saturation procedure familiar from intuitionistic and inquisitive logic (cf. §3.3 of [16] and [20]). See Appendix D for the details.

Step 3. We construct the required state S.

Since $\mathsf{L} \not\Rrightarrow_\Gamma \mathsf{R}$, the rule (CN) guarantees that $\mathsf{L} \nvdash \mathsf{R}$. By Lemma 5.6 we can find a set $\Delta \in \mathcal{R}(\mathsf{L})$ with $\Delta \nvdash \mathsf{R}$. We can now take $S = S_\Delta = \{\Gamma' \in W^c \mid \Delta \subseteq \Gamma'\}$. We need to verify that (i) $\bigcap S \vdash \varphi$, (ii) $\bigcap S \nvdash \psi$ and (iii) $S \in \Sigma^c(\Gamma)$.

- For (i), recall $\varphi \in \mathsf{L}$. Since $\Delta \in \mathcal{R}(\mathsf{L})$, for some $\alpha \in \mathcal{R}(\varphi)$ we have $\alpha \in \Delta$. By Lemma 5.3, $\Delta \vdash \varphi$, and thus by Lemma 7.3 $\bigcap S \vdash \varphi$.
- For (ii), recall $\psi \in \mathsf{R}$. Since $\Delta \nvdash \mathsf{R}$ we have $\Delta \nvdash \psi$. By Lemma 7.3, $\bigcap S \nvdash \psi$.
- For (iii), first note that since $\Delta \nvdash \mathsf{R}$, we have $\Delta \nvdash \bot$, so by Lemma 7.2 $S \neq \emptyset$. Next, suppose $(\chi \Rrightarrow \xi) \in \Gamma$ and $\bigcap S \vdash \chi$. By Lemma 7.3, $\Delta \vdash \chi$. Since by construction $\Delta \nvdash \mathsf{R}$, it follows that $\chi \notin \mathsf{R}$, so $\chi \in \mathsf{L}$. Now we must have $\xi \in \mathsf{L}$ as well, for if we had $\xi \in \mathsf{R}$ it would follow from $(\chi \Rrightarrow \xi) \in \Gamma$ that $\mathsf{L} \Rrightarrow_\Gamma \mathsf{R}$, contrary to what we saw. Since $\xi \in \mathsf{L}$ and $\Delta \in \mathcal{R}(\mathsf{L})$, for some $\alpha \in \mathcal{R}(\xi)$ we have $\alpha \in \Delta$, so by Lemma 5.3 $\Delta \vdash \xi$. Finally, Lemma 7.3 gives $\bigcap S \vdash \xi$.

□

The bridge between derivability in our proof system and semantics in M^c is given by the following support lemma, which generalizes the usual truth lemma.

Lemma 7.6 *For all states* $S \subseteq W^c$ *and all* $\varphi \in \mathcal{L}$*:* $M^c, S \models \varphi \iff \bigcap S \vdash \varphi$*.*

Proof. By induction on φ. The cases for atoms and connectives are standard (cf. pp. 90-91 in [10]). We spell out the inductive step for $\varphi = (\psi \Rrightarrow \chi)$.

Suppose $\bigcap S \vdash (\psi \Rrightarrow \chi)$. Take a world $\Gamma \in S$ and a state $T \in \Sigma^c(\Gamma)$ with $M^c, T \models \psi$. By induction hypothesis we have $\bigcap T \vdash \psi$. Since $\Gamma \in S$ we have $\bigcap S \subseteq \Gamma$, so $\Gamma \vdash (\psi \Rrightarrow \chi)$. Since $(\psi \Rrightarrow \chi) \in \mathcal{L}_!$, it follows that $(\psi \Rrightarrow \chi) \in \Gamma$. By definition of Σ^c, from $(\psi \Rrightarrow \chi) \in \Gamma$ and $\bigcap T \vdash \psi$ we can conclude $\bigcap T \vdash \chi$, which by induction hypothesis gives $M^c, T \models \chi$. Hence, $M^c, S \models (\psi \Rrightarrow \chi)$.

For the converse, suppose $\bigcap S \nvdash (\psi \Rrightarrow \chi)$. Then there is some $\Gamma \in S$ such that $(\psi \Rrightarrow \chi) \notin \Gamma$. By the Existence Lemma (Lemma 7.5) there is a state

$T \in \Sigma^c(\Gamma)$ such that $\bigcap T \vdash \psi$ and $\bigcap T \nvdash \chi$, which by induction hypothesis means that $M^c, T \models \psi$ and $M^c, T \not\models \chi$. Hence, $M^c, S \not\models (\psi \Rrightarrow \chi)$. □

Finally, we use this lemma to establish the strong completeness of our system.

Theorem 7.7 (Completeness) *For all* $\Phi \cup \{\psi\} \subseteq \mathcal{L}$, $\Phi \models \psi$ *implies* $\Phi \vdash \psi$.

Proof. Suppose $\Phi \nvdash \psi$. By Lemma 5.6 we can find a $\Delta \in \mathcal{R}(\Phi)$ with $\Delta \nvdash \psi$. Note that since $\Delta \in \mathcal{R}(\Phi)$, for all $\varphi \in \Phi$ there is some $\alpha \in \mathcal{R}(\varphi)$ with $\alpha \in \Delta$, which by Lemma 5.3 gives $\Delta \vdash \varphi$. Now take $S_\Delta = \{\Gamma' \in W^c \mid \Delta \subseteq \Gamma'\}$. By Lemma 7.3, $\bigcap S_\Delta \vdash \varphi$ for all $\varphi \in \Phi$, while $\bigcap S_\Delta \nvdash \psi$. By the support lemma, in the model M^c the state S_Δ supports all formulas in Φ but not ψ, so $\Phi \not\models \psi$. □

8 Further work

We close by outlining a number of directions for future work. First, it would be natural to relate the expressive power of InqCM to that of first-order logic over neighborhood models, viewed as two-sorted structures. One could define a standard translation and aim for a van Benthem-style characterization of InqCM as the bisimulation-invariant fragment of first-order logic. For inquisitive modal logic over downward closed models, such a result was proved in [13].

Second, it would be interesting to develop modal correspondence theory for InqCM, relating the validity of InqCM-schemata over a neighborhood frame $\langle W, \Sigma \rangle$ to corresponding frame properties. For instance, given an appropriate notion of validity, one can show that the schema $(\varphi \Rrightarrow \psi) \to \boxplus(\varphi \to \psi)$ is valid on a finite neighborhood frame $\langle W, \Sigma \rangle$ if and only if Σ is downward closed, in the sense that $\Sigma = \Sigma^{\downarrow}$ (cf. Definition 2.6). Relatedly, it would be interesting to investigate the logic of frame classes arising from natural constraints on Σ.

One should also explore what happens when we allow empty neighborhoods. Presumably, this makes little difference, provided the clause for $\Rrightarrow$ is explicitly re-cast as "for all *non-empty* $s \in \Sigma(w)$…". Our language $\mathcal{L}$ would, of course, be unable to tell whether or not $\emptyset \in \Sigma(w)$, but this could be fixed by adding a modal atom empty, which is truth-conditional with truth conditions given by $w \models \mathsf{empty} \iff \emptyset \in \Sigma(w)$. We expect that all the results in this paper can be replicated straightforwardly for this extended language.

Finally, it would be interesting to look at concrete interpretations of InqCM. One salient interpretation that we mentioned is in terms of action and ability. This interpretation naturally calls for an extension to the multi-agent setting; in fact, it would be natural to consider not only the abilities of single agents, but also those of groups, leading to an inquisitive extension of coalition logic [25] capable of expressing facts such as "by acting, the coalition of a and b is bound to settle whether p one way or the other" or "the coalition is bound to force p if they want to force q" (see [5,18] for parallel multi-agent extensions of instantial neighborhood logic). Such a logic could, in turn, be further extended with the resources of temporal logic to talk about long-term strategic abilities, leading for instance to an inquisitive extension of ATL [3,19].

Appendix

A Proof of Proposition 3.3 (bisimulation invariance).

In this section we give the proof of Proposition 3.3. It suffices to prove the claim for information states, since the claim about worlds is obtained as a special case (note that $(w \sim w') \iff (\{w\} \sim \{w'\})$ and $(w \leftrightsquigarrow w') \iff (\{w\} \leftrightsquigarrow \{w'\})$).

So, take two in-models M, M'. We will prove that for every $\varphi \in \mathcal{L}$ we have:

$$\forall s \subseteq W\, \forall s' \subseteq W' : (s \sim s') \text{ implies } (M, s \models \varphi \iff M', s' \models \varphi)$$

We show this by induction on φ. The inductive cases for $\wedge$ and $\mathbin{\rotatebox[origin=c]{-90}{$\geqslant$}}$ are immediate; we spell out the remaining cases.

- φ is an atom $p \in \mathcal{P}$. Suppose $s \sim s'$ and let Z be a bisimulation with $s\overline{Z}s'$. Suppose $M, s \models p$. Then $s \subseteq V(p)$. Now take any $w' \in s'$. Since $s\overline{Z}s'$ there is $w \in s$ with wZw', and since Z is a bisimulation, $w \in V(p) \iff w' \in V'(p)$. Since $w \in V(p)$, we have $w' \in V'(p)$. Since w' was arbitrary in s', it follows that $s' \subseteq V(p)$ and so $M', s' \models p$. The converse direction is analogous.
- $\varphi = \bot$. Simply observe that if $s \sim s'$, either s, s' are both empty or neither is.
- $\varphi = (\psi \to \chi)$. Suppose $s \sim s'$. We first show that every $t \subseteq s$ is bisimilar to some $t' \subseteq s'$ and vice versa. Let Z be a bisimulation with $s\overline{Z}s'$. Given $t \subseteq s$, define $t' = \{w' \in s' \mid \exists w \in t : wZw'\}$. We claim that $t\overline{Z}t'$. To see that this is the case, take $w \in t$. Since $s\overline{Z}s'$ we have wZw' for some $w' \in s'$, and then $w' \in t'$ by definition of t'. So for all $w \in t$ there is a $w' \in t'$ with wZw'. The converse is obvious from the definition of t'. So $t \sim t'$. The claim that every subset of s' is bisimilar to some subset of s is proved analogously.

  Now suppose $M, s \not\models \psi \to \chi$. Then there is a state $t \subseteq s$ with $M, t \models \psi$ and $M, t \not\models \chi$. By the previous argument there is $t' \subseteq s'$ with $t \sim t'$. By induction hypothesis we have $M', t' \models \psi$, and $M', t' \not\models \chi$, so $M', s' \not\models \psi \to \chi$. The converse is proved analogously.
- $\varphi = (\psi \Rrightarrow \chi)$. Suppose $s \sim s'$ and let Z be a bisimulation with $s\overline{Z}s'$. Suppose $M, s \not\models (\psi \Rrightarrow \chi)$. So there are a world $w \in s$ and a state $t \in \Sigma(w)$ such that $M, t \models \psi$ but $M, t \not\models \chi$. Since $s\overline{Z}s'$, there is a world $w' \in s'$ with wZw', and then by the Forth condition there is a $t' \in \Sigma'(w')$ with $t\overline{Z}t'$, and thus with $t \sim t'$. By induction hypothesis we have $M', t' \models \psi$ but $M', t' \not\models \chi$, which shows that $M', s' \not\models (\varphi \Rrightarrow \psi)$. The converse is proved analogously. $\square$

B Propositional axioms for InqCM

The propositional axioms of InqCM are all instances of the following schemata, where $\varphi, \psi, \chi \in \mathcal{L}$ and $\alpha \in \mathcal{L}_!$:

- $\varphi \to (\psi \to \varphi)$
- $(\varphi \to (\psi \to \chi)) \to (\varphi \to \psi) \to (\varphi \to \chi)$
- $\varphi \to (\psi \to \varphi \wedge \psi)$
- $\varphi \wedge \psi \to \varphi, \quad \varphi \wedge \psi \to \psi$

- $\varphi \to \varphi \vee\!\!\!\vee \psi, \quad \psi \to \varphi \vee\!\!\!\vee \psi$
- $(\varphi \to \chi) \to ((\psi \to \chi) \to (\varphi \vee\!\!\!\vee \psi \to \chi))$
- $\bot \to \varphi$
- $\neg\neg\alpha \to \alpha$
- $(\alpha \to \varphi \vee\!\!\!\vee \psi) \to (\alpha \to \varphi) \vee\!\!\!\vee (\alpha \to \psi)$

The first seven schemata are simply schemata for intuitionistic logic with $\vee\!\!\!\vee$ identified with intuitionistic disjunction, while the last two schemata are specific to inquisitive logic (see [10] for discussion of their significance).

C Proof of Lemma 5.4

We first show the following claim:

$$\text{If } \vdash \varphi, \text{ then } \vdash \alpha \text{ for some } \alpha \in \mathcal{R}(\varphi) \tag{C.1}$$

We proceed by induction on the length of the shortest proof of φ. If φ is provable with a proof of length 1, φ is an axiom. Suppose φ is a propositional axiom: then we can check case-by-case that some resolution of φ is a classical tautology. By way of example, suppose φ is an instance of $\psi \to (\chi \to \psi)$; then the following resolution of φ is a classical tautology:

$$\bigwedge_{\alpha \in \mathcal{R}(\psi)} (\alpha \to \bigwedge_{\beta \in \mathcal{R}(\chi)} (\beta \to \alpha))$$

But in restriction to declaratives, our system contains a complete set of axioms for classical propositional logic, and so it proves all classical tautologies.

Next, suppose φ is an instance of one of the modal schemata (`Tran`), (`Conj`), or (`Disj`). Then φ is a declarative, and so $\mathcal{R}(\varphi) = \{\varphi\}$ (recall that $\mathcal{R}(\alpha) = \{\alpha\}$ for every declarative α). Thus, (C.1) holds trivially.

Now consider the inductive case. We have only two cases to consider:

(i) $\varphi = (\psi \Rrightarrow \chi)$ can be obtained by (`CN`). In this case, φ is a declarative, so $\mathcal{R}(\varphi) = \{\varphi\}$ and (C.1) holds trivially.

(ii) φ can be obtained by (`MP`) from formulas $\psi \to \varphi$ and ψ which are provable with shorter proofs. By induction hypothesis, the system proves a resolution γ of $\psi \to \varphi$ and a resolution β_0 of ψ. By definition of resolutions of an implication, γ has the form $\bigwedge_{\beta \in \mathcal{R}(\psi)} (\beta \to f(\beta))$ for some $f : \mathcal{R}(\psi) \to \mathcal{R}(\varphi)$. But then, clearly, $f(\beta_0)$ is a provable resolution of φ.

This completes the inductive proof of (C.1). To prove Lemma 5.4, suppose $\Gamma \subseteq \mathcal{L}_!$ and $\Gamma \vdash \varphi$. This means that $\vdash \gamma_1 \wedge \dots \wedge \gamma_n \to \varphi$ for some $\gamma_1, \dots, \gamma_n \in \Gamma$. Let $\gamma := \gamma_1 \wedge \dots \wedge \gamma_n$: since γ is a declarative we have $\mathcal{R}(\gamma) = \{\gamma\}$. Thus, by definition of resolutions of an implication, we have:

$$\mathcal{R}(\gamma \to \varphi) = \{\gamma \to f(\gamma) \mid f : \{\gamma\} \to \mathcal{R}(\varphi)\} = \{\gamma \to \alpha \mid \alpha \in \mathcal{R}(\varphi)\}$$

Since $\vdash \gamma \to \varphi$, it follows from C.1 that $\vdash \gamma \to \alpha$ for some resolution $\alpha \in \mathcal{R}(\varphi)$, which shows that $\Gamma \vdash \alpha$. □

D Proof of Lemma 7.5 (existence lemma).

In this appendix we provide the missing details from the proof of Lemma 7.5.

Step 1. Show that if $\Phi \cup \{\chi\} \Rrightarrow_\Gamma \Psi$ and $\Phi \Rrightarrow_\Gamma \Psi \cup \{\chi\}$, then $\Phi \Rrightarrow_\Gamma \Psi$.

Suppose $\Phi, \chi \Rrightarrow_\Gamma \Psi$ and $\Phi \Rrightarrow_\Gamma \Psi, \chi$. So there are $\varphi_1, \dots, \varphi_n, \varphi_{n+1}, \dots, \varphi_{n+m} \in \Phi$ and $\psi_1, \dots, \psi_h, \psi_{h+1}, \dots, \psi_{h+k} \in \Psi$ such that:

$$((\chi \wedge \bigwedge_{1=i}^{n} \varphi_i) \Rrightarrow \bigvee\!\!\!\!\bigvee_{1=i}^{h} \psi_i) \in \Gamma \qquad ((\bigwedge_{i=1}^{m} \varphi_{n+i}) \Rrightarrow (\chi \vee\!\!\!\vee \bigvee\!\!\!\!\bigvee_{i=1}^{k} \psi_{h+i})) \in \Gamma$$

We show that Γ contains $((\bigwedge_{i \leq n+m} \varphi_i) \Rrightarrow (\bigvee\!\!\!\!\bigvee_{i \leq h+k} \psi_i))$, witnessing $\Phi \Rrightarrow_\Gamma \Psi$. To ease notation, we spell out the details for the case $n = m = h = k = 1$, but the general case is completely analogous. It suffices to show that:

$$(\varphi_1 \wedge \chi \Rrightarrow \psi_1), (\varphi_2 \Rrightarrow \psi_2 \vee\!\!\!\vee \chi) \vdash (\varphi_1 \wedge \varphi_2 \Rrightarrow \psi_1 \vee\!\!\!\vee \psi_2)$$

Since the formulas on the left-hand-side are in Γ and Γ is closed under deduction of declaratives, so is the conclusion.

First note that the following formula is provable in the propositional component of the proof system using the standard axioms for $\wedge$ and $\vee\!\!\!\vee$:

$$\varphi_1 \wedge (\psi_2 \vee\!\!\!\vee \chi) \to \psi_2 \vee\!\!\!\vee (\varphi_1 \wedge \chi) \tag{D.1}$$

In the following derivation, we indicate explicitly only the modal axioms and rules involved in the reasoning, omitting reference to propositional axioms and rules. For simplicity, we use the formulas $(\varphi_1 \wedge \chi \Rrightarrow \psi_1)$ and $(\varphi_2 \Rrightarrow \psi_2 \vee\!\!\!\vee \chi)$ as premises; this is legitimate since we will not use conditional necessitation on these formulas or anything inferred from them. Rewriting the argument with the relevant formulas used as antecedents is tedious but straightforward.

1.	$\varphi_1 \wedge \chi \Rrightarrow \psi_1$	(premise)
2.	$\varphi_2 \Rrightarrow \psi_2 \vee\!\!\!\vee \chi$	(premise)
3.	$\varphi_1 \wedge \varphi_2 \Rrightarrow \varphi_1$	(`CN`) from axiom $\varphi_1 \wedge \varphi_2 \to \varphi_1$
4.	$\varphi_1 \wedge \varphi_2 \Rrightarrow \varphi_2$	(`CN`) from axiom $\varphi_1 \wedge \varphi_2 \to \varphi_2$
5.	$\varphi_1 \wedge \varphi_2 \Rrightarrow \psi_2 \vee\!\!\!\vee \chi$	(`Tran`), 4, 2
6.	$\varphi_1 \wedge \varphi_2 \Rrightarrow \varphi_1 \wedge (\psi_2 \vee\!\!\!\vee \chi)$	(`Conj`), 3, 5
7.	$\varphi_1 \wedge (\psi_2 \vee\!\!\!\vee \chi) \Rrightarrow \psi_2 \vee\!\!\!\vee (\varphi_1 \wedge \chi)$	(`CN`) from (D.1)
8.	$\varphi_1 \wedge \varphi_2 \Rrightarrow \psi_2 \vee\!\!\!\vee (\varphi_1 \wedge \chi)$	(`Tran`), 6, 7
9.	$\psi_1 \Rrightarrow \psi_1 \vee\!\!\!\vee \psi_2$	(`CN`) from axiom $\psi_1 \to \psi_1 \vee\!\!\!\vee \psi_2$
10.	$\psi_2 \Rrightarrow \psi_1 \vee\!\!\!\vee \psi_2$	(`CN`) from axiom $\psi_2 \to \psi_1 \vee\!\!\!\vee \psi_2$
11.	$\varphi_1 \wedge \chi \Rrightarrow \psi_1 \vee\!\!\!\vee \psi_2$	(`Tran`), 1, 9
12.	$\psi_2 \vee\!\!\!\vee (\varphi_1 \wedge \chi) \Rrightarrow \psi_1 \vee\!\!\!\vee \psi_2$	(`Disj`), 10, 11
13.	$\varphi_1 \wedge \varphi_2 \Rrightarrow \psi_1 \vee\!\!\!\vee \psi_2$	(`Tran`), 8, 12

Step 2. We partition $\mathcal{L}$ into two sets L, R with $\varphi \in \mathsf{L}$, $\psi \in \mathsf{R}$, and $\mathsf{L} \not\Rrightarrow_\Gamma \mathsf{R}$.

Fix an enumeration $(\chi_n)_{n \in \mathbb{N}}$ of all formulas in $\mathcal{L}$ (we assume $\mathcal{L}$ is countable, though this is not strictly necessary). We define a sequence of sets $(\mathsf{L}_n)_{n \in \mathbb{N}}$ and $(\mathsf{R}_n)_{n \in \mathbb{N}}$ as follows:

- $\mathsf{L}_0 = \{\varphi\}, \mathsf{R}_0 = \{\psi\}$
- if $\mathsf{L}_n \cup \{\chi_n\} \not\Rrightarrow_\Gamma \mathsf{R}_n$ we let $\mathsf{L}_{n+1} := \mathsf{L}_n \cup \{\chi_n\}$ and $\mathsf{R}_{n+1} = \mathsf{R}_n$
- if $\mathsf{L}_n \cup \{\chi_n\} \Rrightarrow_\Gamma \mathsf{R}_n$ we let $\mathsf{L}_{n+1} := \mathsf{L}_n$ and $\mathsf{R}_{n+1} = \mathsf{R}_n \cup \{\chi_n\}$

We will show by induction on n that $\mathsf{L}_n \not\Rrightarrow_\Gamma \mathsf{R}_n$. For $n = 0$ this amounts to $(\varphi \Rrightarrow \psi) \notin \Gamma$, which is true by assumption. Now suppose this is true for n and consider $n + 1$. If $\mathsf{L}_n \cup \{\chi_n\} \not\Rrightarrow_\Gamma \mathsf{R}_n$, the claim is obvious from the definition. So, suppose $\mathsf{L}_n \cup \{\chi_n\} \Rrightarrow_\Gamma \mathsf{R}_n$. Since by induction hypothesis $\mathsf{L}_n \not\Rrightarrow_\Gamma \mathsf{R}_n$, Step 1 above implies $\mathsf{L}_n \not\Rrightarrow_\Gamma \mathsf{R}_n \cup \{\chi_n\}$, which by definition amounts to $\mathsf{L}_{n+1} \not\Rrightarrow_\Gamma \mathsf{R}_{n+1}$.

Now let $\mathsf{L} = \bigcup_{n\in\mathbb{N}} \mathsf{L}_n$ and $\mathsf{R} = \bigcup_{n\in\mathbb{N}} \mathsf{R}_n$. We have $\mathsf{L} \not\Rrightarrow_\Gamma \mathsf{R}$, since otherwise there would be an n such that $\mathsf{L}_n \Rrightarrow_\Gamma \mathsf{R}_n$. Moreover, L and R form a partition of $\mathcal{L}$: by construction, every formula occurs in either set, and no formula can occur in both (if $\chi \in \mathsf{L} \cap \mathsf{R}$, then since $(\chi \Rrightarrow \chi) \in \Gamma$ we would have $\mathsf{L} \Rrightarrow_\Gamma \mathsf{R}$).

References

[1] Abdou, J. and H. Keiding, "Effectivity Functions in Social Choice," Kluwer, 1991.

[2] Aiello, M., J. Van Benthem and G. Bezhanishvili, *Reasoning about space: the modal way*, Journal of Logic and Computation **13** (2003), pp. 889–920.

[3] Alur, R., T. Henzinger and O. Kupferman, *Alternating-time temporal logic*, Journal of the ACM **49** (2002), pp. 672–713.

[4] van Benthem, J., "Logic in games," MIT press, 2014.

[5] van Benthem, J., N. Bezhanishvili and S. Enqvist, *A new game equivalence, its logic and algebra*, Journal of Philosophical Logic **48** (2019), pp. 649–684.

[6] van Benthem, J., N. Bezhanishvili, S. Enqvist and J. Yu, *Instantial neighbourhood logic*, The Review of Symbolic Logic **10** (2017), pp. 116–144.

[7] Brown, M. A., *On the logic of ability*, Journal of Philosophical Logic **17** (1988), pp. 1–26.

[8] Chellas, B. F., "Modal logic: an introduction," Cambridge university press, 1980.

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

[10] Ciardelli, I., "Questions in logic," Ph.D. thesis, Institute for Logic, Language and Computation, University of Amsterdam (2016).

[11] Ciardelli, I., *Dependence statements are strict conditionals*, in: G. Bezhanishvili, G. D'Agostino, G. Metcalfe and T. Studer, editors, *Advances in Modal Logic (AIML)* (2018), pp. 123–142.

[12] Ciardelli, I., J. Groenendijk and F. Roelofsen, "Inquisitive Semantics," Oxford University Press, 2018.

[13] Ciardelli, I. and M. Otto, *Inquisitive bisimulation*, The Journal of Symbolic Logic **86** (2021), pp. 77–109.

[14] Ciardelli, I. and F. Roelofsen, *Inquisitive logic*, Journal of Philosophical Logic **40** (2011), pp. 55–94.

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

[16] Gabbay, D., "Semantical investigations in Heyting's intuitionistic logic," Springer, 1981.

[17] van Gessel, T., *Action models in inquisitive logic*, Synthese **197** (2020), pp. 3905–3945.

[18] Goranko, V. and S. Enqvist, *Socially friendly and group protecting coalition logics*, in: *17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, The International Foundation for Autonomous Agents and Multiagent Systems, 2018, pp. 372–380.

[19] Goranko, V. and G. van Drimmelen, *Complete axiomatization and decidability of alternating-time temporal logic*, Theoretical Computer Science **353** (2006), pp. 93–117.

[20] Grilletti, G., *Completeness for the classical antecedent fragment of inquisitive first-order logic*, Journal of Logic, Language, and Information **30** (2021), pp. 725–751.

[21] Hansen, H. H., "Monotonic modal logics," Ph.D. thesis, Universiteit van Amsterdam (2003).

[22] Meißner, S. and M. Otto, *A first-order framework for inquisitive modal logic*, The Review of Symbolic Logic (2021), pp. 1–23.

[23] Montague, R., *Universal grammar*, Theoria **36** (1970), pp. 373–398.

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

[25] Pauly, M., *A modal logic for coalitional power in games*, Journal of logic and computation **12** (2002), pp. 149–166.

[26] Punčochář, V. and I. Sedlár, *Epistemic extensions of substructural inquisitive logics*, Journal of Logic and Computation **31** (2021), pp. 1820–1844.

[27] Scott, D., *Advice on modal logic*, in: K. Lambert, editor, *Philosophical Problems in Logic: Some Recent Developments*, Springer Netherlands, 1970 pp. 142–173.