COUNTING TO INFINITY: GRADED MODAL LOGIC WITH AN INFINITY DIAMOND

IGNACIO BELLAS ACOSTA
Independent Scholar
and
YDE VENEMA
Institute for Logic, Language and Computation, Universiteit van Amsterdam

Abstract. We extend the languages of both basic and graded modal logic with the infinity diamond, a modality that expresses the existence of infinitely many successors having a certain property. In both cases we define a natural notion of bisimilarity for the resulting formalisms, that we dub $\mathtt{ML}^\infty$ and $\mathtt{GML}^\infty$, respectively. We then characterise these logics as the bisimulation-invariant fragments of the naturally corresponding predicate logic, viz., the extension of first-order logic with the infinity quantifier. Furthermore, for both $\mathtt{ML}^\infty$ and $\mathtt{GML}^\infty$ we provide a sound and complete axiomatisation for the set of formulas that are valid in every Kripke frame, we prove a small model property with respect to a widened class of weighted models, and we establish decidability of the satisfiability problem.

§1. Introduction. The aim of this paper is to study some properties of a modality $\Diamond^\infty$, that we shall refer to as the *infinity diamond* because of its interpretation in a standard Kripke model $\mathbb{S} = (S, R, V)$:

$$\mathbb{S}, s \Vdash \Diamond^\infty\varphi \text{ if there are infinitely many states } t \text{ such that } Rst \text{ and } \mathbb{S}, t \Vdash \varphi.$$

Although this interpretation makes $\Diamond^\infty$ a quite natural modality, it seems to have received very little attention in the literature.[1]

Rather than investigate the properties of the infinity diamond in isolation, we study the effects of adding $\Diamond^\infty$ to a more basic modal language. In the simplest setting we define $\mathtt{ML}^\infty$ as the language we obtain by adding $\Diamond^\infty$ to the basic modal language that features one modality $\Diamond$. Perhaps a more natural setting is where we add $\Diamond^\infty$ to the language $\mathtt{GML}$ of *graded modal logic*, since this logic already has a so-called *counting modality* $\Diamond^k$ for every natural number k, with the following interpretation:

$$\mathbb{S}, s \Vdash \Diamond^k\varphi \text{ if there are at least } k \text{ many states } t \text{ such that } Rst \text{ and } \mathbb{S}, t \Vdash \varphi.$$

These counting modalities were introduced in the early 1970s by Goble [15] and Fine [14], who gave a complete axiomatisation for the set of valid formulas. The name “graded modalities” originates with the eponymous paper [12] by Fattorosi-

Received: August 2, 2021.
2020 *Mathematics Subject Classification*: Primary 03B45, Secondary 03C80..
Keywords and phrases: graded modal logic, infinity diamond, generalised quantifiers, bisimulation invariance, completeness..

[1] We discuss the few references that we are aware of at the end of this introduction, when we mention related work.

doi:10.1017/S1755020322000247

Barnaba and De Caro, which contains results on completeness (independently of Fine) and decidability, and is the first of a line of papers in which various aspects of counting modalities were studied. Subsequently, graded modal logic increasingly gained attention as a convenient formalism for describing and reasoning about (finite) quantities in a setting of modal or description logics [17], and as an interesting logic in its own right. The literature is by now too extensive to be reviewed here; two notable results are that the finite model property was established by van der Hoek [18], and that Tobies [27] showed that the satisfiability problem for GML can be solved in polynomial space. Of specific interest to our investigations is the work of de Rijke [24] who introduced a notion of bisimulation that is appropriate for graded modal logic in the sense that he showed GML to correspond to the bisimulation-invariant fragment of first-order logic, thus transferring van Benthem's seminal characterisation of basic modal logic to the setting with counting modalities.

Our main goal is to investigate to which degree these results can be extended to the setting where we add the infinite counting modality to graded modal logic. The properties that we will consider here concern expressiveness, axiomatisations and decidability.

Before we continue to discuss our results in more detail, we need to discuss the ambient predicate logic for $\mathtt{GML}^{\infty}$ and $\mathtt{ML}^{\infty}$. Since the semantics of $\Diamond^{\infty}$ cannot be encoded in first-order logic, the natural candidate for such an ambient logic is the extension $\mathtt{FO}^{\infty}$ of first-order logic with the *generalised quantifier* "there exists infinitely many" that we shall denote as $\exists^{\infty}$.

Generalised quantifiers were introduced by Mostowski in [21], and in a more general sense by Lindström in [20], motivated by the observation that the standard first-order quantifiers do not suffice for expressing some basic mathematical concepts. Since then, they have attracted a lot of interest, to the effect that nowadays their study constitutes a well-established field of logic [1, 29]. The infinity quantifier, which is part of the broader family of *cardinal quantifiers*, provides a natural extension of first-order logic, but it was soon discovered that $\mathtt{FO}^{\infty}$ lacks many important metalogical properties, including compactness, finite axiomatisability and the Löwenheim–Skolem property (see [1]), but also Craig interpolation [22]. Probably due to these negative results, $\mathtt{FO}^{\infty}$ lost prominence to better behaved cardinal logics.

Negative properties of a logic, however, are not necessarily inherited by its *fragments*. For instance, the *monadic fragment* of $\mathtt{FO}^{\infty}$ is quite well behaved: Mostowski [21] already proved that this fragment is decidable, and recently, some model-theoretic characterisation theorems were obtained by Carreiro et al. [6].

In the same spirit, the results in this paper show that positive results can be obtained for $\mathtt{GML}^{\infty}$, seen as a fragment of $\mathtt{FO}^{\infty}$. First of all, we consider the (relative) expressiveness of $\mathtt{GML}^{\infty}$, as a language for describing pointed Kripke models. We associate with $\mathtt{GML}^{\infty}$ a natural structural notion of equivalence between pointed models, to be called ∞, #-bisimilarity, and we show that, with respect to this notion,

$$\mathtt{GML}^{\infty} \text{ is the bisimulation-invariant fragment of } \mathtt{FO}^{\infty}.$$

That is, we prove that every $\mathtt{FO}^{\infty}$-formula that is invariant under #, ∞-bisimulations, must be equivalent to some formula in $\mathtt{GML}^{\infty}$. Our proof, which is game-theoretic in nature, basically follows the pattern laid out by Rosen's generalisation [25] of van Benthem's bisimulation invariance result for ordinary modal logic. More specifically, we extend recent work by Otto [23] who took a game-theoretical perspective on de

Rijke's bisimulation-invariance result for graded modal logic. To include the infinity quantifier and diamond, we use the Ehrenfeucht-style model comparison game for monotone generalised quantifiers as designed by Krawczyk and Krynicki [19].

We then move on to questions about axiomatisability and decidability. Taking a standard axiomatisation **G** for graded modal logic as a basis, we add some axioms and a derivation rule to capture the behaviour of the infinity diamond. Thus arriving at a logic $\mathbf{G}^\infty$, we then show that

$\mathbf{G}^\infty$ is a sound and complete axiomatisation of the set of all $\mathtt{GML}^\infty$-validities,

where we call a $\mathtt{GML}^\infty$-formula *valid* if it holds in every state of every Kripke model. To prove this result, we basically apply the canonical model method, which De Caro adapted to the setting of graded modal logic [9] by generalising the Kripke semantics for graded modal logic to that of models in which successors of a given state come with a certain weight. Our perspective on these weighted models will be coalgebraic [7]; that is, we will consider weighted models as coalgebras for a certain endofunctor $\mathtt{B}^\infty$ on the category with sets as objects and functions as arrows. This widening of perspective also allows us to prove that $\mathtt{GML}^\infty$ enjoys a *small model property*, a property that for obvious reasons does not make sense in the setting of Kripke models. Concretely, we will show that

every satisfiable $\mathtt{GML}^\infty$-formula is satisfiable in a weighted model of bounded size,

where "bounded" means that the size of the model is not just finite, but in fact bounded by some parameters of the formula. As an immediate corollary of the latter result, we then prove a *decidability result* for our logic:

it is decidable whether a given $\mathtt{GML}^\infty$-formula is satisfiable in a Kripke model.

Finally, although we only mentioned our results for $\mathtt{GML}^\infty$ here, in each case we have an analogous result for $\mathtt{ML}^\infty$.

In the next section we will introduce the syntax and semantics of the languages $\mathtt{ML}^\infty$ and $\mathtt{GML}^\infty$; we define the appropriate notions of bisimulation for these languages, we discuss the coalgebraic perspective on $\mathtt{GML}^\infty$, and we review the logic $\mathtt{FO}^\infty$ and its Ehrenfeucht-style model comparison games. In Section 3 we prove that $\mathtt{GML}^\infty$ is the bisimulation invariant fragment of $\mathtt{FO}^\infty$. We conclude this paper by proving, in Section 4, the other main results of the paper: completeness, the small coalgebra property, and decidability of the satisfiability problems for both $\mathtt{ML}^\infty$ and $\mathtt{GML}^\infty$.

As mentioned above, the infinity diamond seems to have received little attention in the literature on modal logic; here we briefly discuss the references that we are aware of. In linear temporal logic, the formula GFp may be used to express that p will be true infinitely often; the concept also occurs in branching time logics of the CTL variety, where the notation $\overset{\infty}{F}$ has been used for the infinitely often connective [10].[2] To the best of our knowledge, however, in this context this modality has not been investigated as such. Perhaps the first explicit mention of the infinity modality is in the paper [12] that initiated the *Graded Modalities* series. The modality features explicitly in the sixth paper of the series [13], in the context of an infinitary version of graded modal logic, where countable disjunctions and conjunctions are allowed. In this setting, $\Diamond^\infty$ can

[2] We are grateful to Rob Goldblatt for pointing this out.

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

be expressed by the other modalities: $\Diamond^\infty\varphi \equiv \bigwedge_n \Diamond^n\varphi$, and because of the different nature of this infinitary language, the completeness result by Fattorosi-Barnaba and Grassotti has no direct bearing on our work. This is different in the case of the seventh *Graded Modalities* paper by Fattorosi-Barnaba and Balestrini, which discusses what is essentially our logic $\mathtt{ML}^\infty$. The main result of [11] is a completeness result stating that the logic $\mathbf{K}^\infty$ that we define in Section 4 is a complete axiomatisation of the set of all $\mathtt{ML}^\infty$-validities. As an interesting variation of our logic $\mathtt{ML}^\infty$, we mention the work by van Benthem et al. [4]. They introduce a modality $\bullet$ which is interpreted in Kripke models by stating that $\bullet\varphi$ holds in a state if it has infinitely many *reflexive* successors where φ holds. Among other things, the authors provide a finite axiomatisation and they show that the resulting logic has the Craig interpolation property.

Finally, this paper grew out of the M.Sc. thesis [2], written by the first author under the supervision of the second, which more or less covers the results on the logic $\mathtt{ML}^\infty$ that we report on.

§2. Preliminaries. In this section we formally define the logics $\mathtt{ML}^\infty$ and $\mathtt{GML}^\infty$, we introduce some of the basic concepts that will be used further on in the paper, and we formulate some of our main results. In Section 2.1 we introduce the languages $\mathtt{ML}^\infty$ and $\mathtt{GML}^\infty$, as well as the basic syntactic and semantic concepts related to these languages. In Section 2.2 we adapt the well-known concept of bisimulation to account for the semantic behaviour of the $\Diamond^\infty$ modality. In Section 2.3 we present the reader with an alternative, coalgebraic semantics for our languages, and we relate this to the standard Kripke semantics. We conclude this section with a brief introduction to the basic syntactic and semantic concepts related to the predicate logic $\mathtt{FO}^\infty$.

2.1. Graded modal logic with the infinity diamond. In this subsection we provide some basic syntactic and semantic definitions concerning the modal logics $\mathtt{ML}^\infty$ and $\mathtt{GML}^\infty$. The language of the latter logic will have a family of diamonds, each indexed by either a natural number or infinity; formally, we introduce this index set as follows.

Definition 2.1. *We let $\mathbb{N}^\infty$ denote the set $\mathbb{N}^\infty := \mathbb{N} \cup \{\infty\}$, and we order this set by extending the standard order of the natural numbers in the obvious way, i.e., we declare every natural number to be strictly smaller than ∞.*

As a convention, we will use the Greek letters κ, λ for arbitrary elements of $\mathbb{N}^\infty$, and lower case roman letters $k, l, m, \dots$ for indices in $\mathbb{N}$.

We can now define the languages $\mathtt{ML}^\infty$ and $\mathtt{GML}^\infty$.

Definition 2.2. *Given a set P of proposition letters, we define the language $\mathtt{GML}^\infty(P)$ as follows:*

$$\varphi ::= p \mid \neg\varphi \mid \varphi_0 \vee \varphi_1 \mid \Diamond^\kappa\varphi,$$

where $p \in P$ and $\kappa \in \mathbb{N}^\infty$. The sets $\mathtt{ML}^\infty$, $\mathtt{GML}$ and $\mathtt{ML}$ are defined as those fragments of $\mathtt{GML}^\infty$ where the occurrence of the modalities $\Diamond^\kappa$ is restricted as follows:

$\mathtt{ML}^\infty$	$\mathtt{GML}$	$\mathtt{ML}$
$\kappa \in \{1, \infty\}$	$\kappa \in \mathbb{N}$	$\kappa = 1$

We will use the Boolean connectives $\wedge$, $\to$ and $\leftrightarrow$ as the standard abbreviations, and use modalities $\Box^\kappa$ (for any $\kappa \in \mathbb{N}^\infty$), and $\Diamond^{=k}$ (for any $k \in \mathbb{N}$), given by the following definition:

$$\begin{array}{lcl} \Box^{\kappa}\varphi & := & \neg\Diamond^{\kappa}\neg\varphi, \\ \Diamond^{=k}\varphi & := & \Diamond^{k}\varphi \wedge \neg\Diamond^{k+1}\varphi. \end{array}$$

We occasionally denote $\Diamond^1$ and $\Box^1$ simply as $\Diamond$ and $\Box$, especially in the context of $\mathtt{ML}^{\infty}$. Clearly $\Box^{\kappa}$ is the (Boolean) dual of $\Diamond^{\kappa}$, and $\Diamond^{=k}\varphi$ will have the meaning that the current state has *exactly* k successors where φ holds.

In the sequel we will need three complexity measures for a $\mathtt{GML}^{\infty}$-formula: *size*, *depth* and *rank*. The first two notions are standard, and the *rank* of a formula quantifies the "highest" modality that occurs in the formula (with $\Diamond^{\infty}$ counting as $\Diamond^1$):

Definition 2.3. *Let φ be some arbitrary but fixed $\mathtt{GML}^{\infty}$-formula. The* size $|\varphi|$ *of φ is the number of its subformulas (where the notion of subformula is defined in the standard way). The* (modal) depth *of φ is the maximal height of a sequence of nested modalities occurring in φ.*

The rank *of φ, denoted by $r(\varphi)$, is recursively defined as follows:*

$$\begin{array}{lcl} r(p) & := & 0, \\ r(\neg\varphi) & := & r(\varphi), \\ r(\varphi_0 \vee \varphi_1) & := & \max\{r(\varphi_0), r(\varphi_1)\}, \\ r(\Diamond^{k}\varphi) & := & \max\{k, r(\varphi)\}, \\ r(\Diamond^{\infty}\varphi) & := & \max\{1, r(\varphi)\}. \end{array}$$

The rank *of a finite set Σ of formulas is simply defined as $r(\Sigma) := \max\{r(\varphi) \mid \varphi \in \Sigma\}$.*

We can now introduce the semantics of $\mathtt{GML}^{\infty}$. The intended semantics of this language is that of Kripke models.

Definition 2.4. *A* (Kripke) frame *is a pair (S, R) where $R \subseteq S \times S$ is a binary relation on S. Elements of S will be called* states *or* worlds. *For a fixed Kripke frame, the set $R(s) := \{t \in S \mid (s,t) \in R\}$ will denote the set of successors of s. A P-*valuation *on a set S is a map $V : P \to \wp(S)$; when the set P of proposition letters is understood, we will usually simply speak of a* valuation. *A* (Kripke) model *is a triple $\mathbb{S} = (S, R, V)$ such that (S, R) is a Kripke frame, the* underlying frame *of $\mathbb{S}$, and V is a valuation on S. Finally, a* pointed model *is simply a pair $(\mathbb{S}, s)$ consisting of a Kripke model $\mathbb{S}$ and a point s in $\mathbb{S}$.*

Definition 2.5. *Let $\mathbb{S} = (S, R, V)$ be a Kripke model. By a straightforward induction on the complexity of formulas we define a* satisfaction relation $\Vdash_{\mathbb{S}}$ *between states and $\mathtt{GML}^{\infty}$-formulas, where we write $\mathbb{S}, s \Vdash \varphi$ rather than $s \Vdash_{\mathbb{S}} \varphi$:*

$$\begin{array}{lll} \mathbb{S}, s \Vdash p & \textit{if} & s \in V(p), \\ \mathbb{S}, s \Vdash \neg\varphi & \textit{if} & \mathbb{S}, s \nVdash \varphi, \\ \mathbb{S}, s \Vdash \varphi_0 \vee \varphi_1 & \textit{if} & \mathbb{S}, s \Vdash \varphi_0 \textit{ or } \mathbb{S}, s \Vdash \varphi_1, \\ \mathbb{S}, s \Vdash \Diamond^{\kappa}\varphi. & \textit{if} & |\{t \in R(s) \mid \mathbb{S}, t \Vdash \varphi\}| \geq \kappa. \end{array}$$

In case $\mathbb{S}, s \Vdash \varphi$ we say that φ is true *at* s *(in $\mathbb{S}$) or that* s satisfies φ.

A formula φ is satisfiable *if it is satisfied at some state in some model, and* valid *if it is true at every state in every model.*

We call two pointed models $(\mathbb{S}, s)$ and $(\mathbb{S}', s')$ $\mathtt{GML}^{\infty}$-equivalent, *and we denote this as $(\mathbb{S}, s) \equiv^{\infty,\#} (\mathbb{S}', s')$, if they satisfy the same set of $\mathtt{GML}^{\infty}$-formulas. The relation $\equiv^{\infty}$ of* $\mathtt{ML}^{\infty}$-equivalence *is defined analogously. More generally, given a fragment $\mathtt{L}$ of $\mathtt{GML}^{\infty}$,*

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

we call $(\mathbb{S}, s)$ *and* $(\mathbb{S}', s')$ $\mathtt{L}$*-equivalent if they satisfy the same* $\mathtt{L}$*-formulas; in the case* $\mathtt{L}$ *consists of the atomic formulas we use the phrase* atomic equivalence.

The dual notions of satisfiability and validity are among the key topics of interest in this paper. In Section 4 we will give fairly simple and intuitive finite axiomatisations for the set of validities of $\mathtt{ML}^{\infty}$ and $\mathtt{GML}^{\infty}$, respectively. Here we already formulate the following decidability theorem, which a fortiori applies to $\mathtt{ML}^{\infty}$ since it is a fragment of $\mathtt{GML}^{\infty}$.

Theorem 2.6 (Decidability of satisfiability). *It is decidable in elementary time whether a given* $\mathtt{GML}^{\infty}$*-formula is satisfiable or not.*

For a *proof* of this result, observe that it is an immediate corollary of Corollary 2.18 and Theorem 2.20.

2.2. Bisimilarity. In the previous subsection we introduced the modal logics $\mathtt{GML}^{\infty}$ and $\mathtt{ML}^{\infty}$. We now see how to adapt the standard notion of bisimilarity between pointed models to account for the infinitary nature of these logics. To start with, it will be convenient to use the following notations.

Definition 2.7. *Where X and Y are sets, and X is a subset of Y, we write* $X \subseteq_{\omega} Y$ *if X is finite, and* $X \subseteq_{\infty} Y$ *if X is infinite.*

A suitable notion of bisimilarity for $\mathtt{GML}^{\infty}$ can be defined as follows.

Definition 2.8. *Let* $\mathbb{S}_0 = (S_0, R_0, V_0)$ *and* $\mathbb{S}_1 = (S_1, R_1, V_1)$ *be two Kripke models. A relation* $Z \subseteq S_0 \times S_1$ *is a* $\infty, \#$-bisimulation *if it satisfies the following conditions, for every pair* $(w_0, w_1) \in Z$:

1. w_1 *and* w_2 *are atomically equivalent.*
2. *For every* $X_0 \subseteq_{\omega} R_0(w_0)$ *there is an* $X_1 \subseteq_{\omega} R_1(w_1)$ *such that* $|X_0| = |X_1|$ *and for every* $v_1 \in X_1$ *there is a* $v_0 \in X_0$ *with* $(v_0, v_1) \in Z$.
3. *For every* $X_1 \subseteq_{\omega} R_1(w_1)$ *there is an* $X_0 \subseteq_{\omega} R_0(w_0)$ *such that* $|X_0| = |X_1|$ *and for every* $v_0 \in X_0$ *there is a* $v_1 \in X_1$ *with* $(v_1, v_0) \in Z$.
4. *For every* $X_0 \subseteq_{\infty} R_0(w_0)$ *there is an* $X_1 \subseteq_{\infty} R_1(w_1)$ *such that for every* $v_1 \in X_1$ *there is a* $v_0 \in X_0$ *with* $(v_0, v_1) \in Z$.
5. *For every* $X_1 \subseteq_{\infty} R_1(w_1)$ *there is an* $X_0 \subseteq_{\infty} R_0(w_0)$ *such that for every* $v_0 \in X_0$ *there is a* $v_1 \in X_1$ *with* $(v_1, v_0) \in Z$.

We say that two pointed models $(\mathbb{S}_0, w_0)$ *and* $(\mathbb{S}_1, w_1)$ *are* $\infty, \#$-bisimilar, *denoted by* $\mathbb{S}_0, w_0 \mathrel{\underline{\leftrightarrow}}^{\infty,\#} \mathbb{S}_1, w_1$, *if there is a* $\infty, \#$*-bisimulation linking* w_0 *to* w_1.

For the language $\mathtt{ML}^{\infty}$ we can simplify the definition somewhat.

Definition 2.9. *Let* $\mathbb{S}_0 = (S_0, R_0, V_0)$ *and* $\mathbb{S}_1 = (S_1, R_1, V_1)$ *be two Kripke models. A relation* $Z \subseteq S_0 \times S_1$ *is a* ∞-bisimulation *if it satisfies, for every pair* $(w_0, w_1) \in Z$, *the conditions 1, 4 and 5 of Definition 2.8, together with the standard back and forth conditions*:

6. *For every* $v_0 \in R_0(w_0)$ *there is a* $v_1 \in R_1(w_1)$ *such that* $(v_0, v_1) \in Z$.
7. *For every* $v_1 \in R_1(w_1)$ *there is a* $v_0 \in R_0(w_0)$ *such that* $(v_1, v_0) \in Z$.

Two pointed models $(\mathbb{S}_0, w_0)$ *and* $(\mathbb{S}_1, w_1)$ *are* ∞-bisimilar, *denoted by* $\mathbb{S}_0, w_0 \mathrel{\underline{\leftrightarrow}}^{\infty} \mathbb{S}_1, w_1$, *if there is a* ∞*-bisimulation linking* w_0 *to* w_1.

It is one of the main goals of this paper to show that these notions of bisimilarity are indeed appropriate for the languages $\mathtt{GML}^{\infty}$ and $\mathtt{ML}^{\infty}$. The first observation is that the formulas in these languages are indeed *invariant* for the given notions of bisimulation. The proof of this result is routine.

PROPOSITION 2.10. *Let* $(\mathbb{S}, w)$ *and* $(\mathbb{S}', w')$ *be two pointed models. Then the following holds*:

1. *If* $(\mathbb{S}, w) \underline{\leftrightarrow}^{\infty,\#} (\mathbb{S}', w')$ *then* $(\mathbb{S}, w) \equiv^{\infty,\#} (\mathbb{S}', w')$.
2. *If* $(\mathbb{S}, w) \underline{\leftrightarrow}^{\infty} (\mathbb{S}', w')$ *then* $(\mathbb{S}, w) \equiv^{\infty} (\mathbb{S}', w')$.

The above notions of bisimilarity can be equivalently approached from a game theoretic perspective, via a modification of the standard bisimilarity game.

DEFINITION 2.11. *Let* $\mathbb{S}_0 = (S_0, R_0, V_0)$ *and* $\mathbb{S}_1 = (S_1, R_1, V_1)$ *be two Kripke models*; *as a convention*, *when* $\mathbb{S}_i$ *is one of these two models*, *we write* $\mathbb{S}_{-i}$ *to denote the other. We define the* bisimilarity game $\mathcal{B}^{\infty,\#}(\mathbb{S}_0, \mathbb{S}_1)$ *between two players*, *Spoiler* (*male*) *and Duplicator* (*female*). *A basic position in this game is a pair* (w_0, w_1) *of states from* $\mathbb{S}_0$ *and* $\mathbb{S}_1$, *respectively.*

A single round of the game starts and ends with such a basic position, *and consists of a short interaction between the two players*, *of the following kind.*

- *Let* (w_0, w_1) *be the starting position of the round.*
- *Spoiler chooses a structure* $\mathbb{S}_i$ *and a subset* $X_i \subseteq R_i(w_i)$.
- *Duplicator picks a subset* $X_{-i} \subseteq R_{-i}(w_{-i})$ *such that either* $|X_i| = |X_{-i}| < \omega$ *or* $|X_i|, |X_{-i}| \geq \omega$.
- *Spoiler chooses an element* $u_{-i} \in X_{-i}$.
- *Duplicator responds with an element* $u_i \in X_i$.
- *The new basic position is* (u_0, u_1).

Spoiler wins a match of the game as soon as a basic position (w_0, w_1) *is reached consisting of two states that are not atomically equivalent*, *or if in some round Duplicator fails to come up with a matching set in her response to Spoiler's first move. Duplicator wins a match of the game if*, *during some round of the match*, *Spoiler fails to pick a state in his second move. In addition*, *Duplicator also wins if she manages to survive for* ω *many rounds.*

We write $\mathcal{B}^{\infty,\#}(\mathbb{S}_0, \mathbb{S}_1)@(w_0, w_1)$ *for the* initialised version *of the game which starts at position* (w_0, w_1). *A position* (w_0, w_1) *is* winning *for a player if this player has a winning strategy in the initialised game* $\mathcal{B}^{\infty,\#}(\mathbb{S}_0, \mathbb{S}_1)@(w_0, w_1)$, *i.e.*, *a way of playing that guarantees this player will win*, *no matter how their opponent plays.*

Finally, *we write* $\mathcal{B}^{\infty}(\mathbb{S}_0, \mathbb{S}_1)$ *for the variation of the game where Spoiler can only pick subsets that are either singletons or infinite.*

The following proposition states that these games indeed provide an equivalent perspective on bisimilarity. The proof, which is routine, is left to the reader.

PROPOSITION 2.12. *Let* $(\mathbb{S}_0, w_0)$ *and* $(\mathbb{S}_1, w_1)$ *be two pointed models*, *then the following hold*:

- $(\mathbb{S}_0, w_0) \underline{\leftrightarrow}^{\infty,\#} (\mathbb{S}_1, w_1)$ *iff* (w_0, w_1) *is a winning position for Duplicator in* $\mathcal{B}^{\infty,\#}(\mathbb{S}_0, \mathbb{S}_1)$.
- $(\mathbb{S}_0, w_0) \underline{\leftrightarrow}^{\infty} (\mathbb{S}_1, w_1)$ *iff* (w_0, w_1) *is a winning position for Duplicator in* $\mathcal{B}^{\infty}(\mathbb{S}_0, \mathbb{S}_1)$

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

2.3. Coalgebraic perspective. In this section we extend the semantics of our graded modal logic to so-called weighted models. Intuitively, the idea underlying these structures is that, given a fixed state s, we assign a certain *weight* to every state t in the structure. Thus, Kripke models are a special case of weighted models where every state t receives either the value 1 (indicating that t is a successor of s), or the weight 0 (indicating that it is not). In the general case, weights can take any arbitrary values in the set $\mathbb{N}^\infty$. Before defining the semantics of $\mathtt{GML}^\infty$ in such models, we need to define some *arithmetic* on the set $\mathbb{N}^\infty$.

DEFINITION 2.13. *We extend the addition operation on the natural numbers to the set $\mathbb{N}^\infty$ by setting $\kappa + \lambda := \infty$ whenever $\infty \in \{\kappa, \lambda\}$. Moreover, this binary operation can be extended to a finitary operation $\sum$ on $\mathbb{N}^\infty$ in the obvious way, and subsequently we define, for any set $\{\kappa_i \in \mathbb{N}^\infty \mid i \in I\}$ where I is infinite:*

$$\sum_{i \in I} \kappa_i := \begin{cases} \sum_{i \in I'} \kappa_i, & \text{if } I' := \{i \in I \mid \kappa_i \neq 0\} \text{ is finite,} \\ \infty, & \text{otherwise.} \end{cases}$$

Observe that with this definition we find that $\sum_{i \in I} \kappa_i = \infty$ if and only if ∞ occurs as one of the κ_i or there are infinitely indices i for which κ_i is non-zero.

We can now define weighted models as structures of the form $\mathbb{S} = (S, w, V)$ where S is a set of states, V is a valuation, and $w : W \times W \to \mathbb{N}^\infty$ is a function assigning a weight to every pair of states. The semantics of the graded modalities can then be given by putting $\mathbb{S}, s \Vdash \Diamond^\kappa \varphi$ if the total weight of all φ-states, as seen from s, is at least κ. In fact, this is exactly what we will do, but we will use the notation and terminology from the theory of *coalgebra*.

The theory of Universal Coalgebra [26] provides a mathematical framework for studying various kinds of state-based evolving systems such as finite state automata, Markov chains, or labelled transition systems, in a uniform manner. Kripke structures provide some further prime examples of coalgebra, and in fact the connection between modal logic and coalgebra is very tight [7]. Graded modal logic has been studied from this perspective since the work by D'Agostino and Visser [8].

DEFINITION 2.14. *Given a set S, we let $\mathsf{B}^\infty(S) := (\mathbb{N}^\infty)^S$ denote the set of $\mathbb{N}^\infty$-valued* weight functions *on S, that is, the collection of maps from S to $\mathbb{N}^\infty$. A B^∞-*coalgebra *is a pair (S, σ), where $\sigma : S \to \mathsf{B}^\infty(S)$ assigns, to every state $s \in S$, a weight function, that we will denote as either $\sigma(s)$ or σ_s. A B^∞-*model *is a triple (S, σ, V) where (S, σ) is a B^∞-coalgebra and V is a valuation on S.*

REMARK. *Although we do not stress this point in our paper, in order to fully understand and appreciate the coalgebraic perspective it is essential to realise that B^∞ is part of a* functor *on the category of sets. For readers with some category-theoretic background, we give a brief sketch of the approach. Let* Set *be the category taking sets as objects and functions as arrows, then B^∞ is set to be the* Set*-endofunctor such that for an arbitrary function $f : S \to S'$, the map $\mathsf{B}^\infty f : \mathsf{B}^\infty(S) \to \mathsf{B}^\infty(S')$ is defined as follows. For an arbitrary weight function $w : S \to \mathbb{N}^\infty$ we define $(\mathsf{B}^\infty f)(w) : S' \to \mathbb{N}^\infty$ to be*

$$(\mathsf{B}^\infty f)(w)(t') := \sum_{t \in f^{-1}(t')} w(t),$$

where $t \in S'$.

This functorial perspective is used to define a natural notion of morphism *between* B^∞*-coalgebras: we say that a map* $f : S \to S'$ *is a* morphism *from one* B^∞*-coalgebras* (S, σ) *to another* (S', σ')*, if we have that* $(\mathsf{B}^\infty f) \circ \sigma = \sigma' \circ f$*.*

This is all completely analogous to the coalgebraic perspective on Kripke structures. Kripke frames can be understood as coalgebras for the power set functor P*; here we identify a binary accessibility relation* R *on a set* S *with the function* σ_R *from* S *to* $\mathsf{P}(S)$ *mapping any state* $s \in S$ *to the collection of its successors. It is then straightforward to verify that a map* $f : S \to S'$ *is a* bounded morphism *between the Kripke frames* (S, R) *and* (S', R') *if it satisfies* $(\mathsf{P} f) \circ \sigma_R = \sigma_{R'} \circ f$*.*

Making our earlier remarks precise, we show how B^∞-models provide a natural interpretation for $\mathtt{GML}^\infty$-formulas indeed.

Definition 2.15. *Let* $\mathbb{S} = (S, \sigma, V)$ *be a* B^∞*-model. By a straightforward formula induction we define a* truth relation $\Vdash$ *between states of* $\mathbb{S}$ *and* $\mathtt{GML}^\infty$*-formulas:*

$$\begin{array}{lll} \mathbb{S}, s \Vdash p & \textit{if} & s \in V(p), \\ \mathbb{S}, s \Vdash \neg\varphi & \textit{if} & \mathbb{S}, s \nVdash \varphi, \\ \mathbb{S}, s \Vdash \varphi_1 \vee \varphi_2 & \textit{if} & \mathbb{S}, s \Vdash \varphi_1 \textit{ or } \mathbb{S}, s \Vdash \varphi_2, \\ \mathbb{S}, s \Vdash \Diamond^\kappa \varphi & \textit{if} & \sum_{\mathbb{S}, t \Vdash \varphi} \sigma_s(t) \geq \kappa. \end{array}$$

As we will see now, this approach in fact *extends* the Kripke semantics. We will identify a Kripke frame (S, R) with the B^∞-coalgebra (S, σ_R), where $\sigma_R : S \to \mathsf{B}^\infty(S)$ is given by putting, for every $s \in S$

$$\sigma_R(s)(t) := \begin{cases} 1, & \text{if } t \in R(s), \\ 0, & \text{otherwise.} \end{cases}$$

Introducing some notation, given a Kripke model $\mathbb{S} = (S, R, V)$, we will write $\mathbb{S}^\infty := (S, \sigma_R, V)$. It is then not hard to see that the truth relations of the Kripke model and its coalgebraic incarnation coincide (cf. Proposition 2.17), so that we may say that B^∞-coalgebras extend the Kripke semantics indeed.

Interestingly—in the opposite direction—we may encode B^∞-coalgebras as Kripke models as well. To do so, we simply extend De Caro's original encoding [9] to the setting with the infinity diamond.

Definition 2.16. *Let* (S, σ) *be some* B^∞*-coalgebra, and let* V *be some valuation on* S*. We let* R_σ *be the binary relation on the set* $S \times \mathbb{N}$ *given by*

$$\big((s, k), (t, m)\big) \in R_\sigma \textit{ if and only if } m < \sigma_s(t),$$

and we define the valuation V_∞ *on* $S \times \mathbb{N}$ *by putting*

$$V_\infty(p) := \{(s, k) \in S \times \mathbb{N} \mid s \in V(p)\}.$$

Finally, given a B^∞*-model* $\mathbb{S} := (S, \sigma, V)$*, we define its* associated Kripke model *as the structure*

$$\mathbb{S}_\infty := (S \times \mathbb{N}, R_\sigma, V_\infty).$$

Intuitively, the definition of R_σ ensures that any state (s, k) obtains, as its successors, as many copies of t as it is determined by the weight function σ_s when applied to the state t. Based on this observation, one may easily prove the following equivalences between the two perspectives.

PROPOSITION 2.17. *Let φ be a $\texttt{GML}^\infty$ formula.*

1) *For any Kripke model $\mathbb{S} = (S, R, V)$, and any state $s \in S$ we have that*

$$\mathbb{S}, s \Vdash \varphi \text{ if and only if } \mathbb{S}^\infty, s \Vdash \varphi.$$

2) *For any B^∞-model $\mathbb{S} = (S, \sigma, V)$, any state $s \in S$ and any $k \in \mathbb{N}$ we have that*

$$\mathbb{S}, s \Vdash \varphi \text{ if and only if } \mathbb{S}_\infty, (s, k) \Vdash \varphi.$$

Proof. Both items can be proved by a straightforward induction on the complexity of φ. We omit the details. □

As an immediate consequence of this, the set of Kripke-valid $\texttt{GML}^\infty$-formulas coincides with the collection of coalgebraic validities.

COROLLARY 2.18. *For every $\texttt{GML}^\infty$-formula φ:*

$$\varphi \text{ is valid iff } \varphi \text{ is valid in every } \mathsf{B}^\infty\text{-coalgebra.} \tag{1}$$

At a more conceptual level, we note that the constructions and equivalences discussed above reveal that in some cases B^∞-coalgebras can serve as more compact, and perhaps even finite representations of Kripke structures. This raises the question whether any satisfiable $\texttt{GML}^\infty$-formula can perhaps be satisfied in a *finite* B^∞-coalgebra, where we observe that the analogous question for Kripke models makes no sense at all.

The result below states that indeed we can establish a finite model property for coalgebraic models; we can even recursively *bound* the complexity of such a finite satisfying coalgebra in terms of the satisfiable formula. As a corollary of this, we can show that the satisfiability problem for $\texttt{GML}^\infty$-formulas is *decidable*. To formulate this we need the following size measures of B^∞-coalgebras.

DEFINITION 2.19. *Let $\mathbb{S} = (S, \sigma, V)$ be a finite B^∞-coalgebra. We define its* size $|\mathbb{S}|$ *as the number of its states, $|\mathbb{S}| := |S|$, and its* weight $w(\mathbb{S})$ *as the maximum number reached as a finite value of any weight function*:

$$w(\mathbb{S}) := \max\big(\mathbb{N} \cap \{\sigma_s(t) \mid s, t \in S\}\big).$$

THEOREM 2.20 (Small Model Property). *Let φ be a $\texttt{GML}^\infty$-formula. If φ is satisfiable, then it is satisfiable in some B^∞-coalgebra of size at most $2^{|\varphi|}$ and weight at most $r(\varphi)$.*

For a *proof* of this theorem, we refer to Theorem 4.5. Observe that this result applies to $\texttt{ML}^\infty$ since it is a fragment of $\texttt{GML}^\infty$.

2.4. First-order logic with the infinity quantifier. In this section we introduce the predicate logic $\texttt{FO}^\infty$, the extension of first-order logic with the infinity quantifier $\exists^\infty$. We may, of course, consider the language of $\texttt{FO}^\infty$ over arbitrary signatures. In the context of this paper, however, we restrict attention to a signature that is tailored towards Kripke structures. That is, every proposition letter features as a unary predicate symbol, and in addition our signature has a binary predicate symbol R to denote the accessibility relation.

Having these restrictions in mind, we provide some basic syntactic and semantic definitions. Subsequently, we introduce some useful concepts that link the $\texttt{FO}^\infty$ language with the $\texttt{GML}^\infty$ modal language. We conclude this section by adapting the well-known Ehrenfeucht–Fraïssé game to account for the *infinity quantifier*.

DEFINITION 2.21. *Given a set P of proposition letters, and a disjoint set* $\mathtt{iVar}$ *of individual variables, we define the language* $\mathtt{FO}^{\infty}(P)$ *as follows*:

$$\varphi ::= p(x) \mid R(x, y) \mid x = y \mid \neg\varphi \mid (\varphi_0 \vee \varphi_1) \mid \exists x\, \varphi \mid \exists^{\infty} x\, \varphi,$$

where $p \in P$ and $x, y \in$ $\mathtt{iVar}$.

We use the Boolean connectives $\wedge, \to$ and $\leftrightarrow$ as the standard abbreviations, and define $\forall$ and $\forall^{\infty}$ as the Boolean duals of the quantifiers $\exists$ and $\exists^{\infty}$: $\forall x\, \varphi := \neg\exists x\, \neg\varphi$, and $\forall^{\infty} x\, \varphi := \neg\exists^{\infty} x\, \neg\varphi$. Furthermore, we extend various well-known syntactic notions, such as free and bound variables or quantifier depth, from first-order logic to $\mathtt{FO}^{\infty}$ in the obvious way.

The semantics of $\mathtt{FO}^{\infty}$ in Kripke models is defined as follows.

DEFINITION 2.22. *By induction on the complexity of an* $\mathtt{FO}^{\infty}$*-formula* φ, *we inductively define its* meaning *in a Kripke model* $\mathbb{S} = (S, R, V)$ *under an* assignment $g : \mathtt{iVar} \to S$, *by means of the following* truth relation $\models$:

$$
\begin{array}{lll}
\mathbb{S}, g \models p(x) & \textit{if} & g(x) \in V(p), \\
\mathbb{S}, g \models R(x, y) & \textit{if} & (g(x), g(y)) \in R, \\
\mathbb{S}, g \models x = y & \textit{if} & g(x) = g(y), \\
\mathbb{S}, g \models \neg\varphi & \textit{if} & \mathbb{S}, g \not\models \varphi, \\
\mathbb{S}, g \models (\varphi_0 \vee \varphi_1) & \textit{if} & \mathbb{S}, g \models \varphi_0 \textit{ or } \mathbb{S}, g \models \varphi_1, \\
\mathbb{S}, g \models \exists x\, \varphi & \textit{if} & \textit{there is an } s \in S \textit{ such that } \mathbb{S}, g[x \mapsto s] \models \varphi, \\
\mathbb{S}, g \models \exists^{\infty} x\, \varphi & \textit{if} & \textit{there are infinitely many } s \in S \textit{ such that } \mathbb{S}, g[x \mapsto s] \models \varphi.
\end{array}
$$

Here $g[x \mapsto s]$ *denotes the assignment that is as* g *except that it maps* x *to* s.

It is easy to see that the only variables on which the meaning of a formula φ depends, are the free ones of φ. We adopt the standard notational convention that, where $\overline{x} = x_0 \cdots x_{k-1}$, we write $\varphi(\overline{x})$ to indicate that the free variables of φ are among $x_0, \dots, x_{k-1}$. We will be particularly interested in formulas with a single free variable, since these can naturally be interpreted in *pointed* Kripke models, and may thus be compared with modal formulas.

DEFINITION 2.23. *Let* $\varphi(x) \in \mathtt{FO}^{\infty}$ *be some formula. Given a Kripke model* $\mathbb{S} = (S, R, V)$, *we write* $\mathbb{S} \models \varphi[s]$ *to denote that* $\mathbb{S}, g \models \varphi$, *for some/any assignment* g *such that* $g(x) = s$. *Where* ψ *is a* $\mathtt{GML}^{\infty}$*-formula, we call* φ *and* ψ equivalent, *notation*: $\varphi \equiv \psi$, *if we have that* $\mathbb{S} \models \varphi[s]$ *iff* $\mathbb{S}, s \Vdash \psi$, *for every pointed model* $(\mathbb{S}, s)$.

DEFINITION 2.24. *Let E be some equivalence relation on the class of pointed models. We call an* $\mathtt{FO}^{\infty}$ *formula* $\varphi(x)$ E-invariant *if we have* $\mathbb{S} \models \varphi[s]$ *iff* $\mathbb{S}' \models \varphi[s']$ *for any pair* $(\mathbb{S}, s), (\mathbb{S}', s')$ *of E-equivalent pointed Kripke models.*

We shall be particularly interested in the set of $\underline{\leftrightarrow}^{\infty,\#}$-*invariant* $\mathtt{FO}^{\infty}$-formulas. There is an obvious way to extend the standard translation from modal to first-order logic, to a translation that maps an arbitrary $\mathtt{GML}^{\infty}$-formula φ to an equivalent $\mathtt{FO}^{\infty}$-formula $\varphi'(x)$. This shows that $\mathtt{GML}^{\infty}$ corresponds to some fragment of $\mathtt{FO}^{\infty}$, and clearly this fragment consists of formulas that are all $\underline{\leftrightarrow}^{\infty,\#}$-invariant.

What is not obvious, however, is that this fragment in fact contains, up to equivalence, *all* $\mathtt{FO}^{\infty}$-formulas $\varphi(x)$ that are $\underline{\leftrightarrow}^{\infty,\#}$-invariant. This expressive completeness result is the content of the next bisimulation-invariance theorem.

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

THEOREM 2.25. *Let $\varphi(x)$ be an $\mathtt{FO}^{\infty}$ formula. Then φ is $\underline{\leftrightarrow}^{\infty,\#}$-invariant if and only if it is equivalent to some $\mathtt{GML}^{\infty}$-formula. Similarly, φ is $\underline{\leftrightarrow}^{\infty}$-invariant if and only if it is equivalent to some $\mathtt{ML}^{\infty}$-formula.*

As mentioned already, our proof of Theorem 2.25 will be game-theoretic in nature. For this purpose we adapt the well-known Ehrenfeucht–Fraïssé games for first-order logic, to account for the semantic behaviour of the infinity quantifier. This game can be seen as the adaptation to our specific setting of the Ehrenfeucht-style model comparison game for arbitrary monotone generalised quantifiers originating with Krawczyk and Krynicki [19], and discussed in more detail in Väänänen [28].

DEFINITION 2.26. *Let $\mathbb{S}_0 = (S_0, R_0, V_0)$ and $\mathbb{S}_1 = (S_1, R_1, V_1)$ be two Kripke models. Two k-tuples $\overline{u}_0 = u_{0,0}u_{0,1}\cdots u_{0,k-1}$ and $\overline{u}_1 = u_{1,0}u_{1,1}\cdots u_{1,k-1}$ are* locally isomorphic *if we have, for all i, j, that $u_{0,i} = u_{0,j}$ iff $u_{1,i} = u_{1,j}$, that $u_{0,i} \in V_0(p)$ iff $u_{1,i} \in V_1(p)$, for all $p \in P$, and that $R_0u_{0,i}u_{0,j}$ iff $R_1u_{1,i}u_{1,j}$. The concatenation of two tuples $\overline{u}$ and $\overline{v}$ will usually be denoted by simple juxtaposition (as in $\overline{u}\,\overline{v}$), but for clarity we may also use the symbol "$\cdot$" as an explicit concatenation operator (as in $\overline{u} \cdot \overline{v}$). A* configuration *over $\mathbb{S}_0$ and $\mathbb{S}_1$ is a pair $(\overline{u}_0, \overline{u}_1)$ consisting, for some k, of a tuple $\overline{u}_0 \in S_0^k$ and a tuple $\overline{u}_1 \in S_1^k$.*

DEFINITION 2.27. *Let $\mathbb{S}_0 = (S_0, R_0, V_0)$ and $\mathbb{S}_1 = (S_1, R_1, V_1)$ be two Kripke models; as a convention, when $\mathbb{S}_i$ is one of these two models, we write $\mathbb{S}_{-i}$ to denote the other. We define the game $\mathcal{E}_k^{\infty}(\mathbb{S}_0, \mathbb{S}_1)$ between two players, Spoiler (male) and Duplicator (female). A position in this game is a configuration over $\mathbb{S}_0$ and $\mathbb{S}_1$. The game starts at some initial configuration $(\overline{s}_0, \overline{s}_1)$ and consists of k rounds. A single round of the game consists of a short interaction between the two players, of the following kind.*

First Spoiler chooses to perform one of the following types of moves:

1. *First-order move; the round then consists of the following moves:*
   - *Spoiler chooses a structure $\mathbb{S}_i$ and an element $u_i \in S_i$;*
   - *Duplicator responds with an element $u_{-i} \in S_{-i}$.*
2. *Second-order move; the round now consists of the following moves:*
   - *Spoiler chooses a structure $\mathbb{S}_i$ and an infinite set $U_i \subseteq S_i$;*
   - *Duplicator responds with an infinite set $U_{-i} \subseteq S_{-i}$;*
   - *Spoiler chooses an element $u_{-i} \in U_{-i}$;*
   - *Duplicator responds with an element $u_i \in U_i$.*

The sequences $\overline{u}_0 \in S_0^n$ and $\overline{u}_1 \in S_1^n$ of elements chosen up to round n are then extended to $\overline{u}_0' := \overline{u}_0 \cdot u_0 \in S_0^{n+1}$ and $\overline{u}_1' := \overline{u}_1 \cdot u_1 \in S_1^{n+1}$. Duplicator wins a match of this game if, after the final k-th round, the resulting configuration $(\overline{s}_0\overline{u}_0, \overline{s}_1\overline{u}_1)$ consists of locally isomorphic tuples.

Note that the only items that are recorded in a play of this game are the *objects* picked by the players, not the subsets, and note as well that there are no constraints on the internal structure of the infinite sets chosen by the players. In particular, Duplicator does not have to establish an isomorphism between the sets picked by Spoiler and by herself.

In our context we will mainly be interested in the versions of this Ehrenfeucht–Fraïssé game that are played on *pointed* Kripke models.

DEFINITION 2.28. *Let $(\mathbb{S}, s)$ and $(\mathbb{S}', s')$ be two pointed Kripke models. We let $\mathcal{E}_k^{\infty}(\mathbb{S}, \mathbb{S}')@(s, s')$ denote the k-round Ehrenfeucht–Fraïssé game on $\mathbb{S}$ and $\mathbb{S}'$ that starts* *from the initial configuration* (s, s'). *In this context we define a* k-configuration *to be a pair* $(\overline{u}, \overline{u}')$ *of* $k+1$-*tuples such that* $u_0 = s$ *and* $u'_0 = s'$.

If Duplicator has a winning strategy in this initialized game, we write $\mathbb{S}, s \cong_k^{\infty} \mathbb{S}', s'$.

Note that since the Ehrenfeucht–Fraïssé game on pointed models has an initial configuration consisting of 1-tuples, the configuration reached after k moves of the game is indeed what we defined above as a k-configuration.

We conclude this subsection by mentioning the following adequacy result for $\mathtt{FO}^{\infty}$. For a proof we refer to Väänänen [28], where a more general result (Theorem 10.46) is stated and proved.

Proposition 2.29. *Let* $(\mathbb{S}, s)$ *and* $(\mathbb{S}', s')$ *be two pointed Kripke structures. Then the following are equivalent, for every* n:

(a) $\mathbb{S}, s \cong_n^{\infty} \mathbb{S}', s$;
(b) $\mathbb{S} \models \varphi[s]$ *iff* $\mathbb{S}' \models \varphi[s']$, *for every* $\varphi(x) \in \mathtt{FO}^{\infty}$ *of quantifier depth* $\leq n$.

§3. Bisimulation invariance.

3.1. Introduction. In this section we prove our main expressiveness result, Theorem 2.25, which characterises $\mathtt{GML}^{\infty}$ as the $\underline{\leftrightarrow}^{\infty,\#}$-invariant fragment of $\mathtt{FO}^{\infty}$, and $\mathtt{ML}^{\infty}$ as its $\underline{\leftrightarrow}^{\infty}$-fragment. We will focus on the case of $\mathtt{GML}^{\infty}$, the proof details for the case of $\mathtt{ML}^{\infty}$ can be found in [2].

Bisimulation invariance theorems are model-theoretic characterisation theorems, i.e., results that link certain semantic properties to the syntactic shape of formulas. This line of research started with van Benthem's seminal result, which identifies basic modal logic as the bisimulation invariant fragment of first-order logic [3]. This result has been extended in many directions; for instance, de Rijke was the first to prove a similar result for graded modal logic [24]. Van Benthem's and de Rijke's proofs are rooted in classical model theory, with a prominent role for the notion of compactness.

Our proof will be *game-theoretic* in nature, following the proof strategy that was initiated by Rosen [25], who gave an alternative proof of van Benthem's theorem which does not rely on compactness of first-order logic (and which transfers as well to the setting of *finite* model theory). Otto [23], applying the game-theoretic method to graded modal logic, recently obtained a result that applies in the settings of both general and finite models. Our approach can be seen as an extension of Otto's approach to include the infinity quantifier (with the note that for obvious reasons it does not make sense to study the logics of our interest in the setting of finite models). For more information on the game-theoretic approach we refer to Goranko and Otto's handbook article [16].

Before diving into the proof details, we sketch an overview of the proof. First of all, we will stratify the language $\mathtt{GML}^{\infty}$ into layers:

$$\mathtt{GML}^{\infty} = \bigcup_{q,n} \mathtt{GML}^{\infty}_{q,n},$$

where $\mathtt{GML}^{\infty}_{q,n}$ consists of those formula in $\mathtt{GML}^{\infty}$ of depth at most q and rank at most n. Each fragment $\mathtt{GML}^{\infty}_{q,n}$ comes with an associated bounded notion of bisimilarity, $\underline{\leftrightarrow}^{\infty,\#}_{q,n}$, and the point of this stratification lies in the fact that each set $\mathtt{GML}^{\infty}_{q,n}$ is *finite* up to logical equivalence. This means that for each equivalence class C of $\underline{\leftrightarrow}^{\infty,\#}_{q,n}$ there

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

is a characteristic formula $\chi_C \in \mathtt{GML}^{\infty}_{q,n}$ such that any pointed model satisfies χ_C iff it belongs to C.

The key observation in the proof of Theorem 2.25 is then that

$$\text{if } \varphi(x) \in \mathtt{FO}^{\infty} \text{ is bisimulation invariant, then } \varphi(x) \text{ is } \underline{\leftrightarrow}^{\infty,\#}_{n,n}\text{-invariant,} \tag{2}$$

where $n = 3^{qd(\varphi)}$, $qd(\varphi)$ being the quantifier depth of φ. As we will see, it easily follows from (2) that any bisimulation-invariant formula $\varphi(x)$ is equivalent to the $\mathtt{GML}^{\infty}_{n,n}$-formula $\bigvee \chi_C$, where we take the disjunction over all cells C of $\underline{\leftrightarrow}^{\infty,\#}_{n,n}$ that contain a pointed model satisfying φ.

For the proof of (2) we develop some *model theory* for $\mathtt{FO}^{\infty}$ and $\mathtt{GML}^{\infty}$. In particular, we will first show that every bisimulation-invariant $\mathtt{FO}^{\infty}$-formula $\varphi(x)$ is *local*, in the sense that its truth only depends on a certain *neighbourhood* of the point referred to by x. Our second model-theoretic result states that for *tree models*, we may "upgrade" the relation $\underline{\leftrightarrow}^{\infty,\#}_{n,n}$ to the game-based indistinguishability relation $\cong^{\infty}_{n}$ of Definition 2.28. It then follows from the adequacy of Ehrenfeucht–Fraïssé games (Proposition 2.29) that a bisimulation-invariant $\mathtt{FO}^{\infty}$-formula $\varphi(x)$ of quantifier depth n is $\underline{\leftrightarrow}^{\infty,\#}_{n,n}$-invariant indeed.

In the final subsection we will then see how all these observations can be assembled in order to prove Theorem 2.25.

3.2. A stratification of the language. As mentioned in the introduction to this section, we will start with stratifying the language $\mathtt{GML}^{\infty}$ into layers.

DEFINITION 3.1. *We let* $\mathtt{GML}^{\infty}_{q,n}$ *consist of those formula in* $\mathtt{GML}^{\infty}$ *that have depth at most* q *and rank at most* n.

As was observed by Otto [23] for graded modal logic (i.e., without the infinity diamond), in order to stratify the language $\mathtt{GML}^{\infty}$ into layers that are finite up to logical equivalence, it does not suffice to confine attention to the modal depth of formulas: one needs to take their rank into consideration as well. We leave the (straightforward) verification of the following proposition as an exercise to the reader.

PROPOSITION 3.2. *For each* q, n *the set* $\mathtt{GML}^{\infty}_{q,n}$ *is finite, up to logical equivalence.*

We will adapt the bisimilarity game of Definition 2.11 to the setting of these fragments, by restricting the *size* of the sets that Spoiler can pick, as well as the *length* of the game.

DEFINITION 3.3. *The* $\underline{\leftrightarrow}^{\infty,\#}_{q,n}$-bisimilarity game, *with* $q, n \in \mathbb{N}$, *is the version of the bisimilarity game of Definition 2.11 where (i) Spoiler can only pick subsets that are either infinite or of size* $\leq n$, *and (ii) instead of (potentially) lasting infinitely many rounds, the game now finishes after* q *many rounds.*

As in our previous definitions, two pointed models $(\mathbb{S}, s), (\mathbb{S}', s')$ *are* $\underline{\leftrightarrow}^{\infty,\#}_{q,n}$*-bisimilar, denoted by* $\mathbb{S}, s \,\underline{\leftrightarrow}^{\infty,\#}_{q,n}\, \mathbb{S}', s'$, *if Duplicator has a winning strategy in this version of the game, initialised at position* (s, s').

The following result highlights the close relation between the set of $\mathtt{GML}^{\infty}_{q,n}$-formulas and the bisimilarity equivalence relation $\underline{\leftrightarrow}^{\infty,\#}_{q,n}$.

THEOREM 3.4. *Fix two natural numbers* $q, n \in \mathbb{N}$. *For any pointed Kripke model* $(\mathbb{S}, s)$ *there is a* $\mathtt{GML}^{\infty}_{q,n}$*-formula* $\chi^{\mathbb{S},s}_{q,n}$ *such that the following are equivalent, for any pointed Kripke model* $(\mathbb{S}', s')$:

1. $(\mathbb{S}', s') \leftrightarrows^{\infty,\#}_{q,n} (\mathbb{S}, s)$,
2. $(\mathbb{S}', s') \equiv^{\#,\infty}_{q,n} (\mathbb{S}, s)$,
3. $(\mathbb{S}', s') \Vdash \chi^{\mathbb{S},s}_{q,n}$.

We skip the proof of this result, which proceeds via a routine induction on the modal depth n.

Note that in particular, Theorem 3.4 states that every $\leftrightarrows^{\infty,\#}_{q,n}$-equivalence class of pointed Kripke models is characterised by a single $\mathtt{GML}^{\infty}_{q,n}$ formula. Furthermore, by our earlier observation that the set $\equiv^{\#,\infty}_{q,n}$ is finite up to logical equivalence, we find that the equivalence relation $\leftrightarrows^{\infty,\#}_{q,n}$ has finite *index*. As a corollary of these two observations, we obtain the following result, which we will need in the proof of our bisimulation invariance theorem. Given an equivalence relation E on the class of pointed models, we call a class of pointed models KE-*saturated* if it is a union of E-cells, or, equivalently: if we have $(\mathbb{S}, s) \in K$ iff $(\mathbb{S}', s') \in K$, whenever $(\mathbb{S}, s)$ and $(\mathbb{S}', s')$ are equivalent, according to E.

Proposition 3.5. *Let q and n be two natural numbers, and let K be a $\leftrightarrows^{\infty,\#}_{q,n}$-saturated class of pointed models. Then there is a formula $\chi_K \in \mathtt{GML}^{\infty}_{q,n}$ such that $\mathbb{S}, s \Vdash \chi_K$ iff $(\mathbb{S}, s)$ belongs to K, for any pointed model $(\mathbb{S}, s)$.*

Proof. Given a cell C of $\leftrightarrows^{\infty,\#}_{q,n}$, we fix some formula χ_C in $\mathtt{GML}^{\infty}_{q,n}$, as given by Theorem 3.4; that is, χ_C characterises the pointed models in C.

Let q, n and K be as in the formulation of the proposition. Define $\chi_K := \bigvee \chi_C$, where we take the disjunction over all cells C of $\leftrightarrows^{\infty,\#}_{q,n}$ that contain a pointed model in K. Observe that the formula χ_K is well formed since the relation $\leftrightarrows^{\infty,\#}_{q,n}$ has finite index. We claim that χ_K characterises K, in the sense that, for any pointed model $(\mathbb{S}, s)$:

$$(\mathbb{S}, s) \text{ belongs to } K \text{ iff } \mathbb{S}, s \Vdash \chi_K.$$

To see this, fix a pointed model $(\mathbb{S}, s)$, and let C be its equivalence class under the relation $\leftrightarrows^{\infty,\#}_{q,n}$. It follows by Theorem 3.4 that $\mathbb{S}, s \Vdash \chi_C$.

If $(\mathbb{S}, s)$ belongs to K then by definition of χ_K the formula χ_C will be one of the disjuncts of χ_K, so that we find $\mathbb{S}, s \Vdash \chi_K$ as required. Conversely, if $\mathbb{S}, s \Vdash \chi_K$ then by definition of χ_K we have that $\mathbb{S}, s \Vdash \chi_{C'}$ for some $\leftrightarrows^{\infty,\#}_{q,n}$-cell C' which contains a pointed model $(\mathbb{S}', s')$ in K. It follows from $\mathbb{S}, s \Vdash \chi_{C'}$ that $C' = C$, and so both $(\mathbb{S}, s)$ and $(\mathbb{S}', s')$ belong to this class. This means that $\mathbb{S}, s \leftrightarrows^{\infty,\#}_{q,n} \mathbb{S}', s'$, and from this we find $(\mathbb{S}, s)$ in K by the assumption that K is $\leftrightarrows^{\infty,\#}_{q,n}$-saturated. $\square$

3.3. Locality for bisimulation-invariant $\mathtt{FO}^{\infty}$ ***formulas.*** One of the key concepts that Rosen uses in his proof is that of *locality*: the property that the meaning of a first-order formula $\varphi(x)$ in a structure $\mathbb{S}$ only depends on a small neighbourhood of (the point associated with) x. Our goal here will be to show that this property extends to formulas involving the infinity quantifier, at least to the ones that are bisimulation invariant.

The notion of locality is defined in terms of *Gaifman distance*.

Definition 3.6. *Let $\mathbb{S} = (S, R, V)$ be a Kripke model. Where $E \subseteq S \times S$ is the union of R and its converse, we define the* Gaifman distance *$gaif(s, t)$ between two points s and t to be $n \in \mathbb{N}$ if n is the length of the shortest E-path linking s to t, and we set $gaif(s, t) := \infty$ if there is no such E-path. Given a state $s \in S$ and a natural number n, the n-*neighbourhood *$N^{\mathbb{S}}(s, n)$ of s consists of those states t in $\mathbb{S}$ such that $gaif(s, t) \leq n$; we let $\mathbb{B}^{\mathbb{S}}(s, n)$ denote the induced substructure on $N^{\mathbb{S}}(s, n)$.*

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

An $\mathtt{FO}^\infty$ *formula* φ *is* n-local *if we have* $\mathbb{S} \vDash \varphi(s)$ *iff* $\mathbb{S}' \vDash \varphi(s')$, *for any two pointed Kripke models* $(\mathbb{S}, s)$ *and* $(\mathbb{S}', s')$ *such that* $\mathbb{B}^{\mathbb{S}}(s, n)$ *and* $\mathbb{B}^{\mathbb{S}'}(s', n)$ *are isomorphic.*

We will now prove that every bisimulation-invariant $\mathtt{FO}^\infty$-formula $\varphi(x)$ is 3^q-local, where q is the quantifier depth of φ. In fact, we will prove a slightly stronger statement (as in the case of the bisimulation invariance results for first-order logic): it turns out that the only invariance property that we need to prove locality concerns *disjoint unions*. We assume that the reader is familiar with this model-theoretic construction (the precise definition can be found in [5]). It is easy to see that a pointed model is $\underline{\leftrightarrow}^{\infty,\#}$-bisimilar to its representation inside a disjoint union; from this it follows that bisimulation invariance implies invariance under disjoint unions.

PROPOSITION 3.7. *Let* $\varphi(x) \in \mathtt{FO}^\infty$ *be invariant under taking disjoint unions. Then* φ *is* 3^q*-local, where* $q = qd(\varphi)$.

Proof. Fix a formula $\varphi(x) \in \mathtt{FO}^\infty$ of quantifier depth q, and assume that φ is invariant under taking disjoint unions. Define $n := 3^q$, and let $(\mathbb{S}, s)$ be an arbitrary pointed model. In order to prove the proposition, by invariance under taking disjoint unions it obviously suffices to prove that

$$\mathbb{S} \uplus \mathbb{D} \models \varphi[s] \text{ iff } \mathbb{B} \uplus \mathbb{D} \models \varphi[s], \tag{3}$$

where $\mathbb{B} := \mathbb{B}^{\mathbb{S}}(s, n)$ and $\mathbb{D}$ is the disjoint union of q many copies of $\mathbb{S}$ and q many copies of $\mathbb{B}$. As mentioned, our proof of (3) will be game-theoretic in nature. In view of Proposition 2.29 it suffices to show the following:

$$\text{Duplicator has a winning strategy in } \mathtt{EF}^\infty_q(\mathbb{S} \uplus \mathbb{D}, \mathbb{B} \uplus \mathbb{D})@(s, s). \tag{4}$$

To formulate our proof of (4), we need to introduce some further terminology and notation. To start with, we need some more detail in our notation for the models $\mathbb{S} \uplus \mathbb{D}$ and $\mathbb{B} \uplus \mathbb{D}$. For $i \in \{-q, \dots, q\}$ we let $\mathbb{S}_i$ be the isomorphic copy of $\mathbb{S}$ based on the set $S_i = S \times \{i\}$, where the projection map $\pi_i : S_i \to S$ is the witnessing isomorphism, and similarly for $\mathbb{B}_i$. We identify $\mathbb{S}_0$ with $\mathbb{S}$ and $\mathbb{B}_0$ with $\mathbb{B}$, and define

$$\mathbb{D} := \biguplus_{-q \leq i < 0} \mathbb{S}_i \ \uplus \ \biguplus_{0 < i \leq q} \mathbb{B}_i,$$

so that we may think of $\mathbb{S} \uplus \mathbb{D}$ as the structure $\biguplus_{-q \leq i \leq 0} \mathbb{S}_i \ \uplus \ \biguplus_{0 < i \leq q} \mathbb{B}_i$, and of $\mathbb{B} \uplus \mathbb{D}$ as the structure $\biguplus_{-q \leq i < 0} \mathbb{S}_i \ \uplus \ \biguplus_{0 \leq i \leq q} \mathbb{B}_i$. We will refer to a point u in $\mathbb{S} \uplus \mathbb{D}$ as being $\mathbb{K}$-*type*, with $\mathbb{K} \in \{\mathbb{S}, \mathbb{B}\}$, if u belongs to $\mathbb{K}_i$ for some $i \in \{-q, ..q\}$.

Given a k-configuration $(\overline{u}, \overline{u}')$ arising during a match of the game, we let $N_k(\overline{u})$ and $N_k(\overline{u}')$ be the sets

$$N_k(\overline{u}) := \bigcup_i N^{\mathbb{S} \uplus \mathbb{D}}(u_i, 3^{q-k}), \quad N'_k(\overline{u}') := \bigcup_i N^{\mathbb{B} \uplus \mathbb{D}}(u'_i, 3^{q-k}),$$

respectively. In other words, $N_k(\overline{u})$ is the union of the 3^{q-k}-neighbourhoods of the different worlds in $\overline{u}$. In particular, every point in $N_k(\overline{u})$ is, at most, at a distance 3^{q-k} from one of the worlds in $\overline{u}$. Analogous observations apply to $N'_k(\overline{u}')$.

The key observation in the proof of (4) is that Duplicator can maintain the condition that, where $(\overline{u}, \overline{u}')$ is any configuration reached during a play of $\mathtt{EF}^\infty_q(\mathbb{S} \uplus \mathbb{D}, \mathbb{B} \uplus \mathbb{D})@(s, s)$,

(†) there is a local isomorphism $f : N_k(\overline{u}) \to N'_k(\overline{u}')$ such that $f(u_i) = u'_i$, for all i.

It should be clear that this condition holds at the start of the game, where $k = 0$.

Now assume that, inductively, for some $k < q$ the condition (†) holds for some k-configuration $(\overline{u}, \overline{u}')$, witnessed by some local isomorphism f. We need to prove that to any challenge of Spoiler, Duplicator has a response that guarantees the resulting $k+1$-configuration to again satisfy (†). We show this by distinguishing cases, as to the kind of move Spoiler makes.

Case 1. First assume that Spoiler makes a first-order move. In this case we proceed very much in the same way as Goranko and Otto in [16]. Without loss of generality we may assume that he picks an element, say, u, of the structure $\mathbb{S} \uplus \mathbb{D}$; then we need to identify a point u' in $\mathbb{B} \uplus \mathbb{D}$ as a suitable reply for Duplicator. We make a further case distinction as to whether u is close to one of the points in the existing tuple $\overline{u}$ or not.

Case 1a. $\mathit{gaif}(u, u_i) \leq 2 \cdot 3^{q-(k+1)}$ for some $i \in \{0, \dots, k\}$. The crucial observation in this case is that we have $N_{k+1}(\overline{u} \cdot u) \subseteq N_k(\overline{u})$; the verification of this statement is a simple exercise in Gaifman metric. We then simply define $u' := f(u)$, and it is straightforward to check that f restricts to a local isomorphism from $N_{k+1}(\overline{u} \cdot u)$ to $N'_{k+1}(\overline{u}' \cdot u')$. Since we have $f(u_i) = u'_i$ by assumption and $f(u) = u'$ by definition, this shows that indeed Duplicator can keep the condition (†) for one more round.

Case 1b. $\mathit{gaif}(u, u_i) > 2 \cdot 3^{q-(k+1)}$ for all $i \in \{0, \dots, k\}$. The crucial statement is now that the set $N^{\mathbb{S}\uplus\mathbb{D}}(u, 3^{q-(k+1)})$ does not overlap with any set $N^{\mathbb{S}\uplus\mathbb{D}}(u_i, 3^{q-(k+1)})$, so that $N_{k+1}(\overline{u} \cdot u)$ is the *disjoint* union of the sets $\bigcup_i N^{\mathbb{S}\uplus\mathbb{D}}(u_i, 3^{q-(k+1)})$ and $N^{\mathbb{S}\uplus\mathbb{D}}(u, 3^{q-(k+1)})$. We will assume that u is an $\mathbb{S}$-type point of $\mathbb{S} \uplus \mathbb{D}$ (the case where u is $\mathbb{B}$-type is dealt with analogously). That is, u belongs to $\mathbb{S}_j$ for some j, and we may write $u = (\pi_j(u), j)$ with $\pi_j(u)$ in $\mathbb{S}$.

Since $\overline{u}'$ comprises at most $k+1$ elements, where $k < q$, $\mathbb{D}$ contains q distinct copies of $\mathbb{S}$, and the initial state u'_0 does not belong to $\mathbb{D}$ (but to $\mathbb{B}$), the target structure $\mathbb{B} \uplus \mathbb{D}$ will have a "fresh" copy of $\mathbb{S}$, so to speak. That is, we may consider some r such that $\mathbb{S}_r$ contains none of the u'_i. We now define $u' := (\pi_j(u), r)$; in other words, we use the obvious isomorphism between $\mathbb{S}_j$ and $\mathbb{S}_r$ to find u'. Now consider the following map f' on $N_{k+1}(\overline{u} \cdot u)$:

$$f'(s) := \begin{cases} f(s), & \text{if } s \in \bigcup_i N^{\mathbb{S}\uplus\mathbb{D}}(u_i, 3^{q-(k+1)}), \\ (\pi_j(s), r), & \text{if } s \in N^{\mathbb{S}\uplus\mathbb{D}}(u, 3^{q-(k+1)}). \end{cases} \tag{5}$$

It is then a routine exercise to verify that f' is well defined, and that it provides the required local isomorphism between $N_{k+1}(\overline{u} \cdot u)$ and $N'_{k+1}(\overline{u}' \cdot u')$. Since again we have $f'(u_i) = u'_i$ by assumption and $f'(u) = u'$ by definition, it follows that in this case as well, Duplicator can keep the condition (†) for one more round.

Case 2. Now assume that Spoiler makes a second-order move. Without loss of generality we may assume that he picks an infinite subset, say, U, of the structure $\mathbb{S} \uplus \mathbb{D}$. Similar to Case 1, we make a further case distinction, but now

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

as to whether *infinitely* many elements of U are close to one of the points in the existing tuple $\overline{u}$ or not.

Case 2a. There are infinitely many $u \in U$ such that $gaif(u, u_i) \leq 2 \cdot 3^{q-(k+1)}$ for some $i \in \{0, \dots, k\}$.

Let $X \subseteq U$ be the set of all these elements, then the key observation—as in the first-order move case—is that for each $u \in X$ we have $N_{k+1}(\overline{u} \cdot u) \subseteq N_k(\overline{u})$. Let

$$U' := f[X]$$

be Duplicator's response to Spoiler's move U. Since f is a bijection between $N_k(\overline{u})$ and $N'_k(\overline{u}')$, we immediately see that U' is an infinite subset of $N'_k(\overline{u}')$. Now suppose that, continuing the match, Spoiler picks an element $u' \in U'$; then by definition there is a $u \in X$ such that $u' = f(u)$, and Duplicator will pick this u as her response to Spoiler's move u'. As in Case 1 above, it is then straightforward to verify that the restriction of f to the set $N_{k+1}(\overline{u} \cdot u)$ witnesses that Duplicator has maintained (†) for one more round.

Case 2b. There are infinitely many $u \in U$ such that $gaif(u, u_i) > 2 \cdot 3^{q-(k+1)}$ for all $i \in \{0, \dots, k\}$.

Let $X \subseteq U$ be the set of all these elements, then the key observation is, as in Case 1b, that for each $u \in X$ we have

$$N_{k+1}(\overline{u} \cdot u) \;=\; \bigcup_i N^{\mathbb{S} \uplus \mathbb{D}}(u_i, 3^{q-(k+1)}) \;\uplus\; N^{\mathbb{S} \uplus \mathbb{D}}(u, 3^{q-(k+1)}).$$

Furthermore, by the pigeonhole principle, X must contain either infinitely many $\mathbb{S}$-type points or infinitely many $\mathbb{B}$-type points. We assume the first (again, the other case is very similar). For $j \in \{-q, \dots, 0\}$, let X_j be the set of points in X that belong to $\mathbb{S}_j$; then, once more by the pigeonhole principle, there is an index j such that X_j is infinite. As in Case 1b we may then take some r such that $\mathbb{S}_r$ contains none of the u'_i. We are now ready to define Duplicator's response to Spoiler's move U:

$$U' := \{(\pi_j(u), r) \mid u \in X_j\}.$$

That is, again we use the obvious isomorphism between $\mathbb{S}_j$ and $\mathbb{S}_r$ to find U' in $\mathbb{S}_r$ as the counterpart of X_j in $\mathbb{S}_j$. Clearly then, U' is infinite, so that it constitutes a legitimate move for Duplicator. Now assume that Spoiler continues the match by picking an element $u' \in U$. Then by definition u' is of the form $u' = (\pi_j(u), r)$ for some $u \in X_j$, and since $X_j \subseteq U$, Duplicator may pick this element u as her response to Spoiler.

As discussed in Case 1b, it is now straightforward to verify that the map f' as defined in (5), is a local isomorphism from $N_{k+1}(\overline{u} \cdot u)$ to $N'_{k+1}(\overline{u}' \cdot u')$ that witnesses the fact that Duplicator has carried the condition (†) one round further.

This finishes the proof of (4), and, as discussed, this suffices to prove the proposition. □

3.4. Upgrading. As we saw in the introduction to this section, the second result on which the proof of the key observation (2) rests, is that for *tree models*, the relation $\underline{\leftrightarrow}^{\infty,\#}_{n,n}$

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

may be “upgraded” to the game-based indistinguishability relation $\cong_n^\infty$ of Definition 2.28. This observation is the content of Theorem 3.9, which we will state and prove in this subsection.

Recall that a *tree* is a structure (T, R) such that, for some fixed node r, every node t is reachable from r by a unique finite path. The node r, which is in fact uniquely determined by this property, is called the *root* of the tree. We will call a pointed Kripke model $(\mathbb{T}, r)$ a *tree model* if $\mathbb{T}$ is based on a tree with root r.

To formulate and prove Theorem 3.9 we need to introduce some concepts related to tree models.

DEFINITION 3.8. *Let $\mathbb{T}$ be a tree model with root r. We use the terms* successor/child *and* predecessor/parent *in their usual meaning, and we let the ancestor (descendant) relation be the reflexive-transitive closure of the predecessor (successor) relation. We use $\downarrow(u)$ to denote the set of ancestors of u.*

The height $h(\mathbb{T})$ *of $\mathbb{T}$ is defined as the supremum of the heights of its nodes, where the height $h(u)$ of a node u is defined as its distance to r. Given a tree of finite height, we define the* depth $d(u)$ *of a node u as the maximal distance of u to any descendant which is a leaf of $\mathbb{T}$.*

Finally, given two nodes u and v in $\mathbb{T}$, we let $lca(u, v)$ denote the last common ancestor *of u and v.*

Observe that the root r of a tree of height l is an ancestor of every node u and satisfies $h(r) = 0$ and $d(r) = l$, and that we have $h(u) + d(u) \leq l$ for any node u. The ancestor relation R^* is a partial order on T, which restricts to a linear order on any set of the form $\downarrow(u)$. The least common ancestor map obviously satisfies the conditions $lca(u, u) = u$, and $lca(u, v) = lca(v, u)$, for all nodes u, v. It is also easily verified that $u \in \downarrow(v)$ iff $lca(u, v) = u$. Finally, if $lca(u, v_0)$ and $lca(u, v_1)$ have the same height, then they must be identical.

THEOREM 3.9. *Let $(\mathbb{T}, r)$ and $(\mathbb{T}', r')$ be two tree models of height q. Then $\mathbb{T}, r \,\underline{\leftrightarrow}_{q,q}^{\infty,\#}\, \mathbb{T}', r'$ implies $\mathbb{T}, r \cong_q^\infty \mathbb{T}', r'$.*

Clearly, the goal of our proof for this theorem will be to supply Duplicator with a winning strategy in the Ehrenfeucht–Fraïssé game of length q, between two $\underline{\leftrightarrow}_{q,q}^{\infty,\#}$-equivalent tree models of height q. In the next definition we introduce the notion of q-companionship, which will be a key notion with which Duplicator will link the nodes in the two respective tree models.

DEFINITION 3.10. *Let $\mathbb{T}$ and $\mathbb{T}'$ be tree models with roots r, r', respectively, and let $q \geq h(\mathbb{T})$. We say that two worlds* t *in $\mathbb{T}$ and t' in $\mathbb{T}'$ are q-companions, if $h(t) = h(t')$ and for all $u \in \downarrow(t)$ and $u' \in \downarrow(t')$ with $h(u) = h(u')$ we have $\mathbb{T}, u \,\underline{\leftrightarrow}_{q-h(u),q}^{\infty,\#}\, \mathbb{T}', u'$.*

Here are some basic observations about the notion. First note that nodes can only be q-companions if they have the same height and their respective ancestors are, one by one, also q-companions. In fact, if u and u' are q-companions, then there is a bijection $f : \downarrow(u) \to \downarrow(u')$ mapping every predecessor v of u to a q-companion $f(v)$ of u' of the same height. The bisimilarity condition on these ancestors becomes stronger for nodes that are closer to the respective roots, to the effect that companions of height q (which must be leaves of the tree) are only required to satisfy the same proposition letters, while the two roots are q-companions iff they are $\underline{\leftrightarrow}_{q,q}^{\infty,\#}$-related. Furthermore, the following observation will turn out to be quite useful.

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

Lemma 3.11. *For any $q > 0$, when restricted to pairs $(\mathbb{T}, t)$ where $\mathbb{T}$ is a tree of height at most q, the relation of being q-companions is an equivalence relation of finite index.*

Proof. It is almost immediate from the definition that the relation of being q-companions is an equivalence relation. A routine inductive argument shows that this equivalence relation has finite index if we restrict to nodes of arbitrary but fixed height. From this the proposition is immediate. □

The following lemma states a very useful back-and-forth property for the q-companionship relation.

Lemma 3.12. *Let $\mathbb{T}$ and $\mathbb{T}'$ be two rooted tree-like structures of height q, and let t in $\mathbb{T}$ and t' in $\mathbb{T}'$ be q-companions. Then the following hold.*

1. *Every descendant v of t has a q-companion v' which is a descendant of t'.*
2. *If t has infinitely many descendants, all belonging to the same q-companion class, then t' has infinitely many descendants that all belong to this class as well.*

Proof. Let $\mathbb{T}, \mathbb{T}', t$ and t' be as in the formulation of the lemma.

For the proof of item 1 we first consider the special case where v is a successor of t. By Theorem 3.4 there is a formula $\chi \in \mathtt{GML}^{\infty}_{q-h(v),q}$ characterising the $\underline{\leftrightarrow}^{\infty,\#}_{q-h(v),q}$-cell of $(\mathbb{T}, v)$. It follows from $\mathbb{T}, v \Vdash \chi$ that $\mathbb{T}, t \Vdash \Diamond\chi$. Observe that the formula $\Diamond\chi$ belongs to the set $\mathtt{GML}^{\infty}_{(q-h(v))+1,q} = \mathtt{GML}^{\infty}_{q-h(t),q}$, where the equality holds as $h(t) = h(v) - 1$. By the q-companionship of t and t' we obtain that $\mathbb{T}', t' \Vdash \Diamond\chi$, so that t' must have a successor v' such that $\mathbb{T}', v' \Vdash \chi$. It is then immediate by Theorem 3.4 that $\mathbb{T}, v \,\underline{\leftrightarrow}^{\infty,\#}_{q-h(v),q}\, \mathbb{T}', v'$, and since t and t' are q-companions it easily follows that v and v' must be q-companions as well.

The general case, where v is an arbitrary (not necessarily one-step) descendant of t, is then easily derived from the special case via a straightforward induction on the distance from t to v.

For the proof of item 2, let V be an infinite set of nodes in $\mathbb{T}$ that all belong to the same q-companion class. Clearly then all nodes in V have the same height, a number that we will denote as $h(V)$. Furthermore, since all nodes in V are q-companions of one another, there is a formula $\chi \in \mathtt{GML}^{\infty}_{q-h(V),q}$ characterising the $\underline{\leftrightarrow}^{\infty,\#}_{q-h(V),q}$-cell of the nodes in V. This means in particular that $\mathbb{T}, v \Vdash \chi$ for all $v \in V$.

Now consider an arbitrary common ancestor u of all nodes in V; it is easy to see that u has the same distance to each node in V, so that it makes sense to talk about the distance $d(u, V)$ of u to the set V. We will prove, by induction on this distance, that u has a q-companion u' which has infinitely many descendants that all belong to the q-companion class of V. Clearly this suffices to prove item 2.

In the base step of this induction we assume that $d(u, V) = 1$; that is, $V \subseteq R(u)$. Since V is infinite this means that $\mathbb{T}, u \Vdash \Diamond^{\infty}\chi$. Now observe that as χ belongs to the set $\mathtt{GML}^{\infty,q}_{q-h(V),q}$, we find that $\Diamond^{\infty}\chi \in \mathtt{GML}^{\infty}_{(q-h(V))+1,q} = \mathtt{GML}^{\infty}_{q-h(u),q}$. From this and the assumption on u and u' it follows that $\mathbb{T}', u' \Vdash \Diamond^{\infty}\chi$. In other words, there is an infinite set V' of successors of u', such that each $v' \in V'$ satisfies the characteristic formula χ. By Theorem 3.4 this means that $\mathbb{T}, v \,\underline{\leftrightarrow}^{\infty,\#}_{q-h(v),q}\, \mathbb{T}', v'$, for each $v' \in V'$ and $v \in V$, and since the predecessor u' of any such v' is a q-companion of the predecessor u of any such v, we find that every $v' \in V'$ is a q-companion of every $v \in V$. But then we are done, since V' is an infinite set of descendants of u'.

For the inductive step, assume that $d(u, V) = k + 1$. We make a case distinction.

First assume that u has a successor w such that $|V \cap R^k(w)|$ is infinite. By (the proof of) item 1, u' has a successor w' which is a q-companion of w. Clearly we have $d(w, V) = d(u, V) - 1$, so that we may apply the inductive hypothesis to w. This yields an infinite set V' of descendants of w', all belonging to the same q-companion class as the nodes in V. It is not hard to see that this set V' meets the required conditions with respect to u and V as well.

If, on the other hand, u has no successor w such that $|V \cap R^k(w)|$ is infinite, then it must have an infinite set W of successors, each of which has at least one descendant in V. Without loss of generality we may assume that all nodes in W are q-companions of one another (if not, by Lemma 3.11 W would have an infinite subset with this property, and we could continue with this subset). This means in particular that the nodes in W belong to the same cell of the equivalence relation $\underline{\leftrightarrow}^{\infty,\#}_{q-(h(u)+1),q}$. Writing $\psi \in \mathtt{GML}^{\infty}_{q-(h(u)+1),q}$ for the characteristic formula of this equivalence class, we clearly have that u satisfies the formula $\Diamond^{\infty}\psi$. But since this formula belongs to the set $\mathtt{GML}^{\infty}_{q-h(u),q}$, it must hold at the q-companion u' of u as well, i.e., there is an infinite set W' of successors of u', all satisfying the formula ψ. As before we may derive from this that every $w' \in W'$ is a q-companion of every $w \in W$. But then by item 1 every $w' \in W'$ has a descendant $v'_{w'}$ that is a q-companion of some descendant in V of some successor w of u. From this we easily derive that the set $V' := \{v'_{w'} \mid w' \in W'\}$ has the required properties—note that the set is infinite since W' is infinite and $\mathbb{T}'$ is a tree model. $\square$

Where the q-companion relation concerns pairs of single nodes and their respective ancestors, to provide Duplicator with a winning strategy in the Ehrenfeucht–Fraïssé game we will compare *tuples* of (not necessarily related) points, and the subtrees that they induce.

Definition 3.13. *Let $\mathbb{T} = (T, R)$ be a tree, and let, for $k \geq 1$, the k-tuple $\overline{u} = (u_0, u_1, \dots, u_{k-1})$ be such that u_0 is the root of $\mathbb{T}$. Then we define $T_{\overline{u}} := \bigcup_{0 \leq i < k} {\downarrow}(u_i)$ and we let $\mathbb{T}_{\overline{u}}$ denote the subtree of $\mathbb{T}$ based on the set $T_{\overline{u}}$; we call $\mathbb{T}_{\overline{u}}$ the subtree of $\mathbb{T}$ that is* induced *by $\overline{u}$.*

Let $\mathbb{T}$ and $\mathbb{T}'$ be two tree models and let $\overline{u}$ and $\overline{u}'$ be two k-tuples of nodes in, respectively, $\mathbb{T}$ and $\mathbb{T}'$, such that $k \geq 1$ and u_0 and u'_0 are, respectively, the roots of $\mathbb{T}$ and $\mathbb{T}'$. We say that $\overline{u}$ and $\overline{u}'$ are q-similar, *notation*: $\mathbb{T}, \overline{u} \approx^{\downarrow}_q \mathbb{T}', \overline{u}'$, *if there is an isomorphism f from $\mathbb{T}_{\overline{u}}$ to $\mathbb{T}'_{\overline{u}'}$ such that $f(u_i) = u'_i$, for all i, and x and $f(x)$ are q-companions, for all x in $\mathbb{T}_{\overline{u}}$.*

The main observation in the proof of Theorem 3.9 is that Duplicator can make sure that any configuration arising during a play of the game will consist of two q-similar tuples, provided, of course, that the starting configuration meets this condition. The following lemma is the key result here; it states that Duplicator can maintain the q-similarity condition on configurations for one single round, systematically defending herself from any attack by Spoiler. Recall that we write $X \subseteq_{\omega} Y$ if X is a finite subset of Y.

Lemma 3.14. *Let $\mathbb{T}, \mathbb{T}'$ be two rooted tree-like structures of height q, and let, for some $0 \leq k < q$, $(\overline{u}, \overline{u}')$ be a k-configuration such that $\mathbb{T}, \overline{u} \approx^{\downarrow}_q \mathbb{T}', \overline{u}'$. Then the following hold.*

1. *For every u in $\mathbb{T}$ there is a u' in $\mathbb{T}'$ such that $\mathbb{T}, \overline{u} \cdot u \approx^{\downarrow}_q \mathbb{T}', \overline{u}' \cdot u'$.*
2. *For every $U \subseteq_{\omega} \mathbb{T}$ there is a set $U' \subseteq_{\omega} \mathbb{T}'$ such that for every $u' \in U'$ there is a $u \in U$ with $\mathbb{T}, \overline{u} \cdot u \approx^{\downarrow}_q \mathbb{T}', \overline{u}' \cdot u'$.*

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

Proof. Let $\mathbb{T}, \mathbb{T}', \overline{u}$ and $\overline{u}'$ be as in the formulation of the lemma. Write $\overline{u} = u_0 \cdots u_k$ and $\overline{u}' = u'_0 \cdots u'_k$, and recall that u_0 and u'_0 are the roots of $\mathbb{T}$ and $\mathbb{T}'$, respectively. Let $f : \mathbb{T}_{\overline{u}} \to \mathbb{T}'_{\overline{u}'}$ be an isomorphism witnessing the q-similarity of $\overline{u}$ and $\overline{u}'$.

We first prove item 1. Let u be an arbitrary node in $\mathbb{T}$. The case where u is actually a node of $\mathbb{T}_{\overline{u}}$ is easily dealt with by taking $u' := f(u)$, so assume otherwise.

Consider the set

$$X_u := \{lca(u_i, u) \mid 0 \leq i \leq k\}.$$

Since u does not belong to $T_{\overline{u}}$, clearly each node in X_u is a proper ancestor of u, and X_u is linearly ordered by the ancestor relation. In particular, X_u has a *last* element in this order, which we shall denote as d_u. Let Y_u be the set of nodes u_i in $\overline{u}$ for which $d_u \neq u_i$ but $d_u = lca(u_i, u)$. If $Y_u = \emptyset$ it must be the case that for some i, $d_u = u_i$ while $d_u \notin \downarrow(u_j)$ for any $u_j \neq u_i$. This case can be easily dealt with using Lemma 3.12(1), so in the sequel we assume that Y_u is not empty. Note that u_0, being the root of $\mathbb{T}$, does not belong to Y_u.

Consider the set $R(d_u)$ of successors of d_u. For each $x \in Y_u \cup \{u\}$ we let z_x denote the unique successor of d_u which is an ancestor of x, and define $Z_Y := \{z_y \mid y \in Y_u\}$. It follows from the assumptions that $Z_Y \neq \emptyset$ and that $u_0, z_u \notin Z_Y$. Furthermore we obviously have $|Z_Y| \leq |Y_u|$, and since $Y_u \subseteq \{u_i \mid 0 < i \leq k\}$ we may use the assumption $k < q$ to obtain that $|Z_Y| < q$.

Now we claim that $f(d_u)$ has a successor z'_u which is a q-companion of z_u, but is not of the form $f(z)$ for any $z \in Z_Y$. To see this, let $\chi \in \mathtt{GML}^{\infty}_{q-h(z_u),q}$ be a characteristic formula of the $\leftrightarrows^{\infty,\#}_{q-h(z_u),q}$-cell of z_u, and let m be the number of elements of Z_Y that satisfy χ. Obviously then we have $\mathbb{T}, d_u \Vdash \Diamond^{\geq m+1}\chi$, and since $m \leq |Z_Y| < q$, the formula $\Diamond^{\geq m+1}\chi$ belongs to the set $\mathtt{GML}^{\infty}_{(q-h(z_u))+1,q} = \mathtt{GML}^{\infty}_{q-h(u),q}$. Since d_u and $f(d_u)$ are q-companions, the formula $\Diamond^{\geq m+1}\chi$ must be preserved, so that $f(d_u)$ must satisfy $\Diamond^{\geq m+1}\chi$ as well. Hence $f(d)$ has at least $m+1$ successors satisfying χ. Using the facts that f is an isomorphism between $\mathbb{T}_{\overline{u}}$ and $\mathbb{T}'_{\overline{u}'}$ and that $\mathbb{T}, z \Vdash \chi$ iff $\mathbb{T}', f(z) \Vdash \chi$ for all $z \in R(u)$, we may conclude that $|\{z' \in R'(f(d_u)) \cap f[Z_Y] \mid \mathbb{T}', z' \Vdash \chi\}| = m$. That is, there are exactly m successors of $f(d_u)$ that are in the image of $f[Z_Y]$ and make χ true. But then there must a successor z'_u of $f(d)$ that satisfies χ but does not belong to the set $f[Z_Y]$. This node z'_u is then the successor of $f(d_u)$ that we are looking for (Figure 1).

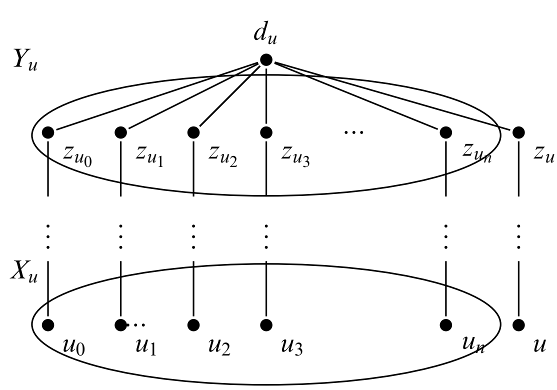

Fig. 1. Lemma 3.14(1).

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

To finish the proof of the first item, observe that as z_u and z'_u are q-companions, it follows from Lemma 3.12(1) that z'_u has a descendant u' which is a q-companion of u. It is then a fairly routine exercise to check that we may extend f to an isomorphism $g : \mathbb{T}_{\overline{u}\cdot u} \to \mathbb{T}'_{\overline{u}'\cdot u'}$ which satisfies in addition that $g(z_u) = z'_u$ and $g(u) = u'$ and that x and $g(x)$ are q-companions, for every x in $\mathbb{T}_{\overline{u}\cdot u}$. This suffices to show that $\mathbb{T}, \overline{u} \cdot u \approx^{\downarrow}_{q} \mathbb{T}', \overline{u}' \cdot u'$.

We now turn to the proof of item 2. Let U be some infinite subset of $\mathbb{T}$. From the constraints on the set U' that we are looking for, it follows that instead of working with U we may work with any subset of U that is guaranteed to be infinite—we will use this observation repeatedly (without explicit reference). For instance, by Lemma 3.11 we may without loss of generality assume that all nodes in U are q-companions of one another.

Similarly, since $T_{\overline{u}}$ is finite we may, without loss of generality assume that $U \cap T_{\overline{u}} = \emptyset$. In addition, we may assume that all nodes in U have the same height and belong to the same q-companion class. Then for each $u \in U$ we define the set X_u, the node d_u, and the set Y_u, exactly as in the proof of the first item. Since each d_u belongs to the (finite) set of nodes of $\mathbb{T}_{\overline{u}}$, and each Y_u is a subset of the finite set $\{u_i \mid 0 < i \leq k\}$, we may without loss of generality assume that $d_u = d_v$ and $Y_u = Y_v$, for all $u, v \in U$. Hence we may drop subscripts and simply write d and Y. If $Y = \emptyset$ it must be the case that $d = u_i$ for some i, and as in the proof of item 1 we leave it for the reader to check that this case can be easily dealt with using Lemma 3.12 (in this case with item (2)).

As in the proof of item 1 of this lemma, we now focus on the set $R(d)$ of direct successors of d. For $x \in Y \cup U$, we define z_x to be the unique element of $R(d)$ which is an ancestor of x, and we set $Z_Y := \{z_y \mid y \in Y\}$ and $Z_U := \{z_u \mid u \in U\}$. Similarly as before we have that $Z_Y \cap Z_U = \emptyset$ and $|Z_Y| < q$. We now make a case distinction.

If Z_U is finite, then we may without loss of generality assume that it is a singleton, say, $Z_U = \{z\}$ where z is an ancestor of every $u \in U$ (Figure 2). In the same way as in the proof of item 1, we may show that $f(d)$ has a successor z' which is a q-companion of z. It then follows by Lemma 3.12(2) that $\mathbb{T}'$ contains an infinite collection U' of nodes that are all descendants of z' and q-companions of the nodes in U. From this it

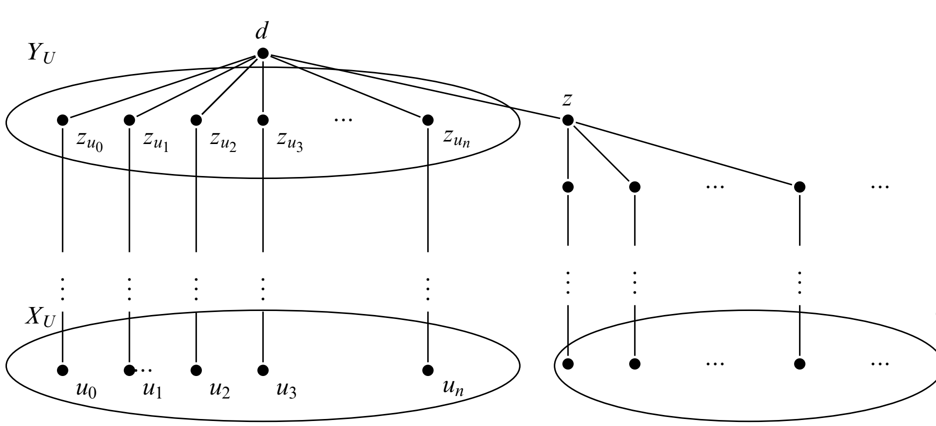

Fig. 2. Lemma 3.14(2) case a.

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

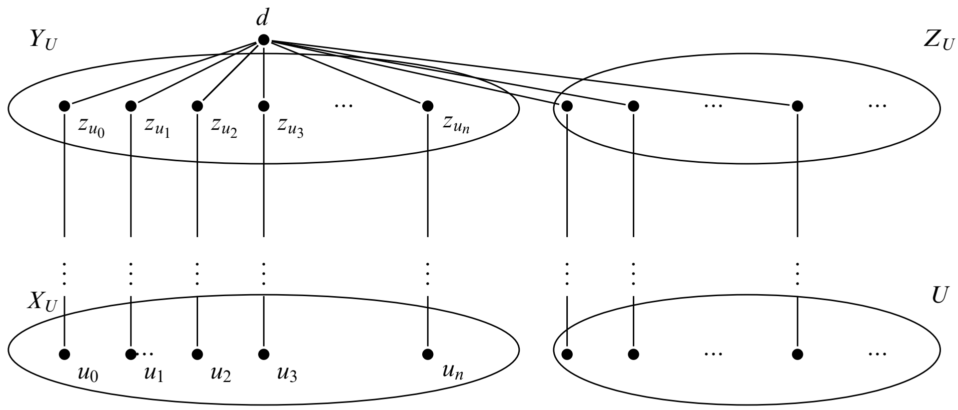

Fig. 3. Lemma 3.14(2) case b.

is easily verified that any $u' \in U'$ and any $u \in U$ satisfy $\mathbb{T}, \overline{u} \cdot u \approx_q^{\downarrow} \mathbb{T}', \overline{u}' \cdot u'$, and this certainly suffices to prove the statement of item 2.

If, on the other hand, Z_U is infinite (Figure 3), our first observation is that all nodes in Z_U belong to the same q-companion class: this easily follows from the assumption (discussed above) that all nodes $u \in U$ are q-companions of one another. Define $\chi \in \mathtt{GML}^{\infty}_{q-h(z),q}$ as the characteristic formula of the $\underline{\leftrightarrow}^{\infty,\#}_{q-h(z),q}$-cell of some/any $z \in Z_U$, then we have $\mathbb{T}, d \Vdash \Diamond^{\infty}\chi$. From this we may use the q-companionship of d and $f(d)$ to derive that $\mathbb{T}', f(d) \Vdash \Diamond^{\infty}\chi$, so that $f(d)$ has an infinite set W' of successors that all satisfy the formula χ. From this it readily follows that each $w' \in W'$ is a q-companion of each $z \in Z_U$. Then by Lemma 3.12 each $w' \in W'$ has a descendant $u'_{w'}$ which is a q-companion of the nodes in U. Finally, it is a routine exercise to show that the set $U' := \{u'_{w'} \mid w' \in W'\}$ meets all the conditions as specified in the lemma. □

Lemma 3.14 gives us all the necessary ingredients to prove Theorem 3.9.

Proof of Theorem 3.9. Fix some natural number $q \geq 1$, let $\mathbb{T}$ and $\mathbb{T}'$ be tree models of height q, with roots r and r', respectively, and assume that $\mathbb{T}, r \underline{\leftrightarrow}^{\infty,\#}_{q,q} \mathbb{T}', r'$. Obviously then, r and r' are q-companions, and from this we readily see that, seen as two 1-tuples, they are q-similar: $\mathbb{T}, r \approx_q^{\downarrow} \mathbb{T}', r'$. It then follows from successive applications of Lemma 3.12 that in the Ehrenfeucht–Fraïssé game $\mathtt{EF}^{\infty}_q$ between $\mathbb{T}, r$ and $\mathbb{T}', r'$, Duplicator has a strategy ensuring that *every* configuration, that is reached in a match of this game, will consist of two q-*similar* tuples. Playing such a strategy, Duplicator thus guarantees that the final q-configuration $(\overline{u}, \overline{u}')$ will be such that $\mathbb{T}, \overline{u} \approx_q^{\downarrow} \mathbb{T}', \overline{u}'$. That is, there is a isomorphism f from $\mathbb{T}_{\overline{u}}$ to $\mathbb{T}'_{\overline{u}'}$ such that $f(u_i) = u'_i$, for all i; in particular then, f is a local isomorphism from $\overline{u}$ to $\overline{u}'$. In other words, any such strategy is a *winning* strategy for her.

This finishes the proof of Theorem 3.9. □

3.5. Proof of bisimulation invariance theorem. We can now collect all the results we obtained in the previous subsections and prove the bisimulation invariance theorem.

Proof of Theorem 2.25. We focus on the forward direction of the proof. The backwards direction can be easily proved by induction on the complexity of the formula ψ.

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

Fix some bisimulation invariant formula $\mathtt{FO}^{\infty}$ formula $\varphi(x)$, and let q be its quantifier depth. Define $n := 3^q$, then it follows from Proposition 3.7 that φ is n-local. We first prove that φ is $\underline{\leftrightarrow}^{\infty,\#}_{n,n}$-invariant (cf. (2)).

For that purpose, let $\mathbb{S}, s$ and $\mathbb{S}', s'$ be two pointed Kripke models such that $\mathbb{S}, s \underline{\leftrightarrow}^{\infty,\#}_{n,n} \mathbb{S}', s'$. We will prove that

$$\mathbb{S} \models \varphi[s] \text{ if and only if } \mathbb{S}' \models \varphi[s']. \tag{6}$$

For a proof of (6), let $\mathbb{U}_{\infty}(\mathbb{S}, s)$ be the unravelling of $(\mathbb{S}, s)$. That is, $\mathbb{U}_{\infty}(\mathbb{S}, s)$ is a tree model whose states are the finite paths of $\mathbb{S}$ that start with s, and whose accessibility relation and valuation are defined in the obvious way. We refer to [5] for the details, and leave it for the reader to verify that $\mathbb{U}_{\infty}(\mathbb{S}, s), s \underline{\leftrightarrow}^{\infty,\#} \mathbb{S}, s$, an observation that we will refer to below as (bisimulation invariance).

We let $\mathbb{U}_n(\mathbb{S}, s)$ be the structure that is based on the n-neighbourhood of s in $\mathbb{U}_{\infty}(\mathbb{S}, s)$—this structure is sometimes referred to as the n-unravelling of $(\mathbb{S}, s)$, since its states correspond to the paths in $\mathbb{S}$ of length at most n that start at s. It follows by locality that $\mathbb{U}_{\infty}(\mathbb{S}, s) \models \varphi[s]$ iff $\mathbb{U}_n(\mathbb{S}, s) \models \varphi[s]$, and of course a similar observation applies to the unravellings of $(\mathbb{S}', s')$.

Furthermore, it is not hard to see that $\mathbb{S}, s \underline{\leftrightarrow}^{\infty,\#}_{n,n} \mathbb{S}', s'$ implies $\mathbb{U}_n(\mathbb{S}, s), s \underline{\leftrightarrow}^{\infty,\#}_{n,n} \mathbb{U}_n(\mathbb{S}', s'), s'$. From this we may derive by Theorem 3.9 that $\mathbb{U}_n(\mathbb{S}, s), s \cong^{\infty}_{n} \mathbb{U}_n(\mathbb{S}', s'), s'$, and so by Proposition 2.29 we find that $\mathbb{U}_n(\mathbb{S}, s) \models \varphi[s]$ iff $\mathbb{U}_n(\mathbb{S}', s') \models \varphi[s']$ $(*)$.

Now consider the following chain of equivalences:

$$\begin{array}{lr}
\mathbb{S} \models \varphi[s] \text{ iff } \mathbb{U}_{\infty}(\mathbb{S}, s) \models \varphi[s] & \text{(bisimulation invariance)} \\
\qquad\quad \text{iff } \mathbb{U}_n(\mathbb{S}, s) \models \varphi[s] & \text{(locality, Theorem 2.25)} \\
\qquad\quad \text{iff } \mathbb{U}_n(\mathbb{S}', s') \models \varphi[s'] & \text{((*) above)} \\
\qquad\quad \text{iff } \mathbb{U}_{\infty}(\mathbb{S}', s') \models \varphi[s'] & \text{(locality, Theorem 2.25)} \\
\qquad\quad \text{iff } \mathbb{S}' \models \varphi[s']. & \text{(bisimulation invariance)}
\end{array}$$

This finishes the proof of (6).

But if φ is $\underline{\leftrightarrow}^{\infty,\#}_{n,n}$-invariant, then the class of pointed models satisfying φ is $\underline{\leftrightarrow}^{\infty,\#}_{n,n}$-saturated (i.e., a union of $\underline{\leftrightarrow}^{\infty,\#}_{n,n}$-cells). It is then immediate from Proposition 3.5 that φ is equivalent to some formula in $\mathtt{GML}^{\infty}_{n,n}$. □

§4. Completeness and small model property. In this section we will prove two results for our logic: the Small Model Property (Theorem 2.20) and a completeness result (Theorem 4.2). In our proofs we will use coalgebras for a set functor $\mathsf{B}^{\bullet}$ which is closely related to the infinitary bag functor B^{∞}.

For an overview of the section: in the first subsection we define and discuss the axiomatisation $\mathbf{G}^{\infty}$, and we formulate our completeness result. In Section 4.2 we introduce the set functor $\mathsf{B}^{\bullet}$ and its coalgebras. Here we also formulate the main theorem of the section, Theorem 4.5, which states that any $\mathtt{GML}^{\infty}$-formula φ is derivable iff it is valid in, respectively, the classes of all Kripke frames, all B^{∞}-coalgebras, all B^{∞}-coalgebras of bounded size, and all $\mathsf{B}^{\bullet}$-coalgebras. In the subsequent subsections we then prove the two nontrivial implications of this theorem: in Section 4.3 we prove $\mathbf{G}^{\infty}$ to be complete for the class of $\mathsf{B}^{\bullet}$-coalgebras, and in Section 4.4 we establish a filtration lemma which shows that any formula that is satisfiable in a $\mathsf{B}^\bullet$-coalgebra is also satisfiable in a B^∞-coalgebra of bounded size.

***4.1. The axiomatisation* $\mathbf{G}^\infty$.** In this subsection we will introduce an axiomatisation for the set of $\mathtt{GML}^\infty$-validities, that is, the formulas that are valid in every Kripke frame.

Since the axiomatisation problem for graded modal logic has been addressed already [7, 11], we can take this as our starting point. The standard axiomatisation **G** for graded modal logic consists of the following axioms:

1. all classical propositional tautologies,
2. $\Box^k \varphi \leftrightarrow \neg\Diamond^k \neg\varphi$,
3. $\Box^1(\varphi \to \psi) \to (\Diamond^k \varphi \to \Diamond^k \psi)$, for every $k \in \mathbb{N}$,
4. $\Diamond^{k+1}\varphi \to \Diamond^k \varphi$, for every $k \in \mathbb{N}$,
5. $\neg\Diamond^1(\varphi \wedge \psi) \wedge \Diamond^{=k_1}\varphi \wedge \Diamond^{=k_2}\psi \to \Diamond^{=(k_1+k_2)}(\varphi \vee \psi)$, for every $k_1, k_2 \in \mathbb{N}$.

In addition, **G** has the following derivation rules:

- modus ponens: from φ and $\varphi \to \psi$ derive ψ;
- $\Box$-necessitation: from φ derive $\Box\varphi$.

Most of these axioms are fairly obvious. The Dual axiom 2 syntactically captures the Boolean duality between the $\Diamond, \Box$ modalities as it was semantically described in Definition 2.2. Axiom 3 expresses the monotonicity of the graded modalities, while axiom 4 concerns the relative strength of the counting modalities, basically stating that "at least $k + 1$" implies "at least k." Finally, axiom 5 states that we may add the numbers of two disjoint groups of successors.

The logic **G** is defined as the smallest set of $\mathtt{GML}$-formulas that contains the previous axioms and is closed under the inference rules described above. As mentioned, for this axiomatisation soundness and completeness has been proved [9, 12], so that **G** consists exactly of the set of $\mathtt{GML}$-validities.

To capture the behaviour of the infinity modality $\Diamond^\infty$ as well, we extend the logic **G** with the following axioms:

6. $\Box^\infty \varphi \leftrightarrow \neg\Diamond^\infty \neg\varphi$,
7. $\Diamond^\infty \varphi \to \Diamond^k \varphi$, for every $k \in \mathbb{N}$,
8. $\Box^\infty(\varphi \to \psi) \to (\Box^\infty \varphi \to \Box^\infty \psi)$.

Furthermore, we add the following derivation rule:

- $\Box^\infty$-necessitation: from φ derive $\Box^\infty\varphi$.

DEFINITION 4.1. $\mathbf{G}^\infty$ *is the smallest set of* $\mathtt{GML}^\infty$*-formulas which contains the axioms 1–8 above, and is closed under the derivation rules of modus ponens and of necessitation, for both* $\Box$ *and* $\Box^\infty$. *A formula* φ *is called a* theorem of $\mathbf{G}^\infty$ *or* derivable in $\mathbf{G}^\infty$ *if it belongs to* $\mathbf{G}^\infty$, *notation*: $\vdash_{\mathbf{G}^\infty} \varphi$ (*or simply* $\vdash \varphi$, *if no confusion is likely*). *The formula* φ *is* derivable from a set of formulas Γ *if there are finitely many formulas* $\psi_1, \ldots, \psi_n$ *in* Γ *such that the formula* $(\psi_1 \wedge \cdots \wedge \psi_n) \to \varphi$ *is derivable in* $\mathbf{G}^\infty$. *A set* Γ *of* $\mathtt{GML}^\infty$*-formulas is* $\mathbf{G}^\infty$-consistent *if* $\bot$ *is not derivable from it*.

Here are some intuitions on the new axioms and derivation rule. Axiom 6 extends the duality expressed in axiom 2 to account for the $\Diamond^\infty$ case. Axiom 7 can be interpreted as an extension of axiom 4 to include the infinite case, expressing that the existence of infinitely many successors of a certain kind implies the existence of an unbounded finite number of these successors. To complete the axiomatisation we make sure that the infinity box is a *normal* modality. This is done by including the K-axiom 8 and the necessitation rule for $\Box^{\infty}$. As a consequence of this, the infinity diamond distributes over disjunctions:

$$\vdash_{\mathbf{G}^{\infty}} \Diamond^{\infty}(p \vee q) \leftrightarrow (\Diamond^{\infty} p \vee \Diamond^{\infty} q).$$

We conclude this subsection by stating one of the main results of this paper, the soundness and completeness theorem for the $\mathtt{GML}^{\infty}$-logic $\mathbf{G}^{\infty}$:

THEOREM 4.2 (Soundness and Completeness). *The logic $\mathbf{G}^{\infty}$ is sound and complete with respect to the class of all Kripke frames. That is, for an arbitrary $\mathtt{GML}^{\infty}$-formula φ we have*

$$\vdash_{\mathbf{G}^{\infty}} \quad \textit{iff } \varphi \textit{ is valid in every Kripke frame.}$$

REMARK. *Let $\mathbf{K}^{\infty}$ be the normal bimodal logic for $\mathtt{ML}^{\infty}$ that we obtain by adding the axioms 6 and 8, as well as the necessitation rule for $\Box^{\infty}$ to the basic modal logic $\mathbf{K}$. It was proved by Fattorosi-Barnaba and Balestrini* [11], *and* (*independently*) *by the first author* [2], *that $\mathbf{K}^{\infty}$ is a sound and complete axiomatisation for the set of all $\mathtt{ML}^{\infty}$-formulas that are valid in every Kripke frame.*

4.2. $\mathsf{B}^{\bullet}$***-coalgebras.*** In this section we introduce the auxiliary structures that play a major role in the proof of the completeness theorem and the small model property, viz., the $\mathsf{B}^{\bullet}$-coalgebras. These structures are very similar in nature to the B^{∞}-coalgebras introduced in Section 2, the difference being that the functor $\mathsf{B}^{\bullet}$ admits a weight ω in between the finite numbers and ∞. This new weight ω allows us to make a subtle difference in the semantics of the infinity diamond, and thus enables us to take care of sets of the form

$$\{\Diamond^{k}\varphi \mid k \in \mathbb{N}\} \cup \{\neg\Diamond^{\infty}\varphi\}$$

which may be consistent (for instance if φ is a proposition letter) but are never satisfiable.

DEFINITION 4.3. *We define $\mathbb{N}^{\omega,\infty} := \mathbb{N} \cup \{\omega, \infty\}$, and extend the ordering $<$ of $\mathbb{N}^{\infty}$ to this set by putting $n < \omega$ for all $n \in \mathbb{N}$, as well as $\omega < \infty$. We will also need to introduce some arithmetic in $\mathbb{N}^{\omega,\infty}$; we first define a binary addition $\oplus$ using the following table*:

$\oplus$	n	ω	∞
m	$m+n$	ω	∞
ω	ω	ω	∞
∞	∞	∞	∞

We now extend this binary addition to an infinitary operation as follows:

$$\bigoplus_{i \in I} \kappa_i := \begin{cases} \infty, & \textit{if } \kappa_i = \infty \textit{ for some } i \in I, \\ \sum_{i \in I'} \kappa_i, & \textit{if } I' := \{i \in I \mid \kappa_i \neq 0\} \textit{ is finite}, \\ \omega, & \textit{otherwise.} \end{cases}$$

Note that the difference between the two operations $\sum_{i \in I}$ and $\bigoplus_{i \in I}$ lies in the particular case where we add infinitely many non-zero natural numbers k_i: this will give $\sum_{i \in I} k_i = \infty$ but $\bigoplus_{i \in I} k_i = \omega$.

The set functor $\mathsf{B}^{\bullet}$ is defined very similarly to B^{∞}, and the same holds for $\mathsf{B}^{\bullet}$-coalgebras and $\mathsf{B}^{\bullet}$-models.

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

DEFINITION 4.4. *For a set S, we define $\mathsf{B}^\bullet S := S \to \mathbb{N}^{\omega,\infty}$, and for a function $f : S \to S'$, we define the map $(\mathsf{B}^\bullet f) : \mathsf{B}^\bullet(S) \to \mathsf{B}^\bullet(S')$ by putting*

$$(\mathsf{B}^\bullet f)(\mu) : s' \mapsto \bigoplus_{s \in f^{-1}(s')} \mu(s).$$

A $\mathsf{B}^\bullet$-coalgebra *is a pair (S, σ) such that $\sigma : S \to \mathsf{B}^\bullet(S)$ assigns a $\mathsf{B}^\bullet$-weight function to each state $s \in S$. A* $\mathsf{B}^\bullet$-model *is a triple $\mathbb{S} = (S, \sigma, V)$ such that (S, σ) is a $\mathsf{B}^\bullet$-coalgebra and V is a valuation on S.*

In these models we interpret formulas in the language $\mathtt{GML}^\infty$ using the obvious definitions for the atomic formulas and Boolean connectives, while the modalities are interpreted using the following clause:

$$\mathbb{S}, s \Vdash^\bullet \Diamond^\kappa \varphi \text{ iff } \kappa \leq \bigoplus_{t \Vdash^\bullet \varphi} \sigma_s(t).$$

Therefore, this means that in a $\mathsf{B}^\bullet$-model we have

$$\mathbb{S}, s \Vdash^\bullet \Diamond^\infty \varphi \text{ iff } \mathbb{S}, t \Vdash^\bullet \varphi \text{ for some } t \text{ with } \sigma_s(t) = \infty.$$

Recall that in a B^∞-model, on the other hand, we can also have $\mathbb{S}, s \Vdash \Diamond^\infty \varphi$ if there are infinitely many t such that $\mathbb{S}, t \Vdash \varphi$ but $0 < \sigma_s(t) < \infty$ for all t.

The $\mathsf{B}^\bullet$-coalgebras may not be very interesting in their own right. However, they do provide the pivotal semantics which link the notions of Kripke validity, derivability and validity in small B^∞-coalgebras, as is witnessed by the following theorem.

THEOREM 4.5. *The following are equivalent, for every formula $\varphi \in \mathtt{GML}^\infty$:*

1. *φ is valid (i.e., valid in every Kripke frame);*
2. *φ is valid in every B^∞-coalgebra;*
3. *φ is valid in every B^∞-coalgebra of size at most $2^{|\varphi|}$ and weight at most $r(\varphi)$;*
4. *φ is valid in every $\mathsf{B}^\bullet$-coalgebra;*
5. *φ is a theorem of $\mathbf{G}^\infty$.*

Proof. The implication from (1) to (2) follows by Corollary 2.18, and the implication from (2) to (3) is trivial. The implications from (3) to (4) and from (4) to (5) will be proved in the Sections 4.4 and 4.3, respectively. Finally, the implication from (5) to (1) states that our axiomatisation is *sound*. The verification of this statement is a routine exercise which we leave to the reader. □

4.3. Completeness for* $\mathsf{B}^\bullet$*-coalgebras. In this subsection we prove one implication of Theorem 4.5, which corresponds to the completeness of our axiom system $\mathbf{G}^\infty$ for the class of $\mathsf{B}^\bullet$-coalgebras.

THEOREM 4.6 (Completeness for $\mathsf{B}^\bullet$-coalgebras). *The following holds, for every formula $\varphi \in \mathtt{GML}^\infty$:*

$$\textit{if } \varphi \textit{ is valid in every } \mathsf{B}^\bullet\textit{-coalgebra, then } \varphi \textit{ is a theorem of } \mathbf{G}^\infty. \tag{7}$$

We will prove Theorem 4.6 using a canonical model construction, somewhat similar to the proof of De Caro in [9] for graded modal logic.

For the rest of the section fix a countably infinite set P of proposition letters. Define a set $\Gamma \subseteq \mathtt{GML}^\infty(P)$ to be *maximally consistent* if Γ is consistent and it does not have a proper consistent extension $\Delta \subseteq \mathtt{GML}^\infty(P)$. We shall abbreviate "maximal consistent set" as MCS. We let C denote the set of all maximally consistent sets of formulas, and we define the canonical valuation V^c in the usual way, that is, $V^c(p) := \{\Phi \in C \mid p \in \Phi\}$.

For the definition of the $\mathsf{B}^\bullet$-coalgebra map $\gamma : C \to \mathsf{B}^\bullet(C)$, consider two MCSs Φ and Ψ. We say that Ψ has *finite multiplicity for* Φ if there is some $n \in \mathbb{N}$ and some formula $\psi \in \Psi$ such that $\Diamond^{=n}\psi \in \Phi$. Clearly, in the situation where such an n exists, there is a unique *minimal* one; we will refer to this number as the *multiplicity* of Ψ for Φ and denote it as $\mu_\Phi(\Psi)$. We now define the map $\gamma : C \to (C \to \mathbb{N}^{\omega,\infty})$ as follows:

$$\gamma_\Phi(\Psi) := \begin{cases} \infty, & \text{if } \Diamond^\infty\psi \in \Phi, \text{ for all } \psi \in \Psi, \\ \mu_\Phi(\Psi), & \text{if } \Psi \text{ has finite multiplicity for } \Phi, \\ \omega, & \text{otherwise.} \end{cases}$$

It is not difficult to see that for any two MCSs Φ and Ψ:

$$\gamma_\Phi(\Psi) = \omega \text{ if and only if } \{\Diamond^k\psi \mid k \in \mathbb{N}\} \cup \{\neg\Diamond^\infty\psi\} \subseteq \Phi \text{ for some } \psi \in \Psi.$$

We shall refer to the triple $\mathbb{C} := (C, \gamma, V^c)$ as the *canonical* $\mathsf{B}^\bullet$*-model*.

The crux of the proof of Theorem 4.6 is the following *Truth Lemma*.

PROPOSITION 4.7. *Let Φ be a $\mathbf{G}^\infty$-MCS. Then we have*

$$\mathbb{C}, \Phi \Vdash^\bullet \varphi \textit{ iff } \varphi \in \Phi,$$

for all formulas $\varphi \in \mathtt{GML}^\infty$.

In order to prove the *Truth Lemma* we need some auxiliary facts that can easily be verified.

FACT 1. *Let Φ be a $\mathbf{G}^\infty$-MCS, and let φ be some $\mathtt{GML}^\infty$-formula. Then*:

1) *if $\Diamond^1\varphi \in \Phi$ then there is some $\mathbf{G}^\infty$-MCS Ψ with $\varphi \in \Psi$ and $\gamma_\Phi(\Psi) \geq 1$;*
2) *if $\Diamond^\infty\varphi \in \Phi$ then there is some $\mathbf{G}^\infty$-MCS Ψ with $\varphi \in \Psi$ and $\gamma_\Phi(\Psi) = \infty$.*

FACT 2. *Let $\Psi_0, \dots, \Psi_k$ be a collection of distinct maximal $\mathbf{G}^\infty$-consistent sets. For every $i \leq k$ we can find a $\mathtt{GML}^\infty$ formula α_i such that $\alpha_i \in \Psi_j$ if and only if $i = j$.*

The key proposition in our proof of the Truth Lemma is the following observation.

PROPOSITION 4.8. *Let Φ be a maximal $\mathbf{G}^\infty$-consistent set of formulas and let κ be any value in $\mathbb{N}^\infty$. Then we have*

$$\Diamond^\kappa\varphi \in \Phi \textit{ iff } \bigoplus_{\Psi \ni \varphi} \gamma_\Phi(\Psi) \geq \kappa$$

for all formulas $\varphi \in \mathtt{GML}^\infty$.

Proof. This proof is an adaptation of Theorem 1 in De Caro's completeness proof [7]. Hence, rather than giving a detailed proof we confine ourselves to a sketch. In the sequel we will abbreviate $m_\Psi := \gamma_\Phi(\Psi)$. The proof is based on a case distinction. Case $\kappa = 0$: This case is trivial since the formula $\Diamond^0\varphi$, being equivalent to $\top$, belongs to every MCS, while on the other hand every $\bigoplus$-type sum in $\mathbb{N}^{\omega,\infty}$ is at least as big as 0.

Case $\kappa \in \{1, \infty\}$: The forward direction follows directly from Fact 1. For the opposite direction, suppose that $\bigoplus_{\Psi \ni \varphi} m_\Psi \geq \kappa$, then by definition of $\bigoplus$ there must be at least one MCS Ψ such that $m_\Psi \geq \kappa$. It now follows from the definition of γ_Φ that $\Diamond^\kappa\varphi \in \Phi$.

Case $\kappa = n \in \mathbb{N}\backslash\{0,1\}$: Let I be the collection of MCSs Ψ such that $\varphi \in \Psi$ and $m_\Psi \geq 1$. If $m_\Psi \geq n$ for some $\Psi \in I$, then the result is obvious. Thus we examine the case where $m_\Psi < n$ for every $\Psi \in I$. We make a further distinction as to the size of I.

Subcase: I is finite. For each $\Psi \in I$ we let α_Ψ be some formula as in Fact 2 and β_Ψ the defining formula that makes Ψ have multiplicity m_Ψ for Φ (as described above). Then, defining ψ_Ψ to be $\alpha_\Psi \wedge \beta_\Psi \wedge \varphi$, we may easily prove that:

1. $\Diamond^{=m_\Psi}\psi_\Psi \in \Phi$, for all Ψ in I,
2. $\vdash \neg(\psi_\Psi \wedge \psi_\Theta)$, for any distinct Ψ and Θ in I,
3. $\vdash \bigvee_{\Psi\in I}\psi_\Psi \to \varphi$.

Combining the observations 1 and 2 with Axiom 5, one may verify that $\Diamond^{=m}\bigvee_{\Psi\in I}\psi_\Psi \in \Phi$, where $m := \sum_{\Psi\in I} m_\Psi$. Furthermore, if we let θ be the formula $\bigwedge_{\Psi\in I}\neg(\alpha_\Psi \vee \beta_\Psi)$, the following facts easily follow:

4. $\vdash \varphi \leftrightarrow \big((\varphi\wedge\theta) \vee \bigvee_{\Psi\in I}\psi_\Psi\big)$,
5. $\neg\Diamond(\varphi\wedge\theta) \in \Phi$.

If we now combine observation 5 and the fact that $\Diamond^{=m}\bigvee_{\Psi\in I}\psi_\Psi \in \Phi$ with Axiom 5, we obtain that $\Diamond^{=m}\big((\varphi\wedge\theta)\vee\bigvee_{\Psi\in I}\psi_\Psi\big) \in \Phi$. Based on this and observation 4 one may conclude that $\Diamond^{=m}\varphi \in \Phi$. From this it easily follows that

$$\Diamond^n\varphi \in \Phi \text{ iff } n \leq m.$$

Furthermore, since I is a finite set, and m_Ψ is a natural number for every $\Psi \in I$, we find that

$$\bigoplus_{\Psi\ni\varphi}\gamma_\Phi(\Psi) = \bigoplus_{\Psi\in I}\gamma_\Phi(\Psi) = \sum_{\Psi\in I}\gamma_\Phi(\Psi) = \sum_{\Psi\in I} m_\Psi = m.$$

From the above two observations the statement of the proposition readily follows.

Subcase: I is infinite. In this case we find $\bigoplus_{\Psi\ni\varphi}\gamma_\Phi(\Psi) = \omega$, so that we need to prove that $\Diamond^n\varphi \in \Phi$. For this purpose we take some set $J \subseteq I$ of size n, and define $m' := \sum_{\Psi\in J} m_\Psi$. Then clearly we have $n \leq m'$. Using similar reasoning as before, we may find a family $\{\psi_\Psi \mid \Psi \in J\}$ of formulas satisfying the items 1–3 above, with J replacing I. Defining $\varphi' := \bigvee_{\Psi\in J}\psi_\Psi$, we may show, again as above, that $\Diamond^{=m'}\varphi' \in \Phi$. But then since $\vdash \varphi' \to \varphi$ and $n \leq m'$ we find $\Diamond^n\varphi \in \Phi$ as required.

This finishes the proof of Proposition 4.8. □

As we will see now, the Truth Lemma and, based on that, the Completeness Theorem itself are immediate consequences of Proposition 4.8.

Proof of Proposition 4.7. We prove the Truth Lemma by a routine induction on the formula φ, taking care of the cases for the modalities using Proposition 4.8. □

Proof of Theorem 4.6. We prove the completeness theorem by showing that every consistent formula is satisfiable in a $\mathsf{B}^\bullet$-model. Let φ be a $\mathtt{GML}^\infty$-formula that is consistent with respect to the proof system $\mathbf{G}^\infty$. By a routine Lindenbaum-type construction we may obtain an MCS Γ containing the formula φ, so that by the Truth Lemma φ is actually true at Γ in the canonical $\mathsf{B}^\bullet$-model. In particular then, φ is satisfiable in a $\mathsf{B}^\bullet$-model. □

4.4. Filtrations: from* $\mathsf{B}^\bullet$*-coalgebras to small* B^∞*-coalgebras. In this subsection we provide the proof of the remaining implication of Theorem 4.5, viz., the statement that validity in bounded-size B^∞-coalgebras implies validity in all $\mathsf{B}^\bullet$-coalgebras. In fact, we will prove the contrapositive statement.

THEOREM 4.9. *Let φ be some formula in $\mathtt{GML}^\infty$. If φ is satisfiable in some $\mathsf{B}^\bullet$-model, then φ is satisfiable in a B^∞-model of size at most $2^{|\varphi|}$ and rank $r(\varphi)$.*

Our proof will be based on a filtration argument, that is, we will show that $\mathsf{B}^\bullet$-models can be quotiented to finite B^∞-models. Our definition of the filtrated model involves choice functions, an idea originating from the literature on coalgebra. We need the notion of a *closed* set of formulas.

DEFINITION 4.10. *A set Σ of $\mathtt{GML}^\infty$-formulas is* closed *if it is closed under taking subformulas and single negations (that is, if $\varphi \in \Sigma$ is not of the form $\neg\psi$ then $\neg\varphi$ also belongs to Σ).*

As usual the filtration of a model will be based on the quotient of the domain of the model under a natural equivalence relation induced by a closed set of formulas.

DEFINITION 4.11. *Let $\mathbb{S} = (S, \sigma, V)$ be a $\mathsf{B}^\bullet$-model. We define the equivalence relation $\equiv_\Sigma$ on S by putting $s \equiv_\Sigma s'$ iff we have that $\mathbb{S}, s \Vdash^\bullet \varphi \Leftrightarrow \mathbb{S}, s' \Vdash^\bullet \varphi$, for all $\varphi \in \Sigma$.*

In the remainder of this section we fix a finite closed set Σ, so that we may write $\equiv$ instead of $\equiv_\Sigma$ without confusion. Recall that a *choice function* on a collection P of nonempty subsets of a set S is any map $c : P \to S$ such that $c(X) \in X$, for all $X \in P$. We are now ready for the definition of a filtration.

DEFINITION 4.12. *Let $\mathbb{S} = (S, \sigma, V)$ be a $\mathsf{B}^\bullet$-model. The equivalence class of $s \in S$ under the relation $\equiv$ will be denoted as $\overline{s}$, and we write $\overline{S} := \{\overline{s} \mid s \in S\}$. The valuation $\overline{V}$ is given by putting $\overline{V}(p) := \{\overline{s} \mid s' \in V(p)$, for some $s' \equiv s\}$.*

Given a choice function $c : \overline{S} \to S$, we define the map $\sigma^c : \overline{S} \to \mathsf{B}^\infty(\overline{S})$ as follows:

$$\sigma^c_{\overline{s}}(\overline{t}) := \begin{cases} \infty, & \text{if } \sigma_{c(\overline{s})}(t') = \infty, \text{ for some } t' \in \overline{t}, \\ \min\left\{r(\Sigma), \bigoplus_{t' \equiv t} \sigma_{c(\overline{s})}(t')\right\}, & \text{otherwise.} \end{cases}$$

Any B^∞-model of the form $\overline{\mathbb{S}} = (\overline{S}, \sigma^c, \overline{V})$ is called a Σ-filtration *of $\mathbb{S}$.*

In words, the weight function σ^c of the filtrated model is defined as follows. In principle, for the weight of a state $\overline{t}$ according to a state $\overline{s}$ we would like to take the $\bigoplus$-sum of the weights of all the members of $\overline{t}$, according to the chosen element $c(\overline{s}) \in \overline{s}$. This candidate for the weight $\sigma^c_{\overline{s}}(\overline{t})$ is capped off, however, by the finitary modal rank of Σ.

The key result about this construction is that the natural quotient map (i.e., the one sending any state to its $\equiv$-cell) preserves the truth of all formulas in Σ.

PROPOSITION 4.13 (Filtration Lemma). *Let Σ be a finite, closed set of $\mathtt{GML}^\infty$-formulas, let $\mathbb{S}$ be a $\mathsf{B}^\bullet$-model, and let $\overline{\mathbb{S}}$ be a Σ-filtration of $\mathbb{S}$. Then we have*

$$\mathbb{S}, s \Vdash^\bullet \varphi \text{ iff } \overline{\mathbb{S}}, \overline{s} \Vdash \varphi,$$

for all $\varphi \in \Sigma$ and all $s \in S$.

Proof. This proposition is proved by induction on the complexity of φ. Let c be the choice function inducing the filtration.

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

Case $\varphi = \Diamond^k \psi$, *with* $k \in \mathbb{N}$. First assume that $\mathbb{S}, s \Vdash^\bullet \Diamond^k \psi$. Since $\Diamond^k \psi \in \Sigma$ and $s \equiv c(\overline{s})$, we find that $\mathbb{S}, c(\overline{s}) \Vdash^\bullet \Diamond^k \psi$. Now, we can distinguish two different cases. First suppose that $\sigma^c_{\overline{s}}(\overline{t}) \geq \mathrm{r}(\Sigma)$, for some $\overline{t} \Vdash \psi$. Since $k \leq \mathrm{r}(\Sigma)$ we are automatically done, as we immediately obtain that $\mathbb{S}, \overline{s} \Vdash \Diamond^k \psi$.We may thus focus on the case where $\sigma^c_{\overline{s}}(\overline{t})$ is strictly smaller than $\mathrm{r}(\Sigma)$, for all $\overline{t}$ in $\overline{S}$ such that $\overline{\mathbb{S}}, \overline{t} \Vdash \psi$. This means in particular that $\sigma^c_{\overline{s}}(\overline{t})$ is finite, and by the definition of σ^c, that

$$\sigma^c_{\overline{s}}(\overline{t}) = \bigoplus_{t' \equiv t} \sigma_{c(\overline{s})}(t') = \sum_{t' \equiv t} \sigma_{c(\overline{s})}(t'),$$

for every such state $\overline{t}$.

Furthermore, it follows from the induction hypothesis and the definition of the relation $\equiv$ that $\mathbb{S}, t' \Vdash^\bullet \psi$ for all $t' \in \overline{t}$ such that $\overline{\mathbb{S}}, \overline{t} \Vdash \psi$. Thus we find that

$$\sum_{\overline{t} \Vdash \psi} \sigma^c_{\overline{s}}(\overline{t}) = \sum_{\overline{t} \Vdash \psi} \sum_{t' \in \overline{t}} \sigma_{c(\overline{s})}(t') = \sum_{t' \Vdash^\bullet \psi} \sigma_{c(\overline{s})}(t') \geq k, \tag{8}$$

where the latter inequality holds simply because $\mathbb{S}, c(\overline{s}) \Vdash^\bullet \Diamond^k \psi$. But from (8) it is immediate that $\overline{\mathbb{S}}, \overline{s} \Vdash \Diamond^k \psi$.

For the opposite direction, assume that $\overline{\mathbb{S}}, \overline{s} \Vdash \Diamond^k \psi$, and suppose for a contradiction that $\mathbb{S}, c(\overline{s}) \nVdash^\bullet \Diamond^k \psi$, i.e., that

$$\sum_{u \Vdash^\bullet \psi} \sigma_{c(\overline{s})}(u) < k.$$

Consider an arbitrary $\overline{t} \in \overline{S}$ such that $\overline{\mathbb{S}}, \overline{t} \Vdash \psi$. By the inductive hypothesis, any t' in $\overline{t}$ satisfies $\mathbb{S}, t' \Vdash^\bullet \psi$. From this it is immediate that

$$\sum_{t' \equiv t} \sigma_{c(\overline{s})}(t') \leq \sum_{u \Vdash^\bullet \psi} \sigma_{c(\overline{s})}(u) < k.$$

Combining these two observations with the (obvious) fact that $k \leq \mathrm{r}(\Sigma)$ we obtain that

$$\sum_{t' \equiv t} \sigma_{c(\overline{s})}(t') < \mathrm{r}(\Sigma).$$

Hence, by the definition of σ^c, we obtain that

$$\sigma^c_{\overline{s}}(\overline{t}) = \sum_{t' \equiv t} \sigma_{c(\overline{s})}(t').$$

Since this holds for every $\overline{t} \in \overline{S}$ such that $\overline{\mathbb{S}}, \overline{t} \Vdash \psi$, we may conclude that

$$\sum_{\overline{t} \Vdash \psi} \sigma^c_{\overline{s}}(\overline{t}) = \sum_{u \Vdash^\bullet \psi} \sigma_{c(\overline{s})}(u) < k,$$

contradicting our assumption that $\overline{\mathbb{S}}, \overline{s} \Vdash \Diamond^k \psi$.

We thus have $\mathbb{S}, c(\overline{s}) \Vdash^\bullet \Diamond^k \psi$, and from this and the fact that $s \equiv c(\overline{s})$ it follows that $\mathbb{S}, s \Vdash^\bullet \Diamond^k \psi$.

Case $\varphi = \Diamond^\infty \psi$. First assume that $\mathbb{S}, s \Vdash^\bullet \Diamond^\infty \psi$, then by definition of $\equiv$ it follows that $\mathbb{S}, c(\overline{s}) \Vdash^\bullet \Diamond^\infty \psi$. By the semantics there must be a state $t \in \mathbb{S}$ such that $\mathbb{S}, t \Vdash^\bullet \psi$ and $\sigma_{c(\overline{s})}(t) = \infty$. From the latter fact and by the definition of

https://doi.org/10.1017/S1755020322000247 Published online by Cambridge University Press

σ^c, it is immediate that $\sigma^c_{\overline{s}}(\overline{t}) = \infty$, while we may use the inductive hypothesis to find that $\overline{\mathbb{S}}, \overline{t} \Vdash \psi$. From this we can derive that $\overline{\mathbb{S}}, \overline{s} \Vdash \Diamond^\infty \psi$.

For the opposite direction, assume that $\overline{\mathbb{S}}, \overline{s} \Vdash \Diamond^\infty \psi$. Since $\overline{S}$ is finite, this can only be the case if there is a state $\overline{t} \in \overline{S}$ such that $\sigma^c_{\overline{s}}(\overline{t}) = \infty$ and $\overline{\mathbb{S}}, \overline{t} \Vdash \psi$. Then the inductive hypothesis yields that $\mathbb{S}, t \Vdash^\bullet \psi$, and by the definition of σ^c we have that $\sigma_{c(\overline{s})}(t') = \infty$, for some $t' \in S$ such that $t' \equiv t$. It then follows from $t' \equiv t$ that $\mathbb{S}, t' \Vdash^\bullet \psi$, so that $\mathbb{S}, c(\overline{s}) \Vdash^\bullet \Diamond^\infty \psi$ by the semantics in $\mathsf{B}^\bullet$-models; finally we conclude from $s \equiv c(\overline{s})$ that $\mathbb{S}, s \Vdash^\bullet \Diamond^\infty \psi$, as required.

This finishes the proof of the Filtration Lemma. □

The proof of Theorem 4.9 is an immediate corollary of the Filtration Lemma.

Proof of Theorem 4.9. Let $\varphi \in \mathtt{GML}^\infty$ be satisfiable in some $\mathsf{B}^\bullet$-model $\mathbb{S}$, say, $\mathbb{S}, s \Vdash^\bullet \varphi$. Define Σ as the *closure of* φ, that is, the smallest closed set containing φ, and let c be an arbitrary choice function on the set of $\equiv_\Sigma$-cells. It follows from the Filtration Lemma that φ is satisfiable at the state $\overline{s}$ of the filtrated model $\overline{\mathbb{S}} = (\overline{S}, \sigma^c, \overline{V})$. Finally, it is straightforward to verify that $\overline{\mathbb{S}}$ is a B^∞-model of the required size and weight. □

Acknowledgment. The authors are grateful to the anonymous referees for various helpful suggestions to improve the presentation of the paper.

BIBLIOGRAPHY

[1] Barwise, J. & Feferman, S., editors (2017). *Model-Theoretic Logics*. Perspectives in Logic. New York: Cambridge University Press.

[2] Bellas Acosta, I. (2020). Studies in the extension of standard modal logic with an infinity modality. Master's Thesis, Institute for Logic, Language and Computation, Universiteit van Amsterdam.

[3] van Benthem, J. (1976). Modal correspondence theory. Ph.D. Thesis, Universiteit van Amsterdam.

[4] van Benthem, J., ten Cate, B., & Väänänen, J. (2007). Lindström theorems for fragments of first-order logic. In Ong, L., editor. *Proceedings of the 22nd IEEE Symposium on Logic in Computer Science (LICS 2007)*. Washington, DC: IEEE Computer Society, pp. 280–289.

[5] Blackburn, P., de Rijke, M., & Venema, Y. (2001). *Modal Logic*. Cambridge Tracts in Theoretical Computer Science, Vol. 53. Cambridge: Cambridge University Press.

[6] Carreiro, F., Facchini, A., Venema, Y., & Zanasi, F. (2022). Model theory of monadic predicate logic with the infinity quantifier. *Archive for Mathematical Logic*, **61**(3–4), 465–502.

[7] Cîrstea, C., Kurz, A., Pattinson, D., Schröder, L., & Venema, Y. (2011). Modal logics are coalgebraic. *The Computer Journal*, **54**, 524–538.

[8] D'Agostino, G. & Visser, A. (2002). Finality regained: A coalgebraic study of Scott-sets and multisets. *Archive for Mathematical Logic*, **41**, 267–298.

[9] De Caro, F. (1988). Graded modalities II: Canonical models. *Studia Logica*, **47**(1), 1–10.

[10] Emerson, E. A. & Halpern, J. Y. (1986). "Sometimes" and "not never" revisited: On branching versus linear time temporal logic. *Journal of the ACM* **33**(1), 151–178.

[11] Fattorosi-Barnaba, M. & Balestrini, U. P. (1999). Graded modalities VII: The modality of finite. *Mathematical Logic Quarterly*, **45**, 471–480.

[12] Fattorosi-Barnaba, M. & De Caro, F. (1985). Graded modalities I. *Studia Logica*, **44**(2), 197–221.

[13] Fattorosi-Barnaba, M. & Grassotti, S. (1995). Graded modalities VI: An infinitary graded modal logic. *Mathematical Logic Quarterly*, **41**, 547–563.

[14] Fine, K. (1972). In so many possible worlds. *Notre Dame Journal of Formal Logic*, **13**, 516–520.

[15] Goble, L. F. (1970). Grades of modality. *Logique et Analyse*, **13**(51), 323–334.

[16] Goranko, V. & Otto, M. (2007). Model theory of modal logic. In Blackburn, P., Benthem, J., and Wolter, F., editors. *Handbook of Modal Logic*. Amsterdam: Elsevier, pp. 249–329.

[17] van der Hoek, W. & de Rijke, M. (1995). Counting objects. *Journal of Logic and Computation*, **5**, 325–345.

[18] van der Hoek, W. & Meyer, J.-J. C. (1992). Graded modalities in epistemic logic. In Nerode, A. and Taitslin, M., editors. *Logical Foundations of Computer Science—Tver'92*. Berlin–Heidelberg: Springer, pp. 503–514.

[19] Krawczyk, A. & Krynicki, M. (1976). Ehrenfeucht games for generalized quantifiers. In Marek, W., Srebrny, and Zarach, A., editors. *Set Theory and Hierarchy Theory*: *A Memorial Tribute to Andrzej Mostowski*. Heidelberg: Springer, pp. 145–152.

[20] Lindström, P. (1966). First order predicate logic with generalized quantifiers. *Theoria*, **32**(3), 186–195.

[21] Mostowski, A. (1957). On a generalization of quantifiers. *Fundamenta Mathematicae*, **44**(1), 12–36.

[22] ———. (1968). Craig's interpolation theorem in some extended systems of logic. In van Rootselaar, B. and Staal, J., editors. *Logic, Methodology and Philosophy of Science III*. Studies in Logic and the Foundations of Mathematics, Vol. 52. Amsterdam: Elsevier, pp. 87–103.

[23] Otto, M. (2019). Graded modal logic and counting bisimulation. Preprint, arXiv:1910.00039.

[24] de Rijke, M. (2000). A note on graded modal logic. *Studia Logica*, **64**(2), 271–283.

[25] Rosen, E. (1997). Modal logic over finite structures. *Journal of Logic, Language and Information*, **6**(4), 427–439.

[26] Rutten, J. (2000). Universal coalgebra: A theory of systems. *Theoretical Computer Science*, **249**, 3–80.

[27] Tobies, S. (2001). PSPACE reasoning for graded modal logics. *Journal of Logic and Computation*, **11**, 85–106.

[28] Väänänen, J. (2011). *Models and Games*. Cambridge Studies in Advanced Mathematics, Vol. 132. Cambridge: Cambridge University Press.

[29] Westerståhl, D. (2016). Generalized quantifiers. In Zalta, E., editor. *The Stanford Encyclopedia of Philosophy (Winter 2016 Edition)*. Stanford, CA: Stanford University.

INDEPENDENT SCHOLAR
E-mail: ignaciobellas15@gmail.com

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION
UNIVERSITEIT VAN AMSTERDAM
P.O. BOX 94242, 1090 GE AMSTERDAM
THE NETHERLANDS
E-mail: y.venema@uva.nl