Tero Tulenheimo

Insights into Modal Slash Logic and Modal Decidability

UNIVERSITY OF TAMPERE
SCHOOL OF INFORMATION SCIENCES
REPORTS IN INFORMATION SCIENCES 6

TAMPERE 2012

UNIVERSITY OF TAMPERE
SCHOOL OF INFORMATION SCIENCES
REPORTS IN INFORMATION SCIENCES 6
JANUARY 2012

Tero Tulenheimo

Insights into Modal Slash Logic and Modal Decidability

SCHOOL OF INFORMATION SCIENCES
FIN-33014 UNIVERSITY OF TAMPERE

ISBN 978-951-44-8725-5

ISSN-L 1799-8158
ISSN 1799-8158

Insights into Modal Slash Logic and Modal Decidability *

Tero Tulenheimo
School of Information Sciences, University of Tampere /
Laboratoire Savoirs, Textes, Langage*, Centre National de la Recherche Scientifique*
tero.tulenheimo@univ-lille3.fr

Abstract. The present paper has a two-fold task. On the one hand, it aims to provide an overview on Independence friendly modal logic as defined in (Tulenheimo, 2003; Tulenheimo, 2004) and studied in a number of subsequent publications. For systematic reasons to be explained, the logic is here referred to as *modal slash logic* (MsL). On the other hand, we take a close look at a syntactic fragment of MsL, to be termed MsL_0, first formulated in (Tulenheimo and Sevenster, 2006). We push the study of this logic deeper at several points: a model-theoretic criterion is presented which serves to tell when a formula of MsL_0 is not truth-equivalent to any formula of basic modal logic (ML); the game-theoretic property of 'bounded quasi-positionality' of MsL_0 is studied in detail; an alternative syntax for MsL_0 is discerned and the logic obtained is shown to enjoy the property of quasi-locality (generalizing the notion of locality familiar from ML); and we formulate an asymmetric bisimulation concept and use it to prove that MsL_0 is not closed under complementation. Drawing from insights provided by the study of MsL_0, we conclude by general observations about claims made on the 'reasons' why various modal logics are computationally well-behaved.

Keywords: complementation, decidability, expressivity, IF logic, independence, modal logic, slash logic

1. Introduction

Independence-friendly (IF) first-order logic (Hintikka and Sandu, 1989; Hintikka, 1996) results from first-order logic by dissociating the notions of syntactic and semantic scope in the following sense: formulas of IF first-order logic are like those of first-order logic, except that in place of plain existential quantifiers, expressions of the form $(\exists x/\forall y_1, \ldots, \forall y_n)$ may appear in a sentence ϕ, provided that in ϕ this expression is syntactically subordinate to (in the syntactic scope of) each of the universal quantifiers $\forall y_1, \ldots, \forall y_n$. The semantic effect of such independence indications '$/\forall y_1, \ldots, \forall y_n$' is taken to be that in order for ϕ to be true, the witness of $\exists x$ may not depend on the values corresponding to the

* I am grateful to Hans van Ditmarsch for having encouraged me to work for a paper on modal slash logic developing the results I presented in a survey talk in *Seminario de Lógica y Lenguaje* at the University of Sevilla on October 19, 2009. Further, I am indebted to Antti Kuusisto for useful comments on an earlier version of this manuscript.

universal quantifiers $\forall y_1, \ldots, \forall y_n$, that is, this value must be provided by a (Skolem) function *not* having the values chosen for those quantifiers among its arguments. Thereby $\forall x(\exists y/\forall x)Rxy$ is true iff $\exists y \forall x Rxy$ is true, and $\forall x \forall y(\exists z/\forall y)(\exists v/\forall x)Sxyzv$ is true iff the Henkin-quantifier sentence

$$\begin{pmatrix} \forall x & \exists z \\ \forall x & \exists v \end{pmatrix} Sxyzv$$

is true. Expressed in terms of two-player evaluation games among players **E** and **A**, the truth of an IF first-order sentence, hence defined, amounts to the existence of a winning strategy for **E** in the correlated evaluation game — supposing that strategy functions of a given player only take as arguments moves made by his or her adversary. However, this is not how strategies are typically understood in game theory: generally a strategy for a player in a two-player game can perfectly well make use of the previous choices of either player, his or her own ones included.

Hodges (2007) points out that in the literature some authors discussing IF first-order logic have opted for formulating the semantics as Hintikka does while others have utilized strategy functions in the standard game-theoretic sense in their semantics. Hodges (1997a) had adopted the notational convention of writing $(\exists x/y)$ where Hintikka writes $(\exists x/\forall y)$, hence marking the difference between semantic games formulated in terms of arbitrary strategy functions and those whose strategy functions are in effect Skolem functions. The variable y in $(\exists x/y)$ might in Hodges's syntax be 'bound' even by a syntactically precedent existential quantifier, whereas such independence of existential quantifiers cannot be syntactically marked in Hintikka's formulation. Hodges (2007, p. 119) writes:

> [W]e refer to the logic with my notation and the general game semantics as slash logic. During recent years many writers in this area (but never Hintikka himself) have transferred the name 'IF logic' to slash logic, often without realising the difference. Until the terminology settles down, we have to beware of examples and proofs that don't make clear which semantics they intend.

The terminology has not shown much tendency of settling down; for instance the recent monograph by Mann, Sandu and Sevenster (2011) discusses in fact first-order slash logic but calls its object of study nevertheless Independence friendly first-order logic.

In the present paper we discuss a modal-like logic which in fact relates to first-order slash logic (FOsL) in the same way as basic modal logic (ML) relates to first-order logic (FOL): just like ML can be translated into FOL, also our modal-like logic — to be referred to as modal slash logic (MsL) — semantically corresponds to a certain fragment of FOsL, as will be explained in Section 3. Different variants of modal slash logic can be discerned by restricting, or liberalizing, the sorts of independence indications allowed; see (Tulenheimo, 2003; Tulenheimo, 2004; Hyttinen and Tulenheimo, 2005; Tulenheimo and Sevenster, 2006; Sevenster, 2006; Tulenheimo and Sevenster, 2007; Tulenheimo and Rebuschi, 2009; Tulenheimo, 2009; Sevenster, 2010). The first formulation of a logic termed 'IF modal logic' was due to Bradfield and Fröschle (2002), who developed further the framework of Bradfield's Henkin modal logic (Bradfield, 2000). The key idea in their research was to use the analysis of quantifier independence in studying transition systems with *concurrency*; for comparison it should be noted that the semantics of MsL is relative to standard modal structures (Kripke models). Recently Väänänen has considered accommodating the framework of his dependence logic (Väänänen, 2007) to the case of modal logic, cf. (Väänänen, 2008; Sevenster, 2009). The logics emerging from the works of Bradfield and Väänänen are not studied in the present paper.

1.1. Basic definitions

We need to lay down some definitions.

1.1.1. *Logics.*

Throughout the paper, **prop** will be a fixed countably infinite set of *propositional atoms*, denoted p, q, r etc. The syntax of *basic modal logic* (or ML) is generated by the grammar $\phi ::= \top \mid \bot \mid p \mid {\sim}p \mid (\phi \vee \phi) \mid (\phi \wedge \phi) \mid \Diamond\phi \mid \Box\phi$, where $p \in \textbf{prop}$. The semantics of ML is defined relative to *models* and their states, i.e., tuples $\mathfrak{M} = (M, R, V)$ and elements $w \in M$, where M is a non-empty domain, R is a binary relation on M termed *accessibility relation*, and $V : \textbf{prop} \to Pow(M)$ is a *valuation function.* If $\mathfrak{M}$ is a model and $w \in M$, $(\mathfrak{M}, w)$ is a *pointed model.* We assume the reader is familiar with the relation $\mathfrak{M}, w \models \phi$, or truth of an ML formula ϕ in $\mathfrak{M}$ at w; see, e.g., (Blackburn *et al.*, 2002, Def. 1.20). The symbols $\top$ and $\bot$ stand for *verum* and *falsum*, respectively: they have a fixed semantics, with $\mathfrak{M}, w \models \top$ and $\mathfrak{M}, w \not\models \bot$ for all pointed models $(\mathfrak{M}, w)$. The *modal depth* of an ML formula ϕ is the maximum number of nested modal operator tokens in ϕ.

Fix a vocabulary τ (i.e., a countable set of constant, relation, and function symbols). We write FOL$[\tau]$ for first-order logic of vocabulary τ; ESO$[\tau]$ stands for existential second-order logic of vocabulary τ, and FOsL$[\tau]$ is first-order slash logic of vocabulary τ. Formulas of ESO$[\tau]$ are strings $\exists f_1 \dots \exists f_n \phi$, where the f_i are function symbols and ϕ is an FOL formula of vocabulary $\tau \cup \{f_1, \dots, f_n\}$. The syntax of FOsL$[\tau]$ is produced by the grammar $\phi \ ::= \ \alpha \mid (\phi \vee \phi) \mid (\phi \wedge \phi) \mid \forall x\phi \mid (\exists x/W)\phi$, where α is an atomic or negated atomic formula of $\mathsf{FOL}[\tau]$ and $W \cup \{x\}$ is a finite set of variables, with $x \notin W$. We agree to write $\exists x$ for $(\exists x/\varnothing)$. If $x_1, \ldots, x_n$ are pairwise distinct variables, we allow writing $(\exists x/x_1, \ldots, x_n)$ for $(\exists x/\{x_1, \ldots, x_n\})$. Expressions produced by the grammar are termed $\mathsf{FOsL}[\tau]$ *strings*. We may wish to distinguish two kinds of free variables in an FOsL string ϕ: those appearing in atomic formulas ($\mathit{Free}_1[\phi]$) and those appearing independence indications ($\mathit{Free}_2[\phi]$). The former sets are recursively defined as in the case of first-order logic, the independence indications playing no role. To define the latter sets, put $\mathit{Free}_2[(\exists x/W)\phi] = W \cup (\mathit{Free}_2[\phi] \setminus \{x\})$, $\mathit{Free}_2[\forall x\phi] = \mathit{Free}_2[\phi]$, $\mathit{Free}_2[\phi \star \psi] = \mathit{Free}_2[\phi] \cup \mathit{Free}_2[\psi]$ for junctions $\star$, and $\mathit{Free}_2[\alpha] = \varnothing$ for (negated) atoms α. A string ϕ is a *formula* if the set $\mathit{Free}_2[\phi]$ is empty. If even the union $\mathit{Free}_1[\phi] \cup \mathit{Free}_2[\phi]$ is empty, ϕ is a *sentence*. For example, $\exists x\exists y(\exists z/x)Rxyz$ is a sentence (and *a fortiori* a formula). The string $\exists x(\exists z/x)Rxyz$ is not a sentence but is a formula. And $\exists y(\exists z/x)Syz$ not even a formula. In the present paper we will assume, for simplicity, that in an FOsL string there will never appear two quantifiers carrying the same variable so that one would be syntactically subordinate to the other: hence, e.g., $\forall x\exists y\forall x(\exists z/x)Rxyz$ is not a string. Thanks to this stipulation, if $(\exists x/W)$ appears in a formula, the variables in the set W refer to *uniquely determined* syntactically preceding quantifiers. To simplify further, we also suppose that if a variable belongs to $\mathit{Free}_1[\phi]$, then no quantifier carrying this variable occurs in ϕ. The notion of *syntactic tree* is defined in the expected way for all logics discussed in the present paper: the nodes of a syntactic tree of ϕ are the operator tokens and tokens of atomic formulas occurring in ϕ, and they are ordered by the relation of syntactic precedence given by the syntax of the logic in question.

To distinguish τ-structures from modal structures, we use calligraphic symbols such as $\mathcal{M}$ and $\mathcal{N}$ for the former. If $\mathcal{M}$ is such a structure, M stands for its domain. For the semantics of $\mathsf{ESO}[\tau]$, see e.g. (Ebbinghaus and Flum, 1999; Väänänen, 2007). As to $\mathsf{FOsL}[\tau]$, given any $\mathsf{FOsL}[\tau]$ *formula* ϕ with $\mathit{Free}_1[\phi] = \{x_1, \ldots, x_k\}$, τ-structure $\mathcal{M}$, and a variable assignment $x_i \mapsto c_i$, there will be an evaluation game $G(\phi, \mathcal{M}, c_1 \cdots c_k)$ between two players, **A** and **E**. There are two kinds of moves: model moves (an element of M is chosen) and junction moves (a term 0 or 1 of a junction is chosen). The positions in the game are triples $(\vec{a}, \vec{i}, \psi)$, where $\vec{a}$ is a tuple keeping track of model moves made in the course of a play, $\vec{i}$ being a tuple keeping track of junction moves made in the course of a play. The initial position of the game is $(c_1 \cdots c_k, \emptyset, \phi)$. If a position $(\vec{a}, \vec{i}, \psi)$ has been reached, the continuation of the play depends on the form of ψ. If $\psi = (\chi_0 \vee \chi_1)$, player **E** selects $j \in \{0,1\}$ and the play continues with the position $(\vec{a}, \vec{i}j, \chi_j)$. It is the player **A** who chooses if $\psi = (\chi_0 \wedge \chi_1)$. If, again, $\psi = (\exists x/W)\chi$, player **E** selects $b \in M$ and the play continues with the position $(\vec{a}b, \vec{i}, \chi)$; note that the rule in no way utilizes the independence indication. If $\psi = \forall x \chi$, player **A** selects an element $b \in M$. Finally, if ψ is (negated) atomic, **E** wins and **A** loses if ψ is satisfied in $\mathcal{M}$ by the variable assignment induced by the tuple $\vec{a}$, else the players receive the reversed payoffs. Note that indeed we may think of $\vec{a}$ as a variable assignment: thanks to our syntactic assumptions, the same variable is never re-interpreted in the course of a play, and therefore the tuples $\vec{a}$ and $\vec{i}$ together uniquely determine an assignment of type $Free_1[\psi] \rightarrow M$ from the input formula ϕ. We could have avoided encoding information about the junction moves in the description of positions, provided that we would have considered the formulas mentioned in positions as formula *tokens*. It serves clarity, however, to explicitly list the junction moves made, since a tuple of junction moves together with a formula type uniquely determines which formula token is meant.

Consider the positions $(\vec{a}, \vec{i}, (\exists x/W)\psi)$ and $(\vec{b}, \vec{i}, (\exists x/W)\psi)$ of one and the same game; note that necessarily the tuples $\vec{a}$ and $\vec{b}$ are of the same length. We say that these positions are *W-equivalent*, given that the following holds: whenever $a_i \neq b_i$, then a_i and b_i are both model moves made for some quantifier Qy with $y \in W$. A strategy of player **E** in game $G(\phi, \mathcal{M})$ is *uniform* if for any W-equivalent positions $(\vec{a}, \vec{i}, (\exists x/W)\psi)$ and $(\vec{b}, \vec{i}, (\exists x/W)\psi)$, the strategy yields the same move. By definition, an FOsL formula ϕ is *satisfied* (*dissatisfied*) in $\mathcal{M}$ under assignment $\vec{x} \mapsto \vec{c}$ if there is a winning strategy for player **E** (player **A** respectively) in game $G(\phi, \mathcal{M}, \vec{c})$. If there is no winning strategy for either player, the formula ϕ is said to be *non-determined* in $\mathcal{M}$ under the assignment $\vec{x} \mapsto \vec{c}$. If ϕ is a sentence, we say that ϕ is *true* (*false*, *non-determined*) in $\mathcal{M}$ when ϕ is satisfied (respectively false, non-determined) in $\mathcal{M}$ by $\emptyset$. For instance, as the reader may readily check, the sentence $\forall x(\exists y/x)\, x = y$ is neither true nor false in any model of size at least 2. We say that sentences ϕ and ψ are *truth equivalent* if they are true in precisely the same structures. Since non-truth does not in general amount to falsity, truth-equivalence does not imply that ϕ and ψ are also false in precisely the same structures. More generally, if ϕ and ψ are formulas, it is said that they are *satisfaction equivalent* if they are satisfied by exactly the same structures and the same suitable variable assignments. For FOsL, see (Hodges, 1997a; Hodges, 2007) and cf. (Väänänen, 2007; Mann *et al.*, 2011).

1.1.2. *Model-theoretic notions.*
If R is a binary relation, write R^+ for the transitive closure of R and R^* for the reflexive transitive closure of R. A modal structure is *tree-like* if its accessibility relation R satisfies: (i) there is a unique element $r \in M$, the root of the model, such that for all $x \in M$, R^*rx; (ii) every element of M distinct from r has a unique R-predecessor; and (iii) R is acyclic, i.e., there is no x such that R^+xx. The *unraveling* of a pointed model (M, R, V, w) is a pointed model (M', R', V', w) such that the domain $M' = \{(x_0, x_1, \ldots, x_n) \in M^{n+1} : x_0 = w$ and $n < \omega$ and Rx_ix_{i+1} for all $0 \leq i < n\}$, the accessibility relation R' satisfies $\langle (w, x_1, \ldots, x_n), (w, y_1, \ldots, y_m) \rangle \in R'$ iff ($m = n + 1$ and $x_i = y_i$ for all $1 \leq i \leq n$ and Rx_ny_m), and the valuation V' satisfies $(w, x_1, \ldots, x_n) \in V'(p)$ iff $x_n \in V(p)$, for all propositional atoms p. Evidently for any pointed model (M, R, V, w) we have that (M', R', V') is tree-like and w is its root. (We identify unit tuples and elements.) Model (N, S, U) is a *submodel* of model (M, R, V) if $N \subseteq M$, S is a restriction of R to N and U is a restriction of V to N. Further, (N, S, U) is a *generated submodel* of (M, R, V) if it is a submodel of (M, R, V) and satisfies the following closure condition: if $w \in N$ and $R(w, v)$, then $v \in N$. Submodel of (M, R, V) *generated by* $S \subseteq M$ is by definition the smallest generated submodel whose domain contains S. The notion of *substructure* is defined similarly for first-order structures of vocabularies containing only one binary but any number of unary relation symbols. A modal logic $\mathfrak{L}$ evaluated over pointed models is said to be *invariant under generated submodels* provided that the following holds: if $\mathfrak{N}$ is a generated submodel of $\mathfrak{M}$, and $w \in N$, then $\mathfrak{M}, w \models \psi$ iff $\mathfrak{N}, w \models \psi$, for all formulas ψ of $\mathfrak{L}$. The notion of invariance under generated substructures extends in an obvious way to abstract logics, in particular FOL. Given a pointed model (M, R, V, w), define a partial map $h : M \rightharpoonup \omega$ as follows: put $h(w) = 0$, and if $h(u) = n$ and Ruv and v has not yet been assigned a value, put $h(v) = n + 1$. This recursive definition assigns a natural number $h(v)$ — the *height* of v — to every state v in M for which there exists a path from w to v along R. That is, the map h is defined on all and only states in the domain of the submodel of (M, R, V, w) generated by the singleton set $\{w\}$. The height of a state v in a tree-like model $\mathfrak{M}$ equals the height of v in the pointed model $(\mathfrak{M}, w)$, where w is the root of $\mathfrak{M}$. If the set $\{h(v) : v \in M\}$ has a maximum, we say that $(\mathfrak{M}, w)$ is of *finite height* and take this quantity to be the *height of the pointed model* $(\mathfrak{M}, w)$. Else $(\mathfrak{M}, w)$ is said to be of infinite height. If $\mathfrak{M} = (M, R, V)$ is tree-like and $k < \omega$, the *restriction of* $\mathfrak{M}$ *to height* k, denoted $(\mathfrak{M} \upharpoonright k)$, is the model (M', R', V'), where M' consists of the states of $\mathfrak{M}$ of height at most k, and R' (respectively V') is the restriction of R (respectively

V) to M'. If v is a state having a finite number, n, of successors along the accessibility relation, we say that the *out-degree* of v is n; if the number of successors has no finite bound, the out-degree of v is said to be infinite.

1.1.3. *Expressivity and bisimulations.*
Suppose the semantics of logics $\mathfrak{L}$ and $\mathfrak{L}'$ are defined relative to pointed models. Logic $\mathfrak{L}$ is *translatable* into logic $\mathfrak{L}'$ (in symbols $\mathfrak{L} \leq \mathfrak{L}'$) if for every $\phi \in \mathfrak{L}$ there is $\psi_\phi \in \mathfrak{L}'$ such that for all $\mathfrak{M}$ and w, we have: $\mathfrak{M}, w \models \phi$ iff $\mathfrak{M}, w \models \psi_\phi$. And $\mathfrak{L}'$ is *strictly more expressive* than $\mathfrak{L}$ (in symbols $\mathfrak{L} < \mathfrak{L}'$) if $\mathfrak{L} \leq \mathfrak{L}'$ but $\mathfrak{L}' \not\leq \mathfrak{L}$. We say that logic $\mathfrak{L}$ is *closed under complementation* — or: *closed under contradictory negation* — if for every formula ϕ of $\mathfrak{L}$ there is a formula $neg(\phi)$ likewise of $\mathfrak{L}$ such that for all pointed models $(\mathfrak{M}, w)$, we have: $\mathfrak{M}, w \models neg(\phi)$ iff $\mathfrak{M}, w \not\models \phi$. Pointed models $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ are $\mathfrak{L}$ *equivalent* if for every $\phi \in \mathfrak{L}$, the following holds: $\mathfrak{M}, w \models \phi$ iff $\mathfrak{M}', w' \models \phi$. For later purposes we recall the notions of bisimulation and n-bisimulation; for more details, see e.g. (Blackburn *et al.*, 2002, Sect. 2.2). Suppose $\mathfrak{M} = (M, R, V)$ and $\mathfrak{M}' = (M', R', V')$ are pointed models, $w \in M$ and $w' \in M'$. A relation $Z \subseteq M \times M'$ is a *bisimulation* between the pointed models $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$, if the following four conditions hold: (0) *Initial coordination:* wZw'; (1) *Atomic harmony:* for all states u and u', if uZu', then for all atoms $p \in \mathbf{prop}$ we have that $u \in V(p)$ iff $u' \in V'(p)$; (2) *Zig:* for all states u, u' and t, if uZu' and $R(u, t)$, then there is t' such that $R'(u', t')$ and tZt'; (3) *Zag:* for all states u, u' and t', if uZu' and $R'(u', t')$, then there is t such that $R(u, t)$ and tZt'. It is well known and easy to prove that bisimulations offer a criterion for ML equivalence. Indeed, if there is bisimulation between $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$, then these pointed models are ML equivalent, i.e., they cannot be distinguished by any ML formula. A sequence $Z_0 \supseteq \ldots \supseteq Z_n$ of relations with $Z_i \subseteq M \times M'$ is an *n-bisimulation* between $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ if we have: (0) *Initial coordination:* wZ_nw'; (1) *Atomic harmony:* for all states u and u', if uZ_0u', then for all atoms $p \in \mathbf{prop}$ we have that $u \in V(p)$ iff $u' \in V'(p)$; (2) *Zig:* for all $i < n$ and for all states u, u' and t, if $uZ_{i+1}u'$ and $R(u, t)$, then there is t' such that $R'(u', t')$ and tZ_it'; (3) *Zag:* for all $i < n$ and for all states u, u' and t', if $uZ_{i+1}u'$ and $R'(u', t')$, then there is t such that $R(u, t)$ and tZ_it'. Using n-bisimulations a criterion for ML equivalence of pointed models up to a modal depth is obtained: if there is an n-bisimulation between $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$, these pointed models cannot be distinguished by any ML formula of modal depth at most n.

1.2. Plan of the paper

Section 2.1 introduces modal slash logic (MsL); the expressive power of this logic is discussed in Section 3. In Section 4 we discern a fragment of MsL, to be denoted MsL_0, and formulate a criterion which serves to tell when a formula of MsL_0 is not truth-equivalent to any formula of basic modal logic (ML). In Section 5 the game-theoretic property of 'bounded quasi-positionality' of MsL_0 is studied in detail. An alternative syntax for MsL_0 is presented in Section 6; the logic obtained is shown to enjoy the property of quasi-locality which is a generalization of the notion of locality familiar from ML. In Section 7 modal slash logic is compared for its expressive power with first-order logic, finite-variable fragments of first-order logic and loosely guarded fragment of first-order logic. Section 8 introduces the notion of breadth of an MsL_0 formula; and it is shown that while the number of non-equivalent MsL_0 formulas of a fixed modal depth is infinite, the number of non-equivalent MsL_0 formulas of a fixed modal depth *and a fixed breadth* is finite. We define an asymmetric bisimulation concept in Section 9 and employ this notion to prove that MsL_0 is not closed under complementation. Section 10 closes the paper by making some general observations about claims made in the literature on the 'reasons' why various modal logics are computationally well-behaved.

2. The logic MsL

We begin by defining modal slash logic (MsL). It is a syntactic extension of basic modal logic (ML). As a matter of fact, MsL is strictly more expressive than ML.

2.1. Syntax of MsL

We need to choose a syntax suitable for indicating the sorts of independence relations we are interested in. Having fixed a set **prop** of propositional atoms, we take as our point of departure basic modal logic in negation normal form (or **ML**): its formulas are obtained from literals (i.e., formulas of the forms p and $\sim p$ with $p \in \mathbf{prop}$) by finitely many applications of the unary operators $\Box$, $\Diamond$ and the binary operators $\wedge$, $\vee$. We wish to study constraints on finding a witness when interested in the *truth* of formulas. This leads us, for reasons that will become clear in the sequel, to allow *slashed diamonds* in the syntax of modal slash logic. By means of a syntactic slashing device ($\Diamond/\ldots$) we will be able to mark diamonds as independent of a selection of syntactically preceding boxes and diamonds (referred to by suitable identifiers written in place of the dots). What this means is that syntactic scope and semantic scope are dissociated: here 'independent of' means 'not in the semantic scope of.' Generalizations of this syntax could be considered, where not only diamonds but also boxes and even conjunction and disjunction symbols could be slashed. These further options will not be explored in the present paper.

In first-order slash logic independence indications are expressed using variables with the aid of suitable further stipulations. If an existential quantifier $\exists y$ is followed by the independence indication '$/x$,' its witness must not depend on the *closest syntactically precedent* quantifier carrying the variable x. (In fact, thanks to our stipulations concerning the syntax, there cannot be more than one quantifier token with the variable x preceding a given quantifier $\exists y$.) Since the syntax of modal logic involves no variables, here we may not use anything like variables to mark independence relations. Several choices are conceivable, one of which would be to employ explicit indexing of tokens of modal operators within a formula (for example by numerals standing for natural numbers) and using those indices in independence indications. This would lead to formulas such as $\Diamond_1\Box_2(\Diamond_3/2)p$. Since the use of indices would mark a difference compared with ML (whose formulas do not carry indexed modal operator tokens), we allow as the only difference between the syntaxes of ML and MsL that diamonds may carry independence indications of the form '$/i_1, \ldots, i_k$,' where the i_j are pairwise distinct positive integers. Semantically these numbers are construed as *relative de Bruijn indices* (to be explained in Definition 2.1); they identify those syntactically preceding modal operator tokens of which the relevant diamond symbol is declared to be independent. For example, we will have formulas such as

$$\Box\Box(\Diamond/1)p \quad \text{and} \quad \Box\Diamond(p \vee (\Diamond/1,2)q).$$

In the first formula, the numeral '1' refers back to the box token which is the immediate syntactic predecessor of the unique diamond symbol. In the second formula, the numeral '1' serves to identify the diamond symbol in whose syntactic scope we find $(p \vee (\Diamond/1,2)q)$, while the numeral '2' refers back to the unique box token in the formula.

The syntax of MsL — or *modal slash logic* — is obtained from that of ML by the following rules:

1. If $\phi \in \mathsf{ML}$ and ϕ' results from replacing in ϕ all tokens of $\Diamond$ by the symbol $(\Diamond/\emptyset)$, then ϕ' is a formula.

2. If ϕ is a formula, $i_1, \ldots, i_k$ is a non-empty strictly increasing tuple of (numerals standing for) positive integers, and α is a token of $(\Diamond/\emptyset)$

that lies in ϕ in the syntactic scope of at least $\max\{i_1, \ldots, i_k\} = i_k$ modal operators, then the result of replacing α in ϕ by the symbol $(\Diamond/i_1, \ldots, i_k)$ is also a formula.

According to the above syntax, all diamond tokens in MsL formulas are of the form $(\Diamond/i_1, \ldots, i_k)$ for some tuple $i_1, \ldots, i_k$. If $k = 0$, the tuple is empty. We abbreviate diamond tokens $(\Diamond/\emptyset)$ by $\Diamond$. We observe that on the basis of this notational convention, all ML formulas are, syntactically, MsL formulas. Examples of further formulas are $\Box(\Diamond/1)p$ and $\Box\Diamond(\Diamond/1,2)q$, as well as

$$\Box((\Diamond/1)p \vee \Diamond(\Diamond/2)q), \ \Box(\Diamond/1)(\Diamond/1)p \quad \text{and} \quad \Box\Box(\Diamond/1)(\Diamond/3)p.$$

By contrast, the strings $\Box(\Diamond/5, 27)p$ and $\Diamond(\Box/1)q$ cannot be produced by the above rules. In the former case this is because there are fewer than 27 modal operator tokens syntactically preceding the diamond symbol in the string $\Box(\Diamond/5, 27)p$; and in the latter case because the syntax does not allow slashing boxes.

As already hinted at, the numerals $i_1, \ldots, i_k$ appearing in independence indications '$/i_1, \ldots, i_k$' are interpreted as *referring* to certain preceding modal operator tokens. The semantics will impose as a condition for finding a witness for the diamond token $(\Diamond/i_1, \ldots, i_k)$ that the witness must not depend on the values chosen for the operators to which the numerals $i_1, \ldots, i_k$ refer. In order to specify which preceding operator is identified by which numeral, we adopt the convention of interpreting them as de Bruijn indices.

DEFINITION 2.1 (Relative de Bruijn index). *Let α and β be modal operator tokens appearing in a given* MsL *formula. We say that β immediately precedes α if β syntactically precedes α and in the relevant formula there is no modal operator token that syntactically precedes α but is syntactically preceded by β. If β immediately precedes α, de Bruijn index of β relative to α is* 1*. And if de Bruijn index of β relative to α is n and γ immediately precedes β, de Bruijn index of γ relative to α is $n+1$.*

For instance, the unique box token in $\Box((\Diamond/1)p \vee \Diamond(\Diamond/1,2)q)$ has de Bruijn index 1 relative to $(\Diamond/1)$, which is why the numeral '1' in $(\Diamond/1)$ by definition refers to that box token. The very same box token has de Bruijn index 2 relative to $(\Diamond/1, 2)$. As this example illustrates, in distinct slashed diamond tokens distinct numerals may refer to the same preceding modal operator token. The numeral '1' in $(\Diamond/1, 2)$ refers to the diamond token immediately preceding $(\Diamond/1, 2)$. By the syntax, for any modal operator token $(\Diamond/i_1, \ldots, i_k)$ appearing in an

MsL formula, there are in that formula preceding modal operator tokens to which the numerals $i_1, \ldots, i_k$ refer.

2.2. Semantics of MsL

The semantics of MsL — like that of ML — is defined relative to models and their states. For later purposes we adopt the following convention: if $\mathfrak{M} = (M, R, V)$ is a model with **prop** $= \{p_1, p_2, \ldots\}$, then $\mathcal{M}$ is the corresponding first-order structure, i.e., the structure $(M, \mathsf{R}^{\mathcal{M}}, \mathsf{P}_1^{\mathcal{M}}, \mathsf{P}_2^{\mathcal{M}}, \ldots)$, with $\mathsf{R}^{\mathcal{M}} = R$ and $\mathsf{P}_i^{\mathcal{M}} = V(p_i)$.

EXAMPLE 2.2. *On what condition should the formula $\Box(\Diamond/1)p$ be true in model $\mathfrak{M}$ at state w? Recalling that the independence indication '/1' is to express that in order for the formula to be true, it must be possible to select a witness for $\Diamond$ independently of the state chosen to interpret $\Box$, we see that $\Box(\Diamond/1)p$ is true in $\mathfrak{M}$ at w iff*

(1) the basic modal formula $\Box\Diamond p$ is true in $\mathfrak{M}$ at w, and

(2) the witness of $\Diamond$ in $\Box\Diamond p$ can be chosen uniformly with respect to the state interpreting $\Box$.

To illustrate, consider two pointed models $(\mathfrak{M}, a)$ and $(\mathfrak{M}', a')$ which both satisfy condition (1), i.e., render $\Box\Diamond p$ true:

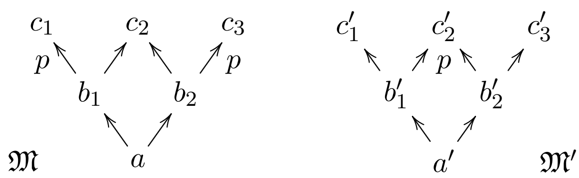

The former fails to satisfy the additional condition (2). For, when $\Box\Diamond p$ is being evaluated on $(\mathfrak{M}, a)$, the possible states interpreting $\Box$ are b_1 and b_2. In the former case the accessible states are c_1 and c_2, in the latter c_2 and c_3. If the witness for $\Diamond$ is to be chosen uniformly, it must be a state accessible both from b_1 and from b_2 and so it must actually be c_2. But if it is also to render the atom p true, it must be either c_1 or c_3. Since no state satisfies both conditions, the formula $\Box(\Diamond/1)p$ is not true in $\mathfrak{M}$ at a. On the other hand, this formula is true in $\mathfrak{M}'$ at a'. Here the diamond token may be witnessed by the state c_2', which renders the atom p true. Indeed, the formula $\Box(\Diamond/1)p$ claims of its state of evaluation w that there exists a state u that is a common successor to all states v accessible from w, and that u makes p true.

For every triple $(\phi, \mathfrak{M}, w)$, where ϕ is an MsL formula and $(\mathfrak{M}, w)$ is a pointed model, we associate a game between two players, **E** and **A**, denoted $G(\phi, \mathfrak{M}, w)$. Player **A** will make choices corresponding to tokens of $\Box$ and $\wedge$ — operators with universal force[1] — while **E** will make choices corresponding to the tokens of $\Diamond$ and $\vee$ — operators with existential force. In order to get an idea of how such games are played, let us take examples of games played relative to the pointed model $(\mathfrak{M}, a)$ of Example 2.2.

EXAMPLE 2.3. *Let us begin by considering an* ML *formula. How is the game for the formula* $\Box\Diamond p$ *played on* $(\mathfrak{M}, a)$*? A 'terminal play'* (a, x, y) *consists of some state* x *accessible from state* a*, chosen by* **A** *corresponding to* $\Box$*, followed by some state* y *accessible from state* x*, chosen by* **E** *and corresponding to* $\Diamond$*.* **E** *wins the play if the state* y *makes the atom* p *true, otherwise* **A** *wins. There are four possible terminal plays:* (a, b_1, c_1) *and* (a, b_1, c_2) *and* (a, b_2, c_2) *and* (a, b_2, c_3)*. There exists a winning strategy for player* **E***: choose* c_1 *if* **A** *chose* b_1*, and choose* c_3 *if* **A** *chose* b_2*.*

What about the game played on the same pointed model but with the MsL *formula* $\Box(\Diamond/1)p$*? The set of possible terminal plays is the* same *as in the case of the* ML *formula* $\Box\Diamond p$*. Also the winning conditions for terminal plays are the same. The difference is that here a restriction is imposed on available winning strategies: a strategy is winning for player* **E** *only if it is uniform in the sense of assigning the same state to* $(\Diamond/1)$ *regardless of the state chosen for* $\Box$*. No uniform winning strategy exists: the only strategy which yields corresponding to* $(\Diamond/1)$ *a state that is available at both states* b_1 *and* b_2 *consists of choosing the state* c_2 *in both cases. But* p *is false at* c_2*.*

We introduce the set of *positions* of game $G(\phi, \mathfrak{M}, w)$ by laying down the relevant game rules. At each position at most one player has to make a move. Depending on the position, he or she must either make a *model move* or a *junction move.* In the former case the player must choose from the domain a state meeting a certain additional condition. In the latter case the choice is syntactic: one of the terms of a junction is chosen. Sequences of positions generated according to the game rules are *plays.* Plays that cannot be further extended are called *terminal* plays. For terminal plays we define conditions under which a player comes out as the winner. By stipulation $G(\phi, \mathfrak{M}, w)$ is a zero-sum game: who does not win, loses.

DEFINITION 2.4 (Positions). *The positions of game* $G(\phi, \mathfrak{M}, w)$ *are generated thus:*

[1] Recall that the negation sign may only appear before an atom.

- *The initial position is* $(\phi, \emptyset, w)$.
- *If the position* $((\psi \vee \chi), \vec{b}, a)$ *has been reached, also* $(\psi, \vec{b}0, a)$ *and* $(\chi, \vec{b}1, a)$ *are positions.*[2] *Player* **E** *chooses one of them.*
- *If the position* $((\psi \wedge \chi), \vec{b}, a)$ *has been reached, also* $(\psi, \vec{b}0, a)$ *and* $(\chi, \vec{b}1, a)$ *are positions. Player* **A** *chooses one of them.*
- *If the position* $((\Diamond/i_1, \ldots, i_k)\psi, \vec{b}, a)$ *has been reached and* $R(a, c)$, *then* $(\psi, \vec{b}, c)$ *is a position. If* a *is* R*-maximal,* $(\psi, \vec{b}, \mathsf{fail})$ *is a position.* **E** *makes the choice. If* $(\psi, \vec{b}, \mathsf{fail})$ *is chosen,* **A** *wins the resulting play.*
- *If the position* $(\Box\psi, \vec{b}, a)$ *has been reached and* $R(a, c)$, *then* $(\psi, \vec{b}, c)$ *is a position. If* a *is* R*-maximal,* $(\psi, \vec{b}, \mathsf{fail})$ *is a position. It is player* **A** *who chooses. If* $(\psi, \vec{b}, \mathsf{fail})$ *is chosen,* **E** *wins the resulting play.*
- *If* ℓ *is a literal and the position* $(\ell, \vec{b}, a)$ *has been reached, then* **E** *wins the play that led to this position if* ℓ *is true at* a *according to the valuation function of the model, else* **A** *wins.*

A position encodes the information about which subformula token has been reached, as well as what the most recent model move has been. Observe that the independence indications play no role whatsoever in the clause for strings of the form $(\Diamond/i_1, \ldots, i_k)\psi$: the rule for the slashed diamond is perfectly symmetric to that of the box. The independence indications will have an effect elsewhere: they regulate winning strategies available to player **E**. For future use, we stipulate that *substrings* of formula ϕ are those expressions that appear as leftmost components of a position of a game $G(\phi, \mathfrak{M}, w)$. We call a string of symbols an MsL *string* if it is a substring of some MsL formula. Not all MsL strings are MsL formulas; e.g., $(\Diamond/27)q$ is not one. If $(\psi, \vec{b}, a)$ is a position, the pair $(\psi, \vec{b})$ identifies a specific *token* of the substring type ψ; indeed the rationale for including the binary tuples $\vec{b}$ in the positions is precisely to keep track of substring tokens (which is important, as we will see, in MsL). The third component of a position $(\psi, \vec{b}, a)$ is either a state belonging to the domain of the relevant model or the symbol fail.

Let us move on to define the notion of strategy and explain what it takes for a strategy to be 'uniform' and 'winning.'

DEFINITION 2.5 (Strategy, extension of a strategy). *If* $\phi \in$ MsL *and* $\alpha_1, \ldots, \alpha_n$ *is a list of all diamond and disjunction tokens appearing in*

[2] We assume that 0 and 1 are objects not belonging to the set M.

ϕ, a strategy of player **E** *in game $G(\phi, \mathfrak{M}, w)$ is a tuple $(f_1, \ldots, f_n)$ of functions, called* strategy functions*. For each i and every position at which a move corresponding to α_i must be made, the strategy function f_i provides a move that respects the game rules, depending on the previous moves made in the play leading to that position. (Additional requirements imposed on a strategy may, in general, cancel out some of those dependencies; the uniformity requirement discussed below is a case in point.)*

By fixing a strategy f, player **E** *leaves open certain plays and excludes others: precisely those are left open that player* **A** *can realize by making suitable moves, given that* **E** *follows the strategy f. Call the set of plays left open by the choice of f the* extension *of f, denoted Ext(f).*

We generalize the notion of strategy as follows. If $(f_{i_1}, \ldots, f_{i_k})$ is the restriction of a strategy $f = (f_1, \ldots, f_n)$ to the set Ext(f) — that is, if $(f_{i_1}, \ldots, f_{i_k})$ is the list of all functions among $f_1, \ldots, f_n$ whose all arguments are plays in Ext(f) — then also the tuple $(f_{i_1}, \ldots, f_{i_k})$ is a strategy*.*

We note that a strategy f in the narrower sense may include strategy functions that never need be used, since the plays in which they would be used lie outside the extension of f. This is what motivates the more general notion of strategy.

Strategy functions are functions that map non-terminal plays to positions. Often it will be more convenient to consider functions which take simply tuples of states as their arguments and which return as values states (case of boxes and diamonds) or objects encoding junction moves (case of junctions). We recall that 0 encodes the choice of the left term of a junction and 1 the choice of its right term. From a suitable function of the latter type it will be straightforward to define a strategy function proper. When no confusion threatens, we will simply take the latter types of functions to be strategy functions.

A strategy f of player **E** is *uniform*, if for every diamond token $(\Diamond/i_1, \ldots, i_k)$ in ϕ, the strategy satisfies the following: whenever two plays that belong to the extension of f lead to a position where a move for this token must be made, and those plays differ at most in the choices made for the preceding modal operator tokens identified by the numerals $i_1, \ldots, i_k$, then the strategy yields the same choice in both cases. For instance, **E**'s strategy f in the game associated with $\Box(\Diamond/1)p$, played relative to the state w, is uniform if there is a state u such that f yields the choice u for $(\Diamond/1)$ no matter which state v accessible from w is chosen by player **A** for $\Box$.

In order for a strategy of player **E** to be *winning*, it must satisfy: (1) all terminal plays in the extension of f are won by **E**, and (2) f

is uniform. The notion of winning strategy of player **A** is otherwise similar but involves no requirement of uniformity. This is because the syntax precludes boxes with independence indications.

DEFINITION 2.6 (Truth, falsity). *Let ϕ be an* MsL *formula. We say that ϕ is* true *in $\mathfrak{M}$ at w, if there exists a winning strategy for* **E** *in game $G(\phi, \mathfrak{M}, w)$. And we say that ϕ is* false *in $\mathfrak{M}$ at w, if there exists a winning strategy for* **A** *in game $G(\phi, \mathfrak{M}, w)$. We write $\mathfrak{M}, w \models \phi$ to indicate that ϕ is true in $\mathfrak{M}$ at w.*

As the following example illustrates, the formula ϕ and the pointed model $(\mathfrak{M}, w)$ can be so chosen that there exists no winning strategy for either player in game $G(\phi, \mathfrak{M}, w)$, i.e., the game is non-determined in the sense of game theory. In such a case we will derivatively say that ϕ is *non-determined* in $\mathfrak{M}$ at w. Observe that the non-determinacy of a formula in $\mathfrak{M}$ at w amounts to its being neither true nor false in $\mathfrak{M}$ at w, given that 'true' and 'false' are understood in the sense of Definition 2.6. It follows from the existence of such formulas that the non-truth of a formula in a pointed model does not guarantee its falsity therein.

EXAMPLE 2.7. *Let us return to the formula $\phi := \Box(\Diamond/1)p$ and the pointed model $(\mathfrak{M}, a)$ of Example 2.3. We already saw that ϕ is not true in $\mathfrak{M}$ at a. We will establish that it is not false in $\mathfrak{M}$ at state a either. Note that there are exactly two possible strategies for* **A** *in the game $G(\phi, \mathfrak{M}, a)$, namely choosing b_1 or choosing b_2 to interpret the box. In the former case* **E** *may choose c_1 and win, while in the latter case she may choose c_3 and win. Therefore no strategy of* **A** *is winning and ϕ is not false in $\mathfrak{M}$ at a.*

We note that since there are formulas non-determined in some pointed models, truth-equivalence of ϕ and ψ does not in general guarantee that these formulas are also false in precisely the same models: there can be models in which one is false and the other is non-determinate. Let us take still further examples of evaluating MsL formulas. These examples will utilize the models $\mathfrak{M}_1$ and $\mathfrak{M}_2$:

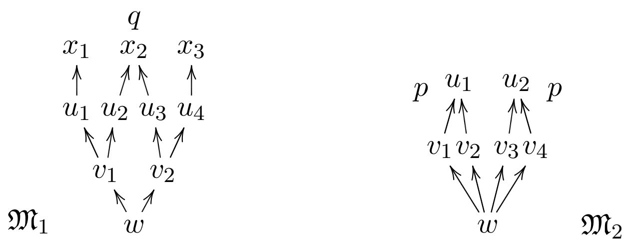

EXAMPLE 2.8. *Consider evaluating the formula $\Box\Diamond(\Diamond/1, 2)q$ in the model $\mathfrak{M}_1$ at state w. Let us first note that the following pair $f :=$*

(f_1, f_2) yields a strategy for **E** *in the relevant game, leading against both possible moves by* **A** *to a terminal play won by* **E***:*

- $f_1(v_1) = u_2$ *and* $f_1(v_2) = u_3$.
- $f_2(v_1, u_2) = x_2 = f_1(v_2, u_3)$.

The functions f_1 and f_2 provide moves that indeed are available according to the game rules, and both terminal plays in the extension of the strategy f are won by **E**. *Second, we may observe that this strategy is uniform: the plays in the extension of f at which a choice for $(\Diamond/1, 2)$ is made are (w, v_1, u_2) and (w, v_2, u_3), and the strategy f indeed maps them both to the same state, viz. x_2. We may conclude that $\mathfrak{M}_1, w \models \Box\Diamond(\Diamond/1, 2)q$.*

EXAMPLE 2.9. *Consider the model $\mathfrak{M}_2$ depicted above. Obviously the formula $\Box(\Diamond/1)p$ considered in Examples 2.3 and 2.7 is not true in $\mathfrak{M}_2$ at w: that would require that there be a common successor to all the states v_1, v_2, v_3, v_4 making p true, and this is not the case. However — and here we see a rather unusual feature of* MsL *— the formula* $\Box((\Diamond/1)p \vee (\Diamond/1)p)$ is *true in $\mathfrak{M}_2$ at w: replacing the string $(\Diamond/1)p$ by the string $((\Diamond/1)p \vee (\Diamond/1)p)$ produces a true formula out of a non-true one. Let us see what is behind this phenomenon.*

Here is a winning strategy $f := (f_1, f_2, f_3)$ for **E** *in the game corresponding to the latter formula, played on the pointed model $(\mathfrak{M}_2, w)$.*

- $f_1(w, v_1) = 0 = f_1(w, v_2)$ *and* $f_1(w, v_3) = 1 = f_1(w, v_4)$
- $f_2(w, v_1, 0) = u_1 = f_2(w, v_2, 0)$
- $f_3(w, v_3, 1) = u_2 = f_3(w, v_4, 1)$.

Not only do the functions f_1, f_2 and f_3 yield a strategy for **E** *such that all plays in its extension are won by* **E***, but also the strategy is uniform. For, consider the plays in the extension of f at which a move must be made corresponding to the leftmost token of $(\Diamond/1)$, i.e., the plays $(w, v_1, 0)$ and $(w, v_2, 0)$. Evidently f maps these plays to the same state, viz. u_1. Similarly, the plays $(w, v_3, 1)$ and $(w, v_4, 1)$ in the extension of f at which a move is to be made for the rightmost token of $(\Diamond/1)$ are likewise mapped to the same state, that is, to u_2. When player* **E** *is making her move for the slashed diamond, she is so to say aware of which one of the two tokens of the slashed diamond she is reacting to, and nothing prevents her from acting differently in the two cases. Considering the totality of the plays, using her preceding disjunctive choice,* **E** *has been able to partition the set of immediate successors* *of the state w in two pieces: those states associated with* 0 *and those associated with* 1*. When it is* **E***'s turn to move for a diamond token, she only needs to be able to provide a common successor to those states which have been associated with the same number,* 0 *or* 1*.*

Example 2.9 shows that if α is a substring of an MsL formula, substituting $(\alpha \vee \alpha)$ for α in that formula is not in general a truth-value preserving operation. However, it is not difficult to see that the phenomenon may only occur if α is not an MsL formula, i.e., only if α is a substring of a formula which cannot itself be produced by the syntax of MsL.

3. The expressivity of MsL

Many properties of MsL are known from the literature. Since we will be interested in a fragment of MsL — the logic MsL_0 to be introduced in Section 4 — and as we wish to contrast this fragment with various logics, among them the logic MsL itself, it is in order to register here some of these properties, notably those concerning to the relative expressive power of MsL.

FACT 3.1 (Tulenheimo 2003). MsL *is more expressive than* ML.

Proof. First, ML is a syntactic fragment of MsL and so the game-theoretic semantics of MsL yields in particular a certain truth condition for each formula of ML. Assuming the Axiom of Choice, the standard semantics of ML and its game-theoretic semantics coincide. The assumption is needed in the direction from the ML semantics to the game-theoretical semantics on models having states with an infinite out-degree: the truth of ϕ in the standard sense only guarantees the existence of a *non-deterministic* strategy in the corresponding game, and for the possibility of turning a non-determinisitc strategy into a deterministic one (as required by the MsL semantics), we need the Axiom of Choice.[3] We conclude that $\mathsf{ML} \leq \mathsf{MsL}$. In order to see that $\mathsf{MsL} \nleq \mathsf{ML}$, consider the pointed models $(\mathfrak{M}, a)$ and $(\mathfrak{M}', a')$ of Example 2.2. They are plainly bisimilar and therefore ML equivalent. Yet, as noted in Example 2.2, we have $\mathfrak{M}, a \not\models \Box(\Diamond/1)p$ but $\mathfrak{M}', a' \models \Box(\Diamond/1)p$. ■

[3] It was observed in (Hodges, 1983, p. 94) and in (Hintikka and Kulas, 1985, pp. 6–7) that one must assume the Axiom of Choice in order to prove the equivalence of the standard Tarskian semantics of first-order logic and its game-theoretical semantics. Hodges (2006) noted that if the weaker notion of strategy is applied — namely that of non-deterministic strategy — we obtain an unconditional correspondence between the Tarskian and the game-theoretical semantics for first-order logic.

It is well-known that ML enjoys the tree-model property: every satisfiable ML formula is true at the root of a tree-like model. The logic MsL, again, lacks this property.

FACT 3.2. MsL *does not have the tree-model property.*

Proof. Consider the MsL formula $\phi := ((\Diamond p \wedge \Diamond \sim p) \wedge \Box(\Diamond/1)\top)$. It is satisfiable: e.g., the pointed model $(\mathfrak{M}, a)$ makes it true:

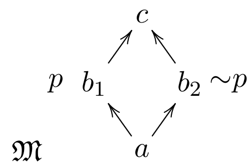

No pointed model $(\mathfrak{N}, v)$ making ϕ true is tree-like: the out-degree of v is at least 2 and the successors of v have a common successor. ■

Just like there exists a well-known straightforward translation of ML into FOL — for this so-called *standard translation*, see e.g. (Blackburn *et al.*, 2002, Def. 2.45) — there exists a straightforward translation of MsL into FOsL. The following examples suffice for seeing how the translation works.

MsL *formula*	*its standard translation*
$\Box(\Diamond/1)p$	$\forall x(\mathsf{R}x_0x \rightarrow (\exists y/x)(\mathsf{R}xy \wedge \mathsf{P}y))$
$\Box((\Diamond/1)q \vee (\Diamond/1)q)$	$\forall x(\mathsf{R}x_0x \rightarrow ((\exists y/x)(\mathsf{R}xy \wedge \mathsf{Q}y) \vee (\exists z/x)(\mathsf{R}xz \wedge \mathsf{Q}z)))$
$\Box\Box(\Diamond/1)(\Diamond/1,3)p$	$\forall x(\mathsf{R}x_0x \rightarrow \forall y(\mathsf{R}xy \rightarrow (\exists z/y)(\mathsf{R}yz \wedge (\exists v/x,z)(\mathsf{R}zv \wedge \mathsf{P}v))))$

For every $\phi \in$ MsL, let $ST_{x_0}(\phi)$ be its standard translation in FOsL obtained in this way, with x_0 as its sole free variable. By the semantics of the two logics, it is evident that $\mathfrak{M}, w \models \phi$ iff $\mathcal{M}, w \models ST_{x_0}(\phi)$.

FACT 3.3. MsL $\leq$ ESO. *That is, for every* $\phi \in$ MsL *there is an* ESO *formula* $\psi_\phi(x)$ *of one free variable such that for every pointed model* $(\mathfrak{M}, w)$*, the following holds:* $\mathfrak{M}, w \models \phi$ *iff* $\mathcal{M}, w \models \psi_\phi$.

Proof. It is well known that FOsL $\leq$ ESO; cf., e.g., (Väänänen, 2007, Subsect. 6.1). Since MsL $\leq$ FOsL, the statement follows. ■

It is quite obvious that Fact 3.3 can be improved: MsL $<$ ESO. Actually already for the fragment FOL of ESO we have that FOL $\not\leq$ MsL. We will obtain a strict proof of this fact in Section 9 (Corollary 9.8).

Alternatively, this could be proven for example by observing that MsL, like ML, is invariant under generated submodels whereas FOL is not.

Basic modal logic is translatable into FOL. Could we perhaps improve Fact 3.3 to the point of showing that actually MsL is translatable already into FOL? The answer is in the negative.

PROPOSITION 3.4 (Tulenheimo & Sevenster 2007). MsL $\not\leq$ FOL.

The idea of proof used in (Tulenheimo and Sevenster, 2007) was as follows. A certain MsL formula χ and a family $\{(\mathfrak{M}_n, a) : n \geq 2\}$ of pointed models were discerned. Then it was shown (1) that $\mathfrak{M}_n, a \models \chi$ iff n is even, and (2) that for every $n \geq 0$ there is a number $k_n \geq 2$ such that the structures $(\mathcal{M}_{k_n}, a)$ and $(\mathcal{M}_{k_n+1}, a)$ are elementarily equivalent up to quantifier rank n. It then follows that χ has no translation into FOL. For suppose it had one, say $\phi_\chi(x)$, and let r be the quantifier rank of $\phi_\chi(x)$. Then $(\mathcal{M}_{k_r}, a)$ and $(\mathcal{M}_{k_r+1}, a)$ would be elementarily equivalent up to quantifier rank r and in particular not distinguished by $\phi_\chi(x)$. But exactly one of the numbers k_r and $k_r + 1$ is even, so $(\mathfrak{M}_{k_r}, a)$ and $(\mathfrak{M}_{k_r+1}, a)$ *are* distinguished by χ. Therefore $(\mathcal{M}_{k_r}, a)$ and $(\mathcal{M}_{k_r+1}, a)$ are distinguished by the translation $\phi_\chi(x)$ of χ. This is a contradiction. In the proof χ was chosen to be the MsL formula $\Box(\Box(\Diamond/1)(\Diamond/1,3)\top \ \vee \ \Box(\Diamond/1)(\Diamond/1,3)\top)$.

The following result concerning contradictory negation follows from Proposition 3.4. In Section 9 this negative result will be improved.

COROLLARY 3.5. MsL *is not closed under contradictory negation. That is, it is not the case that for every* $\phi \in$ MsL *there is* $\psi_\phi \in$ MsL *such that for all pointed models* $(\mathfrak{M}, w)$*, we have:* $\mathfrak{M}, w \models \psi_\phi$ *iff* $\mathfrak{M}, w \not\models \phi$.

Proof. It is well known that contradictory negation is inexpressible in ESO in the following strong sense: if ξ and ζ are two ESO formulas such that for all models $\mathcal{M}$ and all variable assignments γ, we have $\mathcal{M}, \gamma \models \xi$ iff $\mathcal{M}, \gamma \not\models \zeta$, then actually each of ξ and ζ is logically equivalent to an FOL formula; for a proof, see e.g. (Barwise, 1979, pp. 56, 73–74). Now, the MsL formula χ used in the proof of Proposition 3.4 can be translated into ESO but not into FOL. By the fact just mentioned, its contradictory negation has no translation into ESO, and so *a fortiori* not into MsL. ■

4. The fragment MsL_0

We discern a certain simple fragment of MsL, to be dubbed MsL_0. The syntax of MsL_0 is obtained from that of MsL by allowing only formulas containing diamond symbols of the forms $(\Diamond/\emptyset)$ or $(\Diamond/1)$, where '1' identifies a syntactically preceding *box* token. Accordingly, for example $\Box(\Diamond/1)p$ and $\Box(p \vee (\Diamond\Box(\Diamond/1)q \wedge (\Diamond/1)r))$ are formulas, while $\Diamond(\Diamond/1)p$ and $\Box\Diamond(\Diamond/1,2)q$ are not. The logic MsL_0 is a notational variant of the logic $\mathcal{L}_{\mathbf{SD}}$ that was formulated and studied in (Tulenheimo and Sevenster, 2006).

In the present and the five following sections we study the expressive power of MsL_0. In particular, we obtain a model-theoretic criterion for telling when an MsL_0 formula cannot be translated into basic modal logic (the present section); we offer a normal form result for **E**'s winning straregies in MsL_0 games (Section 5); we prove that the semantics of MsL_0 is 'quasi-local' in a sense to be specified (Section 6); we will see that MsL_0 is translatable into FOL but neither to any finite variable fragment of FOL nor to the loosely guarded fragment of FOL (Section 7); it is observed that up to truth-equivalence there are infinitely many MsL_0 formulas of a given modal depth, but only finitely many MsL_0 formulas with a fixed modal depth and a fixed so-called breadth (Section 8); finally, we show that — in a very strong sense — MsL_0 fails to be closed under contradictory negation (Section 9).

4.1. Syntax recursively defined

Consider the following grammars A (the former) and B (the latter); note that A is defined with reference to B:[4]

$$\alpha \;::=\; p \mid\; \sim p \mid (\alpha \vee \alpha) \mid (\alpha \wedge \alpha) \mid \Diamond\alpha \mid \Box\alpha \mid \Box\beta$$

$$\beta \;::=\; (\Diamond/1)\alpha \mid (\alpha \star \beta) \mid (\beta \star \alpha) \mid (\beta \star \beta),$$

with $\star \in \{\vee, \wedge\}$ and $p \in \mathbf{prop}$. The formulas of MsL_0 are recursively generated — as a moment's reflection reveals — by the grammar A. The strings (sequences of symbols) generated by the grammar B are *not* formulas of MsL. However, they are available as building blocks when producing formulas. Indeed, they are substrings of formulas in the sense explained in Section 2.2. Note that for any string β there is a tuple $(s_1, \ldots, s_{n+1})$ of *terms* such that β is obtained from those terms — using them, say, in the order from s_1 up to s_{n+1} — by n applications of $\wedge$ and $\vee$, each term s_i being either an MsL_0 formula or a string of the form $(\Diamond/1)\alpha$, where α is an MsL_0 formula. We say that

[4] Recall that $\Diamond$ abbreviates $(\Diamond/\emptyset)$.

terms of the former kind are *formula terms* and those of the latter kind *non-formula terms*. For later purposes, let us introduce the syntactic notions of *degree* and *breadth* of a formula.

DEFINITION 4.1 (Degree, breadth). *If β is a string produced by the grammar B, its* degree*, denoted* $\deg(\beta)$*, is the number of its non-formula terms. In other words,* $\deg(\beta)$ *is the number of those tokens of the symbol* $(\Diamond/1)$ *in β that are not syntactically subordinate to a token of $\Box$ in β. Given a formula ϕ of* MsL_0*, let $\mathbb{S}_\phi$ be the set of its subformulas (subformula types) generated by the grammar B. By definition, the* breadth *of ϕ, denoted* $\mathrm{bre}(\phi)$*, equals* $\max\{\deg(\beta) : \beta \in \mathbb{S}_\phi\}$.

We will call diamond symbols $\Diamond$ *basic diamonds* and diamond symbols $(\Diamond/1)$ *slashed diamonds*. When reasoning about **E**'s strategies in MsL_0 games, we will make use of the following further definitions.

DEFINITION 4.2 (Contextual vs. basic disjunction tokens). *Suppose $\phi \in \mathsf{MsL}_0$. If $(\chi_1 \vee \chi_2)$ is a substring of ϕ and both of its terms χ_1 and χ_2 are* formulas*, then the token of $\vee$ is* basic*. If, again, at least one of its terms is a non-formula, the token of $\vee$ is* contextual*. That is, the token of $\vee$ is contextual if there is in $(\chi_1 \vee \chi_2)$ a token of slashed diamond not preceded in $(\chi_1 \vee \chi_2)$ by a box token.*

For example, consider the formula $(p \vee \Box((q \vee (r \wedge (\Diamond/1)p')) \wedge (q' \vee r')))$. Counting from left to right, the first token of $\vee$ is basic: p and $\Box((q \vee (r \wedge (\Diamond/1)p')) \wedge (q' \vee r'))$ are formulas. Also the third token of $\vee$ is basic: q' and r' are formulas. However, the second token is contextual: the substring $(r \wedge (\Diamond/1)p')$ is not a formula.

DEFINITION 4.3 (Recollection, recollection state). *Let $\phi \in \mathsf{MsL}_0$. If $((\Diamond/1)\chi, \vec{c}, a)$ is the last position of a play π of game $G(\phi, \mathfrak{M}, w)$, there is in π a unique earlier position of the form $(\Box\beta, \vec{d}, b)$ with $R(b, a)$ such that $(\Diamond/1)\chi$ is a term of β. Similarly, if $((\chi_1 \vee \chi_2), \vec{c}, a)$ is the last position of a play π with $\vee$ contextual, there is in π a unique earlier position of the form $(\Box\beta, \vec{d}, b)$ with $R(b, a)$ such that χ_1 and χ_2 are Boolean combinations of terms of β. We say that the position $(\Box\beta, \vec{d}, b)$ is the* recollection *of the slashed diamond token $(\Diamond/1)$ in π and the element b of M its* recollection state *in π; respectively we say that $(\Box\beta, \vec{d}, b)$ is the* recollection *of the relevant contextual disjunction token in π and b its* recollection state *in π.*

4.2. Criterion for non-properness

A formula of MsL_0 is *proper* if it is not truth-equivalent to any ML formula. We establish a model-theoretic criterion enabling us to tell when a formula is *not* proper. This result will be of use in Section 9.

DEFINITION 4.4 (Global duplication). *Let* $(\mathfrak{M}, w)$ *be a pointed* tree-like *model of finite height. Write* n *for its height.*

(a) Let I *be any subset of* M*. For each* $i \in I$*, fix a set* K_i *of indices; the set* K_i *may have any cardinality. For every* $i \in I$*, suppose that* $J_i = \{(\mathfrak{M}^i_j, w^i_j) : j \in K_i\}$ *is a set of pointed tree-like models having no states in common with* $\mathfrak{M}$ *and having pairwise disjoint domains. Write* $\mathcal{F}$ *for the family* $\{J_i : i \in I\}$*. We define the* sum *of* $(\mathfrak{M}, w)$ *and the family* $\mathcal{F}$*, denoted* $(\mathfrak{M}, w) \bigoplus_I \mathcal{F}$*, to be the tree-like model whose root is* w*, whose domain is* $M \cup \bigcup_{i \in I} \bigcup_{j \in K_i} M^i_j$*, whose accessibility relation is* $R \cup \bigcup_{i \in I} \bigcup_{j \in J_i} R^i_j \cup \{(i, w^i_j) : j \in K_i \text{ and } i \in I\}$ *and whose valuation assigns to every atom* p *the set* $V(p) \cup \bigcup_{i \in I} \bigcup_{j \in J_i} V^i_j(p)$*.*

(b) Let $0 \leq m < n$ *and let* L_m *be the set of states of* $\mathfrak{M}$ *of layer* m *(i.e., the set of states of height* m*). For every state* v *in* L_m *and every* R*-successor* u *of* v*, let* $(\mathfrak{M}^{v,u}, u')$ *be an isomorphic copy of the submodel of* $\mathfrak{M}$ *generated by* u*, having no nodes in common with* $\mathfrak{M}$*. Write* J_v *for the set* $\{(\mathfrak{M}^{v,u}, u') : R(v, u)\}$ *and* $\mathcal{F}$ *for the family* $\{J_v : v \in L_m\}$*. By definition the* duplication of $\mathfrak{M}$ at layer m *equals* $(\mathfrak{M}, w) \bigoplus_{L_m} \mathcal{F}$*. (There is nothing to duplicate at layer* n *— leaves have no successors — hence we only consider values of* m *smaller than* n*.)*

(c) Let $\mathcal{S}_0$ *be the duplication of* $(\mathfrak{M}, w)$ *at layer 0. Further, if* $1 \leq r < n - 1$*, then let* $\mathcal{S}_{r+1}$ *be the duplication of* $\mathcal{S}_r$ *at layer* $r + 1$*. By definition the* global duplication of $(\mathfrak{M}, w)$ *equals* $\mathcal{S}_{n-1}$*.*

Let $(\mathfrak{N}, w)$ be the global duplication of a tree-like pointed model $(\mathfrak{M}, w)$ of finite height, and let R be the accessibility relation of $\mathfrak{N}$. Suppose u is an R-successor of v. Then, due to the way in which $\mathfrak{N}$ was constructed, there is a state u' with $u \neq u'$ such that $R(v, u')$ and the submodel of $\mathfrak{N}$ generated by u is *isomorphic* to the submodel of $\mathfrak{N}$ generated by u'. Further, suppose $\mathcal{K}$ is a class of pointed models closed under global duplications. Let k be any positive integer. If $(\mathfrak{M}, w) \in \mathcal{K}$ is tree-like and of finite height, there is a pointed model $(\mathfrak{N}, w) \in \mathcal{K}$ with the following property: if v and u are states in N, u accessible from v, there are $2^k - 1$ further states in N, likewise accessible from v, such that the submodel of $\mathfrak{N}$ generated by any of these $2^k - 1$ states is isomorphic to the submodel generated by u. Such a pointed model — to be termed the *k-fold duplication* of $(\mathfrak{M}, w)$ — is obtained by applying k times successively the operation of global duplication starting from $(\mathfrak{M}, w)$.

To facilitate discussion, we stipulate that if a pointed model of finite height is *not* tree-like, it is its own global duplication. In this way the operation of global duplication is defined for all pointed models of finite height. We write $Mod(\phi)$ for the class of all pointed models in which ϕ is true and the height of which is at most $md(\phi)$. Nothing of interest is left outside the class $Mod(\phi)$: the class of *all* pointed models of ϕ is evidently the class of those pointed models whose restriction to height $md(\phi)$ belongs to $Mod(\phi)$. The *n-unraveling* of a pointed model is the restriction of its unraveling to height n.

THEOREM 4.5. *Let* $\phi \in \mathsf{MsL}_0$. *If the class* $Mod(\phi)$ *is closed under* $md(\phi)$*-unraveling and global duplication, then* ϕ *is not proper.*

Proof. Let k be a positive integer such that $2^k \geq \text{bre}(\phi)+1$. Suppose $Mod(\phi)$ is closed under $md(\phi)$-unraveling and global duplication. If $Mod(\phi)$ is empty, trivially ϕ is not proper. So suppose there is a pointed model $(\mathfrak{M}, w)$ in $Mod(\phi)$. By assumption the $md(\phi)$-unraveling $(\mathfrak{M}', w)$ of $(\mathfrak{M}, w)$ belongs to $Mod(\phi)$ as well. Now, the $md(\phi)$-unraveling is of finite height, so there is in $Mod(\phi)$ also the k-fold duplication $(\mathfrak{N}, w)$ of $(\mathfrak{M}', w)$.

Let f be **E**'s winning strategy in game $G(\phi, \mathfrak{N}, w)$. Since ϕ is a formula of MsL_0, whenever a position of the form $((\Diamond/1)\psi, \vec{c}, a)$ occurs in a play of this game, its recollection $(\Box\beta, \vec{d}, b)$ occurs in that play as well. Without loss of generality, we may assume that the strategy f avoids positions $((\Diamond/1)\psi, \vec{c}, a)$ whenever possible. What this means can be explained in other words as follows: if the strategy nevertheless leads to such a position against some sequence of moves by **A**, then all *formula terms* χ of the corresponding Boolean combination β are *non-true* at a. Now, given that f is a winning strategy and avoids positions of the form $((\Diamond/1)\psi, \vec{c}, a)$ whenever possible, it actually follows that no position of the form $((\Diamond/1)\psi, \vec{c}, a)$ is reached in the first place. In order to prove this we make essential use of the fact that $(\mathfrak{N}, w)$ is a k-fold duplication of a tree-like pointed model of ϕ.

Suppose for contradiction that following f against some sequence of moves by **A** leads to a position $((\Diamond/1)\psi_0, \vec{c}_0, a_0)$. If $(\Box\beta, \vec{d}, b)$ is the recollection of the token of $(\Diamond/1)$, we know that no formula term of β is true at a_0. By the construction of $\mathfrak{N}$, there are at least bre(ϕ) further states $a_1, \ldots, a_{\text{bre}(\phi)}$, all accessible from b, such that each of them generates a submodel of $\mathfrak{N}$ isomorphic to the submodel of $\mathfrak{N}$ generated by a_0. Due to isomorphism, we may conclude that no formula term of β is true at any of those further states either. Therefore, the winning strategy f must lead against any moves by **A** to such disjunctive choices that for each of those further states a_i, a position $((\Diamond/1)\psi_i, \vec{c}_i, a_i)$ is reached, where $(\Diamond/1)\psi_i$ is a non-formula term of β. Since there are at most $\text{bre}(\phi)$ such non-formula terms to assign to the $\text{bre}(\phi)+1$ states $a_0, a_1, \ldots, a_{\text{bre}(\phi)}$, one of those non-formula terms must be utilized twice. That is, there are i and j with $a_i \neq a_j$ such that $\psi_i = \psi_j$ (meaning that ψ_i and ψ_j are the same string *token*). Since f is a winning strategy, it will map a_i and a_j to their common successor. But $\mathfrak{N}$ is tree-like, so no two states a_i, a_j with $R(b, a_i)$ and $R(b, a_j)$ have a common successor. This is a contradiction. We have just shown that no winning strategy of **E** in game $G(\phi, \mathfrak{N}, w)$ will yield any position of the form $((\Diamond/1)\psi, \vec{c}, a)$.

Now, let ϕ' be the formula obtained from ϕ by replacing by $\bot$ its all substrings of the form $(\Diamond/1)\psi$ which do not lie in ϕ in the syntactic scope of a further diamond symbol $(\Diamond/1)$. Note that by syntactic criteria, ϕ' is a formula of ML. It remains to show that ϕ and ϕ' are truth-equivalent. Trivially, whenever ϕ' is true in $\mathfrak{M}$ at w, also ϕ is true in $\mathfrak{M}$ at w. Namely, the formula ϕ' determines a certain subtree of the syntactic tree of ϕ. The existence of a winning strategy for **E** corresponding to ϕ' guarantees that she can remain within that subtree when playing the game corresponding to ϕ. For any sequence of moves by **A** in game $G(\phi, \mathfrak{M}, w)$, player **E** may respond by the move that her winning strategy yields for the same sequence of moves in game $G(\phi', \mathfrak{M}, w)$; in this way **E** will in particular never end up in a position involving a slashed diamond in game $G(\phi, \mathfrak{M}, w)$. In order to see that conversely, the truth of ϕ in $\mathfrak{M}$ at w suffices for the truth of ϕ' in $\mathfrak{M}$ at w, we may reason as follows. Suppose $\mathfrak{M}, w \models \phi$. Construct from $\mathfrak{M}$ a model $\mathfrak{N}$ as above. Given that $Mod(\phi)$ is closed under $md(\phi)$-unraveling and global duplication, we have $\mathfrak{N}, w \models \phi$. Since **E**'s winning strategy in $G(\phi, \mathfrak{N}, w)$ does not yield any positions of the form $((\Diamond/1)\psi, \vec{c}, a)$, evidently ϕ' is true in $\mathfrak{N}$ at w. Now, observe that $(\mathfrak{M}, w)$ and $(\mathfrak{N}, w)$ are $md(\phi)$-bisimilar. Since $md(\phi') \leq md(\phi)$, we may conclude that ϕ' is true in $\mathfrak{M}$ at w. ■

5. MsL_0 and bounded quasi-positionality

We will next prove a result which serves to deepen our understanding of the semantics of MsL_0. We show that the value of a winning strategy for player **E** at a given play only depends on certain earlier positions

— in fact it depends on at most two earlier positions. In what follows, if π and σ are any sequences, $\pi^\frown\sigma$ stands for the sequence obtained by concatenating π and σ (in this order). We need to agree on some further terminology.

DEFINITION 5.1 (Bounded quasi-positionality). *Let $\phi \in \mathsf{MsL}_0$ and let α be a diamond or disjunction token in ϕ.*

(a) *A strategy function for α is* positional *if its value on a play π depends only on the last model move in π. If α is a slashed diamond, a strategy function for α is* quasi-positional *if its value on π depends only on the recollection state of α in π. And if α is a contextual disjunction token, a strategy function for α has* bounded quasi-positionality *if its value depends only on the following two positions: the last model move of π and the recollection state of α in π.*

(b) *A sequence of strategy functions (in particular, an entire* strategy*) of player* **E** *has* bounded quasi-positionality *if its strategy functions for basic diamonds and basic disjunctions are positional, its strategy functions for slashed diamonds are quasi-positional, and its strategy functions for contextual disjunctions have bounded quasi-positionality.*

(c) *Derivatively, a fragment $\mathcal{L}$ of MsL_0 is said to have* bounded quasi-positionality*, if the following condition holds: for all formulas $\phi \in \mathcal{L}$ and pointed models $(\mathfrak{M}, w)$, if there exists a winning strategy for* **E** *in $G(\phi, \mathcal{M}, w)$, then in this game there exists a winning strategy for* **E** *which has bounded quasi-positionality.*

As the following example illustrates, in MsL_0 we cannot generally hope that the value of a winning strategy for **E** on a play would depend only on one single model move therein.

EXAMPLE 5.2. *Consider the MsL_0 formula $\phi := \Box\Box((\Diamond/1)p \vee (\Diamond/1){\sim}p)$ and the model $\mathfrak{M}$ depicted below.*

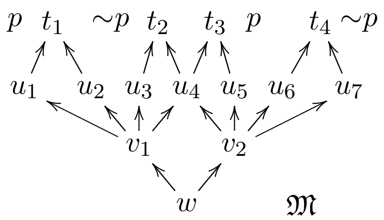

(1) We note that $\mathfrak{M}, w \models \phi$. Indeed, define the functions f, g and h as follows: $f(v_1, u_1) = f(v_1, u_2) = f(v_2, u_4) = f(v_2, u_5) = 0$, $f(v_1, u_3) =$

$f(v_1, u_4) = f(v_2, u_6) = f(v_2, u_7) = 1$, $g(v_1) = t_1$, $g(v_2) = t_3$, $h(v_1) = t_2$, $h(v_2) = t_4$. The functions f, g and h induce strategy functions for the disjunction, the left slashed diamond and the right slashed diamond, respectively; obviously these strategy functions constitute a winning strategy for $\mathbf{E}$ in game $G(\phi, \mathfrak{M}, w)$. We note that the function f yielding the strategy function for the disjunction uses two arguments. (2) We proceed to prove that if d is a strategy function for the disjunction in $G(\phi, \mathfrak{M}, w)$ and the value of d at a play depends (i) only on $\mathbf{A}$'s move for the first occurrence of $\Box$ or (ii) only on $\mathbf{A}$'s move for the second occurrence of $\Box$, then d is not $\mathbf{E}$'s strategy function for the disjunction in any winning strategy.

Suppose d, g, h are functions inducing a winning strategy for $\mathbf{E}$ in $G(\phi, \mathfrak{M}, w)$, with g and h of type $\{v_1, v_2\} \to \{t_1, t_2, t_3, t_4\}$, corresponding to the left and the right slashed diamond respectively. Case (i): suppose d is a function of type $\{v_1, v_2\} \to \{0, 1\}$. If $d(v_1) = 0$, the function g maps in particular both plays (w, v_1, u_2) and (w, v_1, u_3) to the state $g(v_1)$. Since g belongs to a winning strategy, the state $g(v_1)$ is accessible from u_2 and from u_3, which is impossible. Similar reasoning applies if $d(v_1) = 1$, this time using the function h. Case (ii): suppose d is a function of type $\{u_1, \ldots, u_7\} \to \{0, 1\}$, i.e., suppose d is positional. Suppose $d(u_4) = 0$. Since g yields a strategy function for the left slashed diamond and g belongs to a winning strategy, g must map each of the plays (w, v_1, u_4) and (w, v_2, u_4) to a state at which p is true — not because of the uniformity requirement, but because the state chosen must make p true. Hence in fact $g(v_1) = t_3 = g(v_2)$. Now, the state t_3 is not accessible from any of the states u_1, u_2, u_3, u_6, u_7. So d cannot map any of these five states to 0. Thus, in particular $d(u_1) = d(u_3) = 1$. But then h maps both of the plays (w, v_1, u_1) and (w, v_1, u_3) to the state $h(v_1)$. Since h yields a strategy function for the right slashed diamond and h belongs to a winning strategy, $h(v_1)$ is accessible from both u_1 and u_3, which is impossible as the two states do not have a common successor. Similar reasoning applies if $d(u_4) = 1$. We conclude that any strategy function for the disjunction in ϕ belonging to a winning strategy must be a function of two arguments.

Directly by the definition of bounded quasi-positionality, we have:

FACT 5.3. *Let $\phi, \psi \in \mathsf{MsL}_0$ and suppose f is a winning strategy for player $\mathbf{E}$ in game $G(\phi, \mathfrak{M}, w)$ such that f has bounded quasi-positionality. If the position $(\psi, \vec{c}, v)$ appears in a play belonging to the extension of f, then there is a winning strategy for $\mathbf{E}$ in game $G(\psi, \mathfrak{M}, v)$, in fact a winning strategy having bounded quasi-positionality.*

We prove that the whole MsL_0 has bounded quasi-positionality in the sense of Definition 5.1(c). By Example 5.2 this is an optimal result:

we could not always find a winning strategy whose all strategy functions are functions of one argument only. For a less transparent proof of a more general result, see (Tulenheimo, 2009), Appendix §3, Lemma 3.

THEOREM 5.4. MsL_0 *has bounded quasi-positionality.*

Proof. Suppose there is a winning strategy for **E** in game $G(\phi, \mathfrak{M}, w)$, call it f. We will define a certain function g and show that it is a winning strategy for **E** in $G(\phi, \mathfrak{M}, w)$ with bounded quasi-positionality.

Let Π be the set of all plays of $G(\phi, \mathfrak{M}, w)$ at which it is **E**'s turn to move. Before doing anything else, with each play π in Π we associate a tuple P_π of moves in π as follows. (Depending on how the last position of π looks like, this tuple will contain either one or two members.) Suppose the last position of π is $(\psi, \vec{c}, a)$.

- If $\psi = \Diamond\chi$, or $\psi = (\chi_1 \vee \chi_2)$ with $\vee$ basic, then $P_\pi := (a)$.
- If $\psi = (\Diamond/1)\chi$ and b is the recollection state of this token of $(\Diamond/1)$, then $P_\pi := (b)$.
- If $\psi = (\chi_1 \vee \chi_2)$, with $\vee$ contextual, and b is the recollection of this token of $\vee$, then $P_\pi := (b, a)$.

Then define an equivalence relation $\approx$ on Π by setting $\pi \approx \pi'$ iff (π and π' are of the same length, involve the same junction moves, and satisfy $P_\pi = P_{\pi'}$). With these definitions at our disposal, let us proceed to define the function g. We wish to define g in such a way that whenever $\sigma, \sigma' \in Ext(g)$ and $\sigma \approx \sigma'$, there is a play $\pi \in Ext(f)$ with $\sigma \approx \pi \approx \sigma'$ such that $g(\sigma) = f(\pi) = g(\sigma')$. We must of course check that a play $\pi \in Ext(f)$ with $\sigma \approx \pi \approx \sigma'$ can always be found. We define g together with an auxiliary function $h : Ext(g) \cap \Pi \to Ext(f)$ which indeed provides such a play: for any σ on which h is defined, $\sigma \approx h(\sigma)$ and $h(\sigma) = h(\sigma')$ for any $\sigma \approx \sigma'$. Once the function g with the desired properties is defined, we show that it is actually a winning strategy for **E**; its quasi-positionality follows automatically from its definition.

First, if $S \in \Pi/\!\approx$ and all plays in S contain only moves by **A**, select arbitrarily a play $\pi \in S$ and set $h(\sigma) := \pi$ for all $\sigma \in S$. Trivially π belongs to the extension of f. Define $g(\sigma) := (f \circ h)(\sigma)$, for all $\sigma \in S$. Next suppose

$$\sigma = \tau^\frown g(\tau)^\frown \rho$$

is a play in the set Π and constructed using g, where ρ contains only moves by **A** and there is $\pi \in Ext(f)$ such that for all τ' constructed using g and satisfying $\tau' \approx \tau$, we have: $\tau' \approx \pi$ and $h(\tau') = \pi$ and $g(\tau') = (f \circ h)(\tau')$. We wish to find $\pi_0 \in Ext(f)$ such that for all plays

σ' constructed using g with $\sigma' \approx \sigma$, we have $\pi_0 \approx \sigma'$. Once such a play is found we define $h(\sigma') := \pi_0$ and we set $g(\sigma') := (f \circ h)(\sigma')$ for all σ' constructed using g with $\sigma' \approx \sigma$. We will consider cases according to the form of the string ψ at the last position $(\psi, \vec{c}, a)$ of the sequence $g(\tau)^\frown\rho$. (Note that ρ may be empty.) Now, by assumption $\tau \approx h(\tau)$. Since $h(\tau)$ is in $Ext(f)$, also $h(\tau)^\frown(f \circ h)(\tau)^\frown\rho = h(\tau)^\frown g(\tau)^\frown\rho$ is in $Ext(f)$. Write $\pi^* := h(\tau)^\frown g(\tau)^\frown\rho$. We will show that actually $\sigma \approx \pi^*$.

1. The case $\psi := \Diamond\chi$ or $\psi := (\chi_1 \vee \chi_2)$ with basic $\vee$: Since the plays σ and π^* share their last member, we have that $\sigma \approx \pi^*$.

2. The case: $\psi := (\Diamond/1)\chi$. We consider two subcases according to whether the recollection of the slashed diamond token occurs in the sequence τ or in the sequence $g(\tau)^\frown\rho$. (2.1) If it occurs in the sequence $g(\tau)^\frown\rho$, then it is immediate $\sigma \approx \pi^*$. (2.2) If, again, the recollection $(\Box\beta, \vec{d}, b)$ of $(\Diamond/1)$ with $R(b, a)$ occurs in the sequence τ, then we may observe that the substring component of the last position of τ must be a contextual diamond. But then, since $\tau \approx h(\pi)$, the plays τ and $h(\tau)$ share the position $(\Box\beta, \vec{d}, b)$, and it follows that also the plays σ and $h(\tau)^\frown g(\tau)^\frown\tau$ share the position $(\Box\beta, \vec{d}, b)$. That is, $\sigma \approx \pi^*$.

3. The case $\psi := (\chi_1 \vee \chi_2)$ with contextual $\vee$: Again we consider two subcases, according to where the recollection $(\Box\beta, \vec{d}, b)$ of the slashed disjunction token lies; note that here $R(b, a)$. (3.1) If it occurs in the sequence $g(\tau)^\frown\rho$, then the plays π^* and σ share both the position $(\Box\beta, \vec{d}, b)$ and their last model moves, namely the position $(\beta, \vec{d}, a)$ and we have $\sigma \approx \pi^*$. (3.2) If, again, the recollection $(\Box\beta, \vec{d}, b)$ of the diamond token occurs in the sequence τ, we may observe that also the substring component of the last position of τ must be a contextual diamond. But then, since $\tau \approx h(\pi)$, the plays τ and $h(\tau)$ share both the position $(\Box\beta, \vec{d}, b)$ and the position $(\beta, \vec{d}, a)$, and it follows that also the plays σ and π^* share these positions. That is, $\sigma \approx \pi^*$.

Let us now proceed to show that the function g is a winning strategy for player **E**. First we observe that for every play $\sigma \in \Pi$ on which g is defined, the move $g(\sigma)$ respects the game rules, i.e., the sequence $\sigma^\frown g(\sigma)$ is a play. In the case that the last position of σ carries a disjunction (whether basic or contextual) this is trivial, and in the case its last position carries a diamond (whether basic or slashed) this follows immediately from the definition of g. Second, if $\tau^\frown g(\tau)^\frown\rho$ is a terminal play in the extension of g, where ρ contains only moves by **A**, this play is won by $\mathbf{E}$ because the play $h(\tau)^\frown g(\tau)^\frown \rho$ in the extension of f is won by $\mathbf{E}$ (the condition of winning depends only on the last position of a terminal play). Third, we note that g respects the required uniformity conditions. Actually, if $\sigma \approx \sigma$, the function g agrees on these plays. This is more than enough for respecting the corresponding uniformity constraint: the latter would allow the value of the strategy on σ to depend on any model move preceding the earliest model move on which the value of g on σ depends. (That is, the choice made for basic and contextual disjunctions and basic diamonds could depend on any earlier moves, and the choice made for slashed diamonds could depend, in addition to its recollection, on any model move preceding the recollection.) ■

Let us still introduce some more terminology.

DEFINITION 5.5 (Forwards-looking truth-condition). *Let ϕ be a formula of* MsL *(not necessarily a formula of* MsL_0*).*

(1) Suppose f is a strategy for $\mathbf{E}$ in game $G(\phi, \mathfrak{M}, w)$ and $\mathcal{S}$ is a set of positions at which it is $\mathbf{E}$'s turn to move and which appear in the extension of f. We say the set $\mathcal{S}$ is forwards-looking for *f provided that we have: if $\sigma := \pi^\frown(\psi, \vec{c}, v)^\frown\rho$ is any play in Ext(f) on which f is defined and for which $(\psi, \vec{c}, v) \in \mathcal{S}$, the value of f on σ does not depend on any position in π.*

(2) We say that ϕ has a strongly *(respectively,* weakly*)* forwards-looking truth-condition*, if for all $\mathfrak{M}$ and $w \in M$, the following holds: if $\mathfrak{M}, w \models \phi$, there exists a winning strategy g for $\mathbf{E}$ in $G(\phi, \mathfrak{M}, w)$ such that the set $\mathcal{P}_g^{\mathfrak{M},w}$ (respectively $\mathcal{S}_g^{\mathfrak{M},w}$) is forwards-looking for g, where $\mathcal{P}_g^{\mathfrak{M},w}$ is the set of* all *positions in Ext(g) at which it is $\mathbf{E}$'s turn to move, and $\mathcal{S}_g^{\mathfrak{M},w}$ is the subset of $\mathcal{P}_g^{\mathfrak{M},w}$ consisting of positions $(\psi, \vec{c}, v)$ in which ψ is a* formula *(and not merely a substring).*

OBSERVATION 5.6. *By Theorem 5.4, all formulas of* MsL_0 *have a weakly forwards-looking truth-condition, though not in general a strongly forwards-looking one. The formulas of* ML *indeed have a strongly forwards-looking truth-condition.*

When discussing the possibility for a language to have a compositional semantics, Hintikka has stressed the property of *semantic context-independence*: semantic attributes of a complex expression depend only on the semantic attributes of its constituent expressions, plus its structure — not on the sentential context in which the expression is embedded (Hintikka and Kulas, 1983; Hintikka, 1996). As subsequent work by Hodges in effect highlights, a language — Hodges discusses precisely first-order slash logic — may enjoy semantic context-independence relative to one type of semantic attribute while failing to enjoy this property relative to another type of semantic attribute. He (Hodges, 1997a; Hodges, 1997b) showed that first-order slash logic actually *does* admit of a compositional semantics, given that formulas are evaluated relative to *sets* of variable assignments, instead of being evaluated relative to one assignment only. On the other hand, Cameron and Hodges (2001, Cor. 6.2) proved that no 'Tarskian' compositional semantics exists for first-order slash logic, interpreting a string $\phi(x)$ by a subset of M^n for some $n < \omega$ rather than by a subset of $Pow(M)$. That is, any compositional semantics of first-order slash logic is bound to employ sets of assignments rather than single assignments. Hintikka used to take it for granted that context-independence unavoidably fails for first-order slash logic (IF first-order logic). This is, then, correct if attention is confined to the semantic attribute *satisfaction under an assignment*, but not if the attribute *satisfaction under a set of assignments* is considered. Even without resorting to the non-trivial result of Cameron and Hodges, it is plain that at least there is no obvious way in which the semantics of expressions such as $(\exists x/\forall y)\phi$ could have a context-independent semantics in the 'Tarskian' sense. Actually, from the viewpoint of game-theoretical semantics, such strings do not by themselves have any satisfaction conditions at all, so in particular not satisfaction conditions that could be used for determining whether a larger formula is satisfied. For a discussion on compositionality and first-order slash logic (IF first-order logic), see (Sandu and Hintikka, 2001).

In connection with the semantic attribute of *truth*, the above notion of strongly forwards-looking truth-condition provides a formulation of the notion of semantic context-independence utilizing the notion of strategy. The fact that MsL_0 formulas have a weakly forwards-looking truth-condition suggests that through a suitable syntactic reformulation we might be able to turn MsL_0 into an equivalent language whose formulas have a strongly forwards-looking truth-condition and indeed manifest semantic context-independence. Such a language will actually be generated by the grammar C, which we now move on to introduce.

6. Grammar C and the quasi-locality of the semantics

6.1. Grammar C

Let us take another look at strings β generated by the grammar B introduced in Subsection 4.1. There are $n^2 \cdot (n-1)!$ trees with n nodes of out-degree 2, the rest being of out-degree 0, such that every node of out-degree 2 is labeled either with $\vee$ or with $\wedge$. We may call such trees *binary syntactic structures of size* n. In such a tree the number of leaves is $n+1$. We may take the nodes to be binary strings, the root of the tree being the empty string and the successors of a node $\vec{b}$ of out-degree 2 being $\vec{b}0$ and $\vec{b}1$. The leaves may be enumerated by stipulating that $\vec{b}_1$ precedes $\vec{b}_2$ iff in the first position in which these strings differ, $\vec{b}_1$ carries 0 while $\vec{b}_2$ carries 1. Binary syntactic structures of size n can be understood as $(n+1)$-ary connectives. When such a connective C_n is applied to a tuple of $n+1$ syntactic objects $\xi_1, \ldots, \xi_{n+1}$, the result is one syntactic object, namely the binary syntactic structure of size n whose leaf ℓ_i is labeled with ξ_i for every $1 \leq i \leq n+1$. We may write $C_n(\xi_1, \ldots, \xi_{n+1})$ for the resulting syntactic object. If for example $\xi_1, \xi_2, \xi_3, \xi_4$ are formulas of propositional logic and C_3 is a binary syntactic structure of size 3, then $C_3(\xi_1, \xi_2, \xi_3, \xi_4)$ is the syntactic tree of a formula of propositional logic obtained from $\xi_1, \xi_2, \xi_3, \xi_4$ by exactly 3 applications of connectives from the set $\{\vee, \wedge\}$. Now, for every string β produced by grammar B there is a positive interger n, a set $S \subseteq \{1, \ldots, n+1\}$ and a binary syntactic structure $C_{n,S}$ of size n such that $C_{n,S}(\xi_1, \ldots, \xi_{n+1})$ is the syntactic tree of β, given that for all $1 \leq i \leq n+1$, we have: $\xi_i \in \mathsf{MsL}_0$ if $i \notin S$, while ξ_i is a string of the form $(\Diamond/1)\phi_i$ with $\phi_i \in \mathsf{MsL}_0$ if $i \in S$.

If $\phi_1, \ldots, \phi_{n+1}$ are formulas of MsL_0, define $D_{n,S}(\phi_1, \ldots, \phi_{n+1}) := C_{n,S}(\xi_1, \ldots, \xi_{n+1})$, where $\xi_i := \phi_i$ if $i \notin S$ and $\xi_i := (\Diamond/1)\phi_i$ if $i \in S$. For example, $D_{2,\{1,3\}}(p, q, r) = C_{2,\{1,3\}}((\Diamond/1)p, q, (\Diamond/1)r)$. That is, $D_{n,S}$ is a binary syntactic structure of size n which, when applied to an $(n+1)$-tuple of MsL_0 formulas, returns a syntactic tree whose leaves with an order position in S are labeled with strings of the form $(\Diamond/1)\phi$, while leaves having their order position outside S are labeled with MsL_0 formulas. Labeled trees such as those denoted by expressions $D_{n,S}(\phi_1, \ldots, \phi_{n+1})$ can be identified with the corresponding MsL_0 strings. Let us define $E_{n,S}(\phi_1, \ldots, \phi_{n+1}) := \Box D_{n,S}(\phi_1, \ldots, \phi_{n+1})$. Any MsL_0 formula of the form $\Box\beta$, with β generated by grammar B, has as its notational variant an expression of the form $E_{n,S}(\phi_1, \ldots, \phi_{n+1})$, the ϕ_i being MsL_0 formulas. The following grammar C yields an alternative notation for the syntax of MsL_0:

$$\phi \quad ::= \quad p \mid {\sim}p \mid (\phi \vee \phi) \mid (\phi \wedge \phi) \mid \Diamond\phi \mid \Box\phi \mid E_{n,S}(\phi_1, \ldots, \phi_{n+1}).$$

Here $p \in \mathbf{prop}$ and $E_{n,S} = \Box D_{n,S}$, with $D_{n,S}$ a binary syntactic structure of size n for some $n \geq 1$ and $S \subseteq \{1, \ldots, n+1\}$. The notion of immediate subformula of a formula generated by the grammar C is defined in the obvious way. Syntactically formulas generated by C can be considered in their own right, but at the same time they serve as a metalogical notation denoting MsL_0 formulas. Continuing with the above example, the formula $E_{2,\{1,3\}}(p, q, r)$ equals by definition $\Box D_{2,\{1,3\}}(p, q, r)$ which equals by definition $\Box C_{2,\{1,3\}}((\Diamond/1)p, q, (\Diamond/1)r)$. While the immediate subformulas of $E_{2,\{1,3\}}(p, q, r)$ are p, q and r, it stands for an MsL_0 formula which is syntactically composed of expressions not all of which are MsL_0 formulas. In particular the Boolean combination denoted by $C_{2,\{1,3\}}((\Diamond/1)p, q, (\Diamond/1)r)$ is not an MsL_0 formula.

6.2. Quasi-locality

Let ϕ be a formula generated by grammar C. The semantics of ϕ is said to be *quasi-local* provided that the following condition holds. If $\mathfrak{M} = (M, R, V)$ is any model and $w \in M$, then in order to determine whether ϕ is true in $\mathfrak{M}$ at w, it is sufficient to consider the set $\{x : x = v$ or $R(v, x)$ or $(R \circ R)(v, x)\}$ of states and to know for every immediate subformula ψ of ϕ and all states u in that set whether ψ is true in $\mathfrak{M}$ at u. Quasi-locality amounts to semantic context-independence relative to the semantic attribute of truth at states at most 2 steps away from the current state along the relevant accessibility relation. If it is actually enough to consider the smaller set $\{x : x = v$ or $R(v, x)\}$ of states, the semantics of ϕ is *local*. This is the case with formulas ϕ of ML. The following result is a consequence of Theorem 5.4.

COROLLARY 6.1. *The semantics of all formulas generated by grammar C is quasi-local.*

Proof. Suppose there is a map F such that for all immediate subformulas ψ of ϕ and all states u in the set $\{x : x = v$ or $R(v, x)$ or $(R \circ R)(v, x)\}$, we have $F(\psi, u) = 1$ if $\mathfrak{M}, u \models \psi$ and $F(\psi, u) = 0$ otherwise. We show that the map F allows us to determine whether $\mathfrak{M}, w \models \phi$. If $(\psi, \vec{b}, u)$ is a position of game $G(\phi, \mathfrak{M}, w)$ belonging to the extension of $\mathbf{E}$'s strategy g, we say $(\psi, \vec{b}, u)$ is a *g-winning position for* $\mathbf{E}$, if player $\mathbf{E}$ wins all terminal plays in $Ext(g)$ containing this position. (We note that a strategy f is winning iff the initial position of the game is an f-winning position, but a position may be g-winning without g being a winning strategy.)

Let us consider different cases according to the form of ϕ. The case of literals is trivial, and so is the case of junctions. If $\phi := \Box\chi$, then

$\mathfrak{M}, w \models \phi$ iff for some winning strategy g of $\mathbf{E}$ in $G(\phi, \mathfrak{M}, w)$ having bounded quasi-positionality we have that for all u with $(w, u) \in R$, the position $(\chi, \emptyset, u)$ is a g-winning position for $\mathbf{E}$ iff $F(\chi, u) = 1$ for all u with $(w, u) \in R$. In the first equivalence, the direction from left to right holds by Theorem 5.4, while in the second equivalence, the direction from left to right holds because the truth-condition of ϕ is weakly forwards-looking and ϕ is a formula. The case of $\phi := \Diamond\chi$ is perfectly analogous. Finally, consider the case $\phi := E_{n,S}(\chi_1, \ldots, \chi_{n+1}) = \Box D_{n,S}(\chi_1, \ldots, \chi_{n+1})$. Recall how the binary syntactic structure $D_{n,S}$ was defined from a certain binary syntactic structure $C_{n,S}$, the latter being applied to an $(n + 1)$-tuple of strings such that the strings in places $i \in S$ are of the form $(\Diamond/1)\chi_i$ and those in places $i \notin S$ are of the form χ_i. Write m for the number of conjunction symbols in the binary syntactic structure $C_{n,S}$. Now, $\mathfrak{M}, w \models \phi$ iff for some winning strategy g of $\mathbf{E}$ in $G(\phi, \mathfrak{M}, w)$ having bounded quasi-positionality, there are states $t_1, \ldots, t_{|S|}$ such that for all u with $R(v, u)$ and every tuple $\vec{c} \in \{0, 1\}^m$ of choices for conjunctions, there are choices for disjunctions $\vec{d}$ such that the substring token χ_i corresponding to the tuple $\vec{b}$ determined by these binary choices $\vec{c}, \vec{d}$ satisfies: if $i \notin S$, then $(\chi_i, \vec{b}, u)$ is a g-winning position for $\mathbf{E}$ — whereas if $i \in S$, we have that $R(u, t_i)$ and $(\chi_i, \vec{b}, t_i)$ is a g-winning position for $\mathbf{E}$ iff there are states $t_1, \ldots, t_{|S|}$ such that for all u with $R(v, u)$ and every tuple $\vec{c} \in \{0, 1\}^m$ of choices for conjunctions, there are choices for disjunctions $\vec{d}$ such that the substring token χ_i corresponding to these binary choices $\vec{c}, \vec{d}$ satisfies: if $i \notin S$, then $F(\chi_i, u) = 1$ — whereas if $i \in S$, we have $R(u, t_i)$ and $F(\chi_i, t_i) = 1$. Again, Theorem 5.4 grants the left-to-right direction of the first equivalence; in the second equivalence, the direction from left to right uses the fact that the truth-condition of ϕ is weakly forwards-looking and the χ_i are formulas. ∎

Note that Corollary 6.1 is not literally about MsL_0; it is about formulas generated by grammar C. Thanks to this result we may, derivatively, say that the semantics of MsL_0 itself is quasi-local. What this really means is that when MsL_0 formulas are suitably syntactically analyzed — when they are considered as generated by grammar C — suitable subformulas can be identified so that in terms of the semantic values of those subformulas (truth or non-truth relative to certain nearby states), the semantic value of a given formula at a given state can be determined.

7. MsL_0 compared with FOL

7.1. Translatability into FOL

Trivially $\mathsf{MsL}_0 \leq \mathsf{MsL}$. Unlike MsL itself, MsL_0 can actually be translated into FOL. Let us first take an example.

EXAMPLE 7.1. *Consider the* MsL_0 *formula* $\Box(\Diamond/1)p$. *Its standard translation is* $\theta_1 := \forall x(Rx_0x \to (\exists y/x)(Rxy \wedge Py))$. *This* FOsL *formula is satisfaction-equivalent to the* FOL *formula* $\theta_2 := \exists y \forall x(Rx_0x \to (Rxy \wedge Py))$, *obtained from* θ_1 *by erasing the expression* $(\exists y/x)$ *and placing the quantifier* $\exists y$ *in front of* $\forall x$.

In order to see that θ_1 *and* θ_2 *are indeed satisfaction-equivalent, suppose first that* (f, g) *is* **E***'s winning strategy in game* $G(\theta_1, \mathcal{M}, w)$. *Then in particular* g *can be taken to be a function of type* $M \to M$ *which is uniform in its sole argument. If the constant value of* g *equals* c, *it is obvious that* (f, c) *is* **E***'s winning strategy in game* $G(\theta_2, \mathcal{M}, w)$. *Conversely, if* (f, c) *is* **E***'s winning strategy in game* $G(\theta_2, \mathcal{M}, w)$, *letting a map* g *be defined by putting* $g(a) = c = g(a')$ *for all* $a, a' \in M$, *we have that* (f, g) *is* **E***'s winning strategy in game* $G(\theta_1, \mathcal{M}, w)$.

Similarly it is seen that the MsL_0 *formula* $\Box((\Diamond/1)p \vee (\Diamond/1)p)$ *admits of an* FOL *translation: its standard translation* $\forall x(Rx_0x \to ((\exists y/x)(Rxy \wedge Py) \vee (\exists z/x)(Rxz \wedge Pz))$ *is satisfaction-equivalent to the* FOL *formula* $\exists y \exists z \forall x(Rx_0x \to ((Rxy \wedge Py) \vee (Rxz \wedge Pz))$.

It was proven in (Tulenheimo and Sevenster, 2006) that MsL_0 is translatable into FOL, by utilizing a suitable skolemization procedure; in the outputs of the procedure all function symbols were nullary, i.e., the outputs were notational variants FOL formulas not using function symbols. We prove the fact here by more direct means.

THEOREM 7.2 (Tulenheimo & Sevenster 2006). $\mathsf{MsL}_0 \leq \mathsf{FOL}$.

Proof. Let us define a map $T_{x_0} : \mathsf{MsL}_0 \to \mathsf{FOL}$ recursively, relative to the syntax given by grammar C discerned in Section 6. If y is a variable and the map T_{x_0} is defined on ϕ, we write $T_{x_0/y}(\varphi)$ for the result of having first changed, if necessary, variables in $T_{x_0}(\phi)$ so that in the resulting formula, x_0 does not appear free in the syntactic scope of any quantifier $\mathsf{Q}y$, and having then substituted y for x_0 in that resulting formula. For literals, disjunctions, conjunctions and formulas with the prefix $\Diamond$ or $\Box$, let the map T be defined in the same way as the standard translation of basic modal logic. Finally, let $T_{x_0}[E_{n,S}(\phi_1, \ldots, \phi_{n+1})] :=$

$$\exists x_{i_1} \ldots \exists x_{i_{|S|}} \forall x(\mathsf{R}x_0x \to C_{n,S}(\xi_1, \ldots, \xi_{n+1})),$$

where $\xi_j = (\mathsf{R}xx_j \wedge T_{x_0/x_j}[\phi_j])$ if $j \in S$, and $\xi_j = T_{x_0/x}[\phi_j]$ if $j \notin S$, given that the map $j \mapsto i_j$ is an enumeration of the elements of S. We claim: for all MsL_0 formulas ϕ, models $\mathfrak{M}$ and states a of M, there is a winning strategy for player **E** in game $G(\phi, \mathfrak{M}, a)$ iff $\mathcal{M}, a \models T_{x_0}[\phi]$.

The claim holds evidently for literals. It is also clear that it holds for formulas of the forms $\wedge, \vee, \Box, \Diamond$ if it holds for their immediate subformulas. (For reasons mentioned in the proof of Fact 3.1 above, we need to assume the Axiom of Choice in the direction from right to left.) As to the remaining case, let $\psi := E_{n,S}(\phi_1, \ldots, \phi_{n+1})$. Let $(\mathfrak{M}, a)$ be an arbitrary pointed model and assume inductively that the claim holds for the formulas ϕ_i. Suppose first that there is a winning strategy f for **E** in game $G(\psi, \mathfrak{M}, a)$. The strategy f induces a map $\sharp$ such that for any state b with $R(a, b)$ and any tuple $\vec{d} \in \{0,1\}^m$ of left/right choices for conjunctions in the tree $C_{n,S}$, a certain substring token $(\zeta_i, \vec{e}_i)$ is determined with $i = \sharp(f, \vec{d}, b)$ and $1 \leq i \leq n+1$ such that $\zeta_i = \phi_i$ if $i \notin S$ and $\zeta_i = (\Diamond/1)\phi_i$ if $i \in S$. Now, if indeed $i \notin S$, then ζ_i is itself a formula of MsL_0 and by Observation 5.6, there is a winning strategy for player **E** in game $G(\zeta_i, \mathfrak{M}, b)$. Then, by the inductive hypothesis we have $\mathcal{M}, b \models T_{x_0}[\zeta_i]$. If, again, we have $i \in S$, then ζ_i is a non-formula. Letting c_i with $R(b, c_i)$ be the state that the strategy f yields at the position $((\Diamond/1)\phi_i, \vec{e}, b)$ of game $G(\psi, \mathfrak{M}, a)$, we may by Observation 5.6 infer that there is a winning strategy for **E** in $G(\phi_i, \mathfrak{M}, c_i)$ and so by the inductive hypothesis that $\mathcal{M}, c_i \models T_{x_0}[\phi_i]$. It follows that $\mathcal{M}, a \models T_{x_0}[\psi]$, since $\mathcal{M}, \gamma \models (\mathsf{R}x_0x \to C_{n,S}(\xi_1, \ldots, \xi_{n+1}))$ with $\gamma(x_0) = a$ and $\gamma(x_i) = c_i$ for $i \in S$. Conversely, suppose $\mathcal{M}, a \models T_{x_0}[\psi]$. Then there are witnesses $c_{i_1}, \ldots, c_{i_{|S|}}$ such that for all b with $R(a, b)$ we have $\mathcal{M}, \gamma \models C_{n,S}(\xi_1, \ldots, \xi_n)$, with $\gamma(x) = b$ and $\gamma(x_{i_k}) = c_{i_k}$. Let $\vec{d} \in \{0,1\}^m$ be a tuple of conjunctive choices in the tree $C_{n,S}$. Then there is a corresponding substring token $(\xi_j, \vec{e}_j)$ such that $\mathcal{M}, \gamma \models \xi_j$. If $j \notin S$, by the inductive hypothesis there is a winning strategy h_j for **E** in game $G(\xi_j, \mathfrak{M}, b)$. If, again, $j \in S$, there is by the inductive hypothesis a winning strategy h_j for **E** in game $G(\xi_j, \mathfrak{M}, c_j)$. Define a strategy g for **E** in $G(\psi, \mathfrak{M}, a)$ as follows: if **A** chooses b for $\Box$ and he produces the string $\vec{c}$ of conjunctive choices, let **E** respond in such a way that the string $(\xi_j, \vec{e}_j)$ is reached with $j = \sharp(f, \vec{d}, b)$. If **E** arrives at a position $(\phi_j, \vec{e}_j, b)$ with $j \notin S$, let **E** continue by applying strategy h_j; if again she arrives at a position $((\Diamond/1)\phi_j, \vec{e}_j, b)$ with $j \in S$, let her pick the state c_j and then continue by applying the corresponding strategy h_j. Clearly g is winning for **E** in $G(\psi, \mathfrak{M}, a)$. ■

It was noted in connection with Fact 3.3 that FOL cannot be translated into MsL, though we postponed the proof (we will obtain one via Corollary 9.8). *A fortiori*, then, FOL cannot be translated into the fragment MsL_0 of MsL. It follows by Theorem 7.2 that $\mathsf{MsL}_0 < \mathsf{FOL}$. As to the relation of MsL_0 and MsL, we have:

COROLLARY 7.3. $\mathsf{MsL}_0 < \mathsf{MsL}$.

Proof. Trivially $\mathsf{MsL}_0 \leq \mathsf{MsL}$. By Theorems 3.4 and 7.2, MsL cannot be translated into FOL but MsL_0 can. So $\mathsf{MsL} \not\leq \mathsf{MsL}_0$. ■

7.2. Non-translatability into FOL^n

It is well known that ML can be translated into the two-variable frament FOL^2 of FOL. Now, MsL_0 can be translated into FOL; does perhaps some finite number of variables suffice for carrying out this translation? As it happens, the answer is in the negative.

THEOREM 7.4 (Tulenheimo & Sevenster 2006). *For all positive integers n, we have:* $\mathsf{MsL}_0 \not\leq \mathsf{FOL}^n$.

Proof. Let $m \geq 2$ be arbitrary and write ϕ_m for the formula

$$\Box(\underbrace{(\Diamond/1)\top \vee \ldots \vee (\Diamond/1)\top}_{m-1 \text{ times}}).$$

The formula ϕ_m can be translated into FOL^{m+1} by the formula $\chi_m :=$

$$\exists z_1 \ldots \exists z_{m-1} \forall y (Rx_0 y \rightarrow (Ryz_1 \vee \ldots \vee Ryz_{m-1})).$$

If $k \geq 1$, define a structure $\mathcal{N}_k = (N_k, R^{\mathcal{N}_k})$ by putting $N_k = \{w, v_1 \ldots, v_{2k}, u_1, \ldots, u_k\}$ and $R^{\mathcal{N}_k} = \{(w, v_i) : 1 \leq i \leq 2k\} \cup \bigcup_{1 \leq i \leq k} \{(v_i, u_i), (v_{k+i}, u_i)\}$. The structure $\mathcal{N}_k$ consists of three mutually disjoint layers: the root w, a layer of $2k$ elements and a layer of k elements. The second layer can be partitioned into cells of size two such that the elements of each cell have a common successor on the third layer. We observe that $\mathfrak{N}_k, w \models \phi_m$ iff $1 \leq k < m$. We proceed to prove that χ_m is not equivalent to any formula of FOL^m. To this end, we claim that the structures $(\mathcal{N}_{m-1}, w)$ and $(\mathcal{N}_m, w)$ are FOL^m equivalent. From this we will be able to conclude that χ_m is not equivalent to any formula of FOL^m, since $\mathcal{N}_{m-1}, w \models \chi_m$ but $\mathcal{N}_m, w \not\models \chi_m$. The proof uses a pebble game argument. For using pebble games $G^n_m(\mathcal{M}, \vec{a}, \mathcal{N}, \vec{b})$ to characterize equivalence of structures up to quantifier rank $\leq r$ relative to FOL^r, see, e.g., (Ebbinghaus and Flum, 1999; Väänänen, 1999).

The claim follows if we show that there is a winning strategy for *Duplicator* in the pebble game $G^m(\mathcal{N}_{m-1}, w, \mathcal{N}_m, w)$. Let the relevant m pairs of pebbles be $(\alpha_0, \beta_0), \ldots, (\alpha_{m-1}, \beta_{m-1})$, with α_0 placed initially on the root of $\mathcal{N}_{m-1}$ and β_0 on the root of $\mathcal{N}_m$. Clearly an optimal strategy for *Duplicator* would consist of considering successively some $m-1$ elements of the third layer from $\mathcal{N}_m$, say $u_1, \ldots, u_{m-1}$, and placing the pebble β_i on u_i; *Duplicator* may respond to these moves by placing the pebble α_i on the element u_i of $\mathcal{N}_{m-1}$. Should the players be granted one extra pair of pebbles, (α_m, β_m), *Spoiler* could place β_m on the second-layer element v_{2m} of $\mathcal{N}_m$, and to this *Duplicator* would have no response, since the third layer successor u_m of v_{2m} would have no pebble placed on it, but all third layer elements of $\mathcal{N}_{m-1}$ would already carry a pebble. We conclude that the structures $(\mathcal{N}_{m-1}, w)$ and $(\mathcal{N}_m, w)$ are FOL^m equivalent. ■

7.3. Non-translatability into LGF

Consider the result of translating ML into FOL. In the fragment in question any existential quantifier appears in a context $\exists y(Rxy \wedge \psi(y))$. We may say that the atomic formula Rxy is the *guard* of the quantifier $\exists y$. Similarly all universal quantifiers in the fragment are guarded by an atomic formula: they all appear in contexts of the form $\forall y(Rxy \rightarrow \psi(y))$. The guards have a double function. On the one hand they render the quantifiers of the relevant fragment relativized: they impose a condition that the value of the quantified variable must satisfy. On the other hand, they dictate a syntactic property that the matrix formula ψ must satisfy: its only free variable y appears free in the guard as well. In the *guarded fragment* (GF) this setting is generalized. In GF quantifiers may only appear in contexts of the form $Q\vec{y}(G \star \psi)$, where G is *atomic* and the free variables of ψ are among those of G. Let τ be a relational vocabulary. The syntax of $\mathsf{GF}[\tau]$ is specified by the following grammar:

$$\psi \quad ::= \quad x_1 = x_2 \mid R\vec{x} \mid \neg\psi \mid (\psi \wedge \psi) \mid (\psi \vee \psi) \mid Q\vec{x}(G \star \psi),$$

where $R \in \tau$ is an n-ary relation symbol from τ, $Q \in \{\forall, \exists\}$, the relativizer G is an atomic formula, $(Q, \star) \in \{(\forall, \rightarrow), (\exists, \wedge)\}$, x_1, x_2 are variables, $\vec{x}$ is a finite tuple of variables, and $\mathit{Free}(\psi) \subseteq \mathit{Free}(G)$. Let us mention also a generalization of GF due to van Benthem(1997), called the *loosely guarded fragment* (LGF). Also here it is required that all quantifiers have relativizers, but the relativizers need not be atomic. Further, a syntactic condition on the distribution of the variables of the tuple $\vec{y}$ over the relativizer G is imposed. The syntax of LGF is obtained from that of GF by keeping the definition otherwise intact but requiring only that G be a *conjunction* of atomic formulas (not necessarily a single atomic formula as in GF), and, moreover, stipulating that for every y_i in the tuple $\vec{y}$ and every $z \in \mathit{Free}(G)$, at least one conjunct of G contains both y_i and z.

As is obvious and well known, ML can be translated into the guarded fragment of first-order logic (GF). Andréka, van Benthem and Németi (1998) famously suggested that the distinguishing feature of 'modal fragments' of FOL is that quantifiers appear *guarded* in them. Notably, they proposed that this fact provides an 'explanation' of the decidability of basic modal logic. Indeed GF is decidable. The loosely guarded fragment (LGF) extends GF but remains decidable (van Benthem, 1997; Grädel, 1999a). So it is potentially capable of 'explaining' even a larger variety of modal decidability results than GF. Postponing comments on the computational properties of MsL_0 to Section 10, we proceed to prove that MsL_0 *cannot* be translated into LGF and consequently not into GF.

We need some definitions. If the α_i with $1 \leq i \leq m$ are atomic first-order formulas of some vocabulary τ and $(x_1, \ldots, x_s)$ is a tuple of variables, the conjunction $\alpha_1 \wedge \ldots \wedge \alpha_m$ is said to *guard* the tuple $(x_1, \ldots, x_s)$ if for all $1 \leq i, i' \leq s$ there is $1 \leq j \leq m$ such that x_i and $x_{i'}$ appear both in α_j. Given that $\mathcal{N}$ is a τ-structure, a subset X of N is said to be *loosely k-guarded* in $\mathcal{N}$, if for some $s \leq k$ there is an assignment $\gamma : \{x_1, \ldots, x_s\} \to N$ such that $X \subseteq Im(\gamma)$, and there are atomic formulas α_i with $1 \leq i \leq m$ of vocabulary τ, such that $\mathcal{N}, \gamma \models \alpha_1 \wedge \ldots \wedge \alpha_m$ and the conjunction $\alpha_1 \wedge \ldots \wedge \alpha_m$ guards the tuple $(x_1, \ldots, x_s)$. A set is *loosely guarded* if it is loosely k-guarded for some k. In order to formulate the notion of loosely guarded k-bisimulation, we adopt the following notation. If $\mathcal{F}$ is a set of finite partial τ-isomorphisms between τ-structures $\mathcal{M}$ and $\mathcal{N}$, and if $f : X \to Y$ is a map in the set $\mathcal{F}$ such that $X \subseteq M$ and $Y \subseteq N$ are loosely k-guarded sets, we write $\mathbf{Zig}(f, X)$ for the claim "for every loosely k-guarded set $X' \subseteq M$ there is in the set $\mathcal{F}$ a partial isomorphism $g : X' \to Y'$ such that f and g agree on $X \cap X'$." Similarly, we write $\mathbf{Zag}(f, Y)$ for the claim "for every loosely k-guarded set $Y' \subseteq N$ there is in the set $\mathcal{F}$ a partial isomorphism $g : X' \to Y'$ such that f^{-1} and g^{-1} agree on $Y \cap Y'$." A *loosely guarded k-bisimulation* between $\mathcal{M}$ and $\mathcal{N}$ is any non-empty set $\mathcal{F}$ of finite partial isomorphisms between $\mathcal{M}$ and $\mathcal{N}$ such that for all $f : X \to Y$ in $\mathcal{F}$ with X, Y loosely k-guarded sets, we have $\mathbf{Zig}(f, X)$ and $\mathbf{Zag}(f, Y)$. The *width* of an LGF formula is the maximum number of free variables (variable types) in its subformulas.

The notion of loosely guarded k-bisimulation offers, in particular, a criterion for equivalence of two structures $(\mathcal{M}, a)$ and $(\mathcal{N}, b)$ with respect to LGF formulas of one free variable and width at most k: if there is a loosely guarded k-bisimulation $\mathcal{F}$ between $\mathcal{M}$ and $\mathcal{N}$ with $\{(a, b)\} \in \mathcal{F}$, then for all LGF formulas ψ of width at most k and with exactly one free variable, x, we have: $\mathcal{M}, a \models \psi$ iff $\mathcal{N}, b \models \psi$. For this criterion and the requisite definitions presented above, cf. (Grädel, 1999b, Sect. 3).

THEOREM 7.5. $\mathsf{MsL}_0 \not\leq \mathsf{LGF}$.

Proof. Let R be a binary relation symbol. Consider the structures $\mathcal{M} = (\{a, b, c\}, R^{\mathcal{M}})$ and $\mathcal{M}' = (\{a', b'_1, b'_2, c'_1, c'_2\}, R^{\mathcal{M}'})$ depicted below:

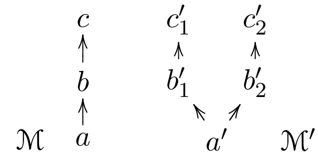

We show that the structures $(\mathcal{M}, a)$ and $(\mathcal{M}', a')$ are LGF-equivalent by showing that for every positive integer $k \geq 2$, there exists a loosely guarded k-bisimulation between $(\mathcal{M}, a)$ and $(\mathcal{M}', a')$. Now, for all $k \geq 2$, there are 6 loosely k-guarded subsets of M, namely $\varnothing$, $\{a\}$, $\{b\}$, $\{c\}$, $\{a, b\}$ and $\{b, c\}$; and there are 10 loosely k-guarded subsets of M', that is, $\varnothing$, $\{a'\}$, $\{b'_1\}$, $\{b'_2\}$, $\{c'_1\}$, $\{c'_2\}$, $\{a', b'_1\}$, $\{a', b'_2\}$, $\{b'_1, c'_1\}$ and $\{b'_2, c'_2\}$. Let $\mathcal{F}$ be the set consisting of the following maps of type $M \to M'$: $\varnothing$, $\{(a, a')\}$, $\{(b, b'_1)\}$, $\{(c, c'_1)\}$, $\{(b, b'_2)\}$, $\{(c, c'_2)\}$, $\{(a, a'), (b, b'_1)\}$, $\{(a, a'), (b, b'_2)\}$, $\{(b, b'_1), (c, c'_1)\}$ and $\{(b, b'_2), (c, c'_2)\}$. These maps are partial isomorpisms. Further, clearly for every $f : X \to Y$ in $\mathcal{F}$ we have both $\mathbf{Zig}(f, X)$ and $\mathbf{Zag}(f, Y)$. For all $k \geq 2$, the set $\mathcal{F}$ is, then, a loosely guarded k-bisimulation between $\mathcal{M}$ and $\mathcal{M}'$ containing the map $\{(a, a')\}$. It follows that the $\{R\}$-structures $(\mathcal{M}, a)$ and $(\mathcal{M}', a')$ satisfy the same $\mathsf{LGF}[\{R\}]$-formulas of exactly one free variable. On the other hand, we note that the MsL_0-formula $\Box(\Diamond/1)\top$ is true in $\mathfrak{M}$ at a but not true in $\mathfrak{M}'$ at a'. Therefore there cannot exist a map T from MsL_0 to the formulas of $\mathsf{LGF}[\{R\}]$ of exactly one free variable such that for all formulas $\phi \in \mathsf{MsL}_0$ and all pointed models $(\mathfrak{N}, w)$, we would have $\mathfrak{N}, w \models \phi$ iff $\mathcal{N}, w \models T(\phi)$. ■

8. Finite modal depth and finite breadth

Restricting attention to a finite set of propositional atoms, the number of pairwise non-equivalent ML formulas of a given modal depth is finite; for a proof, see, e.g., (Blackburn *et al.*, 2002, Prop. 2.29). Not so for the logic MsL_0.

FACT 8.1 (Tulenheimo & Sevenster 2006). *Suppose the set* **prop** *of available propositional atoms is finite. For any $m \geq 2$, the set of pairwise non-truth-equivalent* MsL_0 *formulas of modal depth m is infinite.*

Proof. If $k \geq 0$ and $m \geq 2$, let $\phi^k_m := \Box \ldots \Box \phi_m$, where ϕ_m is defined as in the proof of Theorem 7.4, and the number of boxes preceding ϕ_m

is k. Let $r \geq 2$ be arbitrary, and consider the set of formulas $\{\phi^{r-2}_m : m < \omega\}$. All formulas in this set are of modal depth r. Yet whenever $m_1 < m_2$, the formulas $\phi^{r-2}_{m_1}$ and $\phi^{r-2}_{m_2}$ fail to be truth-equivalent. This follows from Theorem 7.4: the former formula can be translated into FOL^{m_1+r-1} but the latter cannot. ■

There is, however, a way to restore an analogy with basic modal logic. This requires taking into consideration an additional parameter which plays no role in ML but does play a role here: the breadth of an MsL_0 formula (see Definition 4.1).

THEOREM 8.2. *Suppose the set* **prop** *is finite. Let $m, k < \omega$. The set of pairwise non-truth-equivalent* MsL_0 *formulas of modal depth m and breadth k is finite.*

Proof. We prove the claim by double induction on the parameters m and k. The base case of literals (which are both of modal depth and of breadth zero) holds by the assumption that **prop** is finite. We proceed in two steps. (1) First observe that if ϕ is a formula of modal depth $\leq m + 1$ and breadth $\leq k$, it is obtained by some finite number (possibly zero) of applications of $\vee$ and $\wedge$ from formulas of the forms ϕ or $\Box\phi$ or $\Diamond\phi$ or $E_{n,S}(\phi_1, \ldots, \phi_n)$ satisfying the following conditions:

- $md(\phi) \leq m, \quad |S| \leq k, \quad bre(\phi) \leq k$
- $md(\phi_i) \leq m$ if $i \notin S, \quad md(\phi_i) \leq m - 1$ if $i \in S, \quad bre(\phi_i) \leq k.$

Now, let k be fixed and assume that there are (up to truth-equivalence) only finitely many formulas of modal depth m and breadth k. We prove that in that case there are (up to truth-equivalence) likewise only finitely many formulas of modal depth $m+1$ and breadth k. From the assumption it immediately follows that there are only finitely many formulas of each of the forms ϕ or $\Box\phi$ or $\Diamond\phi$ or $E_{n,S}(\phi_1, \ldots, \phi_n)$, where ϕ and the ϕ_i satisfy the relevant syntactic conditions. But the rest of the formulas of modal depth $\leq m + 1$ and breadth $\leq k$ are truth-functions of these formulas. And out of finitely many formulas only finitely many truth-functions can be formed. The claim follows.

(2) Next, note that if ϕ is a formula of modal depth $\leq m$ and breadth $\leq k+1$, it is obtained by some finite number (possibly zero) of applications of $\vee$ and $\wedge$ from formulas of the forms ϕ or $E_{n,S}(\phi_1, \ldots, \phi_n)$ with $|S| \leq k + 1$, where

- $md(\phi) \leq m, \quad md(\phi_i) \leq m - 1$ if $i \notin S, \quad md(\phi_i) \leq m - 2$ if $i \in S$
- $bre(\phi), \ bre(\phi_i) \leq k.$

Let m be fixed. Assume there are (up to truth-equivalence) only finitely many formulas of modal depth m and breadth k. We show that in that case there are (up to truth-equivalence) only finitely many formulas of modal depth m and breadth $k+1$. Directly by assumption there are only finitely many formulas of the forms ϕ or $E_{n,S}(\phi_1,\ldots,\phi_n)$, where ϕ and the ϕ_i satisfy the appropriate syntactic conditions. Since the other formulas of modal depth $\leq m$ and breadth $\leq k+1$ are truth-functions of these formulas, the claim follows. ■

9. Inexpressibility of contradictory negation

In Corollary 3.5 it was seen that MsL is not closed under contradictory negation; the proof turned on the fact that some MsL formulas cannot be translated into FOL. Since the fragment MsL_0 *is* so translatable, we do not know on the basis of the mentioned corollary whether also MsL_0 fails to be closed under contradictory negation. We proceed to prove that as a matter of fact this is the case. Let us first introduce a generalization of the notion of ML bisimulation.

DEFINITION 9.1 (Asymmetric MsL_0 bisimulation). *Let* $(\mathfrak{M},w)$ *and* $(\mathfrak{M}',w')$ *be pointed models, with* $\mathfrak{M}=(M,R,V)$ *and* $\mathfrak{M}'=(M',R',V')$. *A relation* $Z\subseteq M\times M'$ *is an* asymmetric MsL_0 bisimulation *from the pointed model* $(\mathfrak{M},w)$ *to the pointed model* $(\mathfrak{M}',w')$, *symbolically* $Z:(\mathfrak{M},w)\looparrowright(\mathfrak{M}',w')$, *if* Z *satisfies the conditions* (0) *to* (4) *of an* ML *bisimulation and the following additional condition holds:*

(5) *Suppose* $k<\omega$ *and* vZv' *and* $(v,t_i)\in(R\circ R)$ *for all* $1\leq i\leq k$. *Then there are* $t'_1,\ldots,t'_k$ *such that (a) and (b) both hold:*

   (a) $t_iZt'_i$ *and* $(v',t'_i)\in(R'\circ R')$;

   (b) If I *is a subset of* $\{1,\ldots,k\}$ *and* $(v',u')\in R'$ *and* $(u',t'_i)\notin R'$ *for all* $i\in I$, *there is* u *with* uZu' *such that* $(v,u)\in R$ *and* $(u,t_i)\notin R$ *for all* $i\in I$.

We write $(\mathfrak{M},w)\looparrowright(\mathfrak{M}',w')$ *to indicate that there is a relation* Z *with* $Z:(\mathfrak{M},w)\looparrowright(\mathfrak{M}',w')$.

Some comments are in order. Subclause (5b) allows us to select any index set I and to consider the corresponding states t_i with $i\in I$ from the fixed tuple $(t_1,\ldots,t_k)$, and it asserts that if there is in $\mathfrak{M}'$ a successor of v' which is not among the predecessors of any of the states t'_i with $i\in I$, then there is likewise a successor of v in $\mathfrak{M}$ which is not among the predecessors of any of the states t_i with $i \in I$. Positively this condition means that if all successors of v in $\mathfrak{M}$ have at least one of the t_i among their successors, then likewise all successors of v' in $\mathfrak{M}'$ have at least one of the t'_i among their successors. Clause (5) is manifestly asymmetric. Subclause (5a) asserts the existence of a certain tuple of states in $\mathfrak{M}'$ provided that a certain tuple of states exist in $\mathfrak{M}$; there is no clause expressing the analogous condition in the direction from $\mathfrak{M}'$ to $\mathfrak{M}$. And subclause (5b) asserts the existence of a certain state in $\mathfrak{M}$ given that a certain state exists in $\mathfrak{M}'$; no clause expresses the analogous condition in the direction from $\mathfrak{M}$ to $\mathfrak{M}'$.

EXAMPLE 9.2. *Let us consider the models $\mathfrak{M}$, $\mathfrak{M}'$, $\mathfrak{N}$ and $\mathfrak{N}'$:*

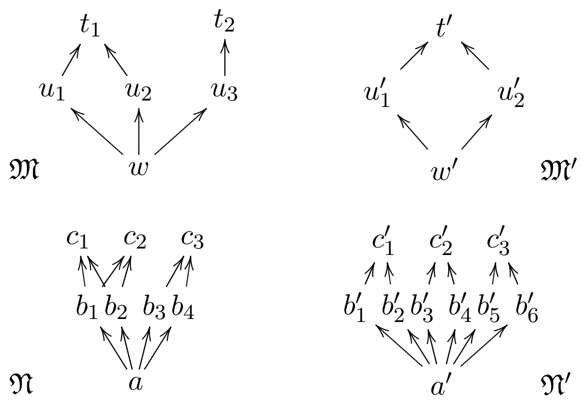

(a) *Let Z be the following relation: $\{(w, w'), (u_1, u'_1), (u_2, u'_1), (u_3, u'_2), (t_1, t'), (t_2, t')\}$. Evidently Z is a bisimulation. Does Z satisfy clause (5)? Given that wZw' and $(w, t_1), (w, t_2) \in (R \circ R)$, we may take as an input of rule (5) one of these three tuples of states: $(t_1), (t_2), (t_1, t_2)$. In each case subclause (a) is satisfied: all states of the relevant set are related via Z to t'. Subclause (b) is trivially met as there is in $\mathfrak{M}'$ no successor of w which would not have t' as its successor. We conclude that $Z : (\mathfrak{M}, w) \looparrowright (\mathfrak{M}', w')$.*

(b) *Conversely, we do* not *have $(\mathfrak{M}', w') \looparrowright (\mathfrak{M}, w)$. For, taking (t') as an input of the rule (5), the state t' would have to be correlated either with t_1 or with t_2. In neither case would subclause (b) be satisfied. Note that the MsL_0 formula $\Box(\Diamond/1)\top$ is true in $\mathfrak{M}'$ at w' but not in $\mathfrak{M}$ at w.*

(c) *We also do* not *have $(\mathfrak{N}, a) \looparrowright (\mathfrak{N}', a')$. To see this, observe first that for the input (c_1, c_3) of rule (5), we must choose a pair (x_1, x_2) of states from the set $\{c'_1, c'_2, c'_3\}$. At least one of the states c'_1, c'_2, c'_3 does not appear in this pair. Hence there is a state b'_j in $\mathfrak{N}'$ having no successor among x_1 and x_2 (actually there are exactly two* *such states). But this violates subclause (5b), since all states among* b_1, b_2, b_3, b_4 *have one of the states* c_1 *and* c_3 *as their successor in* $\mathfrak{N}$*. We note that the formula* $\Box((\Diamond/1)\top \vee (\Diamond/1)\top)$ *is true in* $\mathfrak{N}$ *at* a *but not in* $\mathfrak{N}'$ *at* a'*.*

(d) *We have* $(\mathfrak{N}', a') \looparrowright (\mathfrak{N}, a)$*. Let* Z *be the following relation:* $\{(a, a'), (b_1, b_1'), (b_1, b_3'), (b_2, b_2'), (b_2, b_4'), (b_3, b_5'), (b_4, b_6'), (c_1, c_1'), (c_1, c_2'), (c_2, c_1'), (c_3, c_3')\}$*. Now,* Z *is a bisimulation. Further, for any tuple* $\vec{x} \in \cup_{1 \leq m \leq 3}\{c_1', c_2', c_3'\}^m$*, the corresponding tuple* $\vec{y} \in \cup_{1 \leq m \leq 3}\{c_1, c_2, c_3\}^m$ *(uniquely) determined by the relation* Z *clearly satisfies clause (5).*

Using the notion of asymmetric bisimulation just defined, we obtain a criterion for the preservation of truth of an MsL_0 formula when moving from a given pointed model to another. The criterion precisely does *not* guarantee the preservation of truth also in the direction from the latter to the former pointed model. This asymmetry is crucial when proving that MsL_0 is not closed under contradictory negation. For proofs of this kind, cf., e.g., (Immerman, 1999, *Ch.* 8).

We will use the following terminology. If R is a binary relation, $(U_1, \ldots, U_m)$ is a partition of the set $\{u : R(v, u)\}$ and there are pairwise distinct states $s_1, \ldots, s_n$ with $n \in \{m, m-1\}$ such that s_i is a common R-successor to all states in U_i, we say that the partition is *indexed* by the tuple $(s_1, \ldots, s_n)$. If $n = m-1$, the partition is *with remainder* and the cell U_m is a *remainder cell.* If $n = m$, it is *without remainder.* If $(t_1, \ldots, t_k)$ is any tuple of states (not necessarily pairwise distinct), let $(s_1, \ldots, s_n)$ be a tuple of states such that $\{s_1, \ldots, s_n\} = \{t_1, \ldots, t_k\}$, the states s_i are pairwise distinct (hence $n \leq k$), and s_i precedes s_j iff there are i', j' with $i' < j'$ such that $s_i = t_{i'}$ and $s_j = t_{j'}$. We say derivatively that a partition $(U_1, \ldots, U_m)$ is indexed by the tuple $(t_1, \ldots, t_k)$ if it is indexed by the tuple $(s_1, \ldots, s_n)$. Note that there may in general be many ways of partitioning the set of successors of a state v in such a way as to be indexed by a tuple: this is the case when u_1, u_2, u_3 are successors of v, t is a successor of u_1, u_2 and t' is a successor of u_2, u_3.

THEOREM 9.3. *If* $(\mathfrak{M}, w) \looparrowright (\mathfrak{M}', w')$*, then for all* MsL_0 *formulas* ϕ*, we have: if* $\mathfrak{M}, w \models \phi$*, then* $\mathfrak{M}', w' \models \phi$*.*

Proof. Suppose $Z : (\mathfrak{M}, w) \looparrowright (\mathfrak{M}', w')$. We show that for all states $v \in M$ and $v' \in M'$ with vZv' the following holds: for all $\phi \in \mathsf{MsL}_0$, if $\mathfrak{M}, v \models \phi$, then $\mathfrak{M}', v' \models \phi$. We may consider MsL_0 formulas as generated by the grammar C. It is evident that the claim holds for atoms, and that it can be inductively proven to hold for formulas of the forms

$\vee, \wedge, \Diamond, \Box$; for the cases of $\Diamond$ and $\Box$ we will be able to use the inductive hypothesis thanks to the fact that Z has the properties of usual bisimulation between $(\mathfrak{M}, w)$ to $(\mathfrak{M}', w')$. What remains to be checked are the formulas of the form $\Box D_{n,S}(\xi_1, \ldots, \xi_{n+1})$. Without loss of generality we may assume that there is $1 \leq k \leq n+1$ such that $S = \{1, \ldots, k\}$. Let T_{x_0} be the translation of MsL_0 into FOL discussed in the proof of Theorem 7.2. Now, instead of $\phi := \Box D_{n,S}(\xi_1, \ldots, \xi_{n+1})$, we may just as well consider its FOL translation $\theta := \exists x_1 \ldots \exists x_k \forall y(Rx_0y \to D_{n,S}(\xi'_1, \ldots, \xi'_{n+1}))$, where ξ'_i is the FOL formula $(Ryx_i \wedge T_{x_0/x_i}[\phi_i])$ if $i \in S$ and $\xi_i = (\Diamond/1)\phi_i$, while ξ'_i is the FOL formula $T_{x_0/y}[\xi_i]$ if $i \notin S$. We may note that among the formulas $\xi'_1, \ldots, \xi'_{n+1}$, there is for every variable x_i with $1 \leq i \leq k$ exactly one formula ξ'_j in which this variable occurs free. Now, there are positive integers m and $r_1, \ldots, r_m$ such that θ is logically equivalent to an FOL formula χ of the form

$$\exists x_1 \ldots \exists x_k \bigwedge_{1 \leq i \leq m} \forall y(Rx_0y \to \bigvee_{1 \leq j \leq r_i} \zeta_{ij}),$$

where each ζ_{ij} is one of the formulas $\xi'_1, \ldots, \xi'_{n+1}$. Here one and the same variable x_i with $1 \leq i \leq k$ may well occur free in several conjuncts.

Assume inductively that for any states r and r', if rZr' and $\mathcal{M}, r \models \zeta_{ij}$, then $\mathcal{M}', r' \models \zeta_{ij}$. Now, suppose that vZv' and that for fixed (but not necessarily pairwise distinct) states $t_1, \ldots, t_k$ and for all $1 \leq i \leq m$, we have:

$$\mathcal{M}, v, t_1, \ldots, t_k \models \forall y(Rx_0y \to \bigvee_{1 \leq j \leq r_i} \zeta_{ij}).$$

By the fact that $Z : (\mathfrak{M}, w) \leftrightarrow (\mathfrak{M}', w')$, if we take the tuple $(t_1, \ldots, t_k)$ as the input of rule (5), a certain tuple $(t'_1, \ldots, t'_k)$ is obtained with $t_l Z t'_l$ for all $1 \leq l \leq k$. We proceed to prove that for every $1 \leq i \leq m$,

$$\mathcal{M}', v', t'_1, \ldots, t'_k \models \forall y(Rx_0y \to \bigvee_{1 \leq j \leq r_i} \zeta_{ij}).$$

Let i be fixed. Since $\forall y(Rx_0y \to \bigvee_{1 \leq j \leq r_i} \zeta_{ij})$ is satisfied in $(\mathcal{M}, v, t_1, \ldots, t_k)$, there is $I \subseteq \{1, \ldots, k\}$ such that the set $\{u : R(v, u)\}$ has a partition $\mathcal{P}$ indexed by the tuple $\langle t_i : i \in I \rangle$. The partition $\mathcal{P}$ has at most $r_i \leq k+1$ cells. By subclause (5a) in the definition of asymmetric MsL_0 bisimulation, there is, then, a partition $\mathcal{P}'$ of the set $\{u' : R'(v', u')\}$ indexed by the tuple $\langle t'_i : i \in I \rangle$. Let $l \in I$ be arbitrary. Then there is j such that the state t_l satisfies the formula $T_{x_0/x_l}[\phi_j]$ in $\mathcal{M}$. On the other hand, $t_l Z t'_l$. Thus, by the inductive hypothesis t'_l satisfies $T_{x_0/x_l}[\phi_j]$ in $\mathcal{M}'$. It follows that $\forall y(Rx_0y \to \bigvee_{1 \leq j \leq r_i} \zeta_{ij})$ is satisfied in $(\mathcal{M}', v', t'_1, \ldots, t'_k)$. If, again, the partition $\mathcal{P}$ is with remainder, $\mathcal{P}'$ may or may not be itself with remainder. If it is without remainder, the claim follows immediately by the inductive hypothesis. If it is with remainder, we need to check that every state in the remainder cell of $\mathcal{P}'$ satisfies in $\mathcal{M}'$ some formula ξ'_l with $l \notin S$. Let u' be any state with $R'(v', u')$ in the remainder cell of $\mathcal{P}'$. By clause (5a) we find, then, a state u with $R(v, u)$ in the remainder cell of $\mathcal{P}$ such that uZu'. Because $\mathcal{M}, v, t_1, \ldots, t_k \models \forall y(Rx_0y \rightarrow \bigvee_{1 \leq j \leq r_i} \zeta_{ij})$, there is $l \notin S$ such that $\mathcal{M}, u \models \xi'_l$ (recall that every ζ_{ij} equals ξ'_l for some $1 \leq l \leq n+1$). By the inductive hypothesis we may conclude that $\mathcal{M}', u' \models \xi'_l$, and again the claim follows. ■

In the above proof it was crucial that the states $t'_1, \ldots, t'_k$ provided by the asymmetric MsL_0 bisimulation Z were obtained once and for all as witnesses for certain slashed diamonds, and that it was possible to repeatedly choose a selection I of those states when considering the relevant different conjuncts. Indeed it would have led to too weak a condition if subclause (5b) had been imposed directly on the states $t'_1, \ldots, t'_k$ and not to a separately chosen selection thereof; that would have corresponded to choosing separate witnesses for each conjunct.

We use the following lemma in proving the main result of the present section, i.e., that MsL_0 is not closed under complementation.

LEMMA 9.4. *Let $h, m \geq 0$. If $(\mathfrak{N}, w)$ is the h-fold duplication of the m-unraveling of $(\mathfrak{M}, w)$, then $(\mathfrak{N}, w) \looparrowright (\mathfrak{M}, w)$.*

Proof. First note that the following clearly holds. If $Z_1 : (\mathfrak{M}_1, w_1) \looparrowright (\mathfrak{M}_2, w_2)$ and $Z_2 : (\mathfrak{M}_2, w_2) \looparrowright (\mathfrak{M}_3, w_3)$, then the composite relation $Z_1 \circ Z_2$ satisfies: $Z_1 \circ Z_2 : (\mathfrak{M}_1, w_1) \looparrowright (\mathfrak{M}_3, w_3)$. Let $(\mathfrak{M}, w)$ be a pointed model, $(\mathfrak{M}', w)$ its m-unraveling and $(\mathfrak{N}, w)$ the h-fold duplication of $(\mathfrak{M}', w)$. By what just noted, in order to prove the lemma, it suffices to show that $(\mathfrak{N}, w) \looparrowright (\mathfrak{M}', w)$ and $(\mathfrak{M}', w) \looparrowright (\mathfrak{M}, w)$. Write $\mathfrak{M} = (M, R, V)$, $\mathfrak{M}' = (M', R', V')$ and $\mathfrak{N} = (N, S, U)$.

We define a relation $Z_1 \subseteq M' \times M$ by using the fact that there is a natural correlation of the states of the models $\mathfrak{M}'$ and $\mathfrak{M}$, given the way in which $\mathfrak{M}'$ is produced from $\mathfrak{M}$. For every state v' in $\mathfrak{M}'$, there is a path $(w, v_1, \ldots, v_n)$ along R such that $v' = (w, v'_1, \ldots, v'_n)$. Put $v' Z_1 v_n$. Plainly Z_1 is a bisimulation between $(\mathfrak{M}', w)$ and $(\mathfrak{M}, w)$. We claim that Z_1 also satisfies clause (5) of an asymmetric bisimulation. So suppose $v' Z_1 v$ and let the tuple $(t'_1, \ldots, t'_k)$ of elements of $\mathfrak{M}'$ be the input of clause (5). For every t'_i there is a uniquely determined t_i such that $t'_i Z_1 t_i$. Clearly $(R \circ R)(v, t_i)$. Suppose, then, that there is $I \subseteq \{1, \ldots, k\}$ and u with $R(v, u)$ such that no t_i with $i \in I$ is R-accessible from u. We must show that there is an R'-successor u' of v' from which no state t'_i with $i \in I$ is R'-accessible and which satisfies $u' Z_1 u$. Now, if there was no such state, there would be u' with $u' Z_1 u$

such that $R'(v', u')$ and $R'(u', t'_j)$ for some $j \in I$. Thus, there would be x such that $t'_j Z_1 x$ and $R(u, x)$; indeed t'_j is a certain finite path along R and x is its last member. Since there is only one state from $\mathfrak{M}$ correlated with t'_j via Z_1, it follows that we would have $x = t_j$. Yet x is R-accessible from u, but t_j is not. This is a contradiction. We conclude that $Z_1 : (\mathfrak{M}', w) \looparrowright (\mathfrak{M}, w)$.

To complete the proof, we define a relation $Z_2 \subseteq N \times M'$. By the construction of $(\mathfrak{N}, w)$, there are models $\mathfrak{T}_0, \ldots, \mathfrak{T}_{h \cdot m}$ with the following properties: $\mathfrak{T}_0 = \mathfrak{M}'$; $\mathfrak{T}_{h \cdot m} = \mathfrak{N}$; each $\mathfrak{T}_i$ is a tree of height m; and there is a map $f_i : T_i \to Pow(T_{i+1})$ mapping distinct states to disjoint sets of states such that for every state $x \in T_i$ and every state $y \in f_i(x)$, the submodel of $\mathfrak{T}_i$ generated by x is isomorphic to the submodel of $\mathfrak{T}_{i+1}$ generated by y, and for every state $z \in T_{i+1}$ there is $x \in T_i$ such that $z \in f_i(x)$. We set $v Z_2 v'$ iff there are states $x_0, \ldots, x_{h \cdot m}$ with $x_0 = v'$ and $x_{h \cdot m} = v$ such that for all $0 \leq i < h \cdot m$, we have: $x_{i+1} \in f_i(x_i)$. That is, we have $v Z_2 v'$ iff v is one of the 'copies' of v' created by the process of h-fold duplication of $(\mathfrak{M}', w)$. Trivially Z_2 is a bisimulation between $(\mathfrak{N}, w)$ and $(\mathfrak{M}', w)$. Suppose $v Z_2 v'$ and let the tuple $(s_1, \ldots, s_k)$ of elements of $\mathfrak{N}$ be the input of clause (5). For every s_i there is a uniquely determined s'_i such that $s_i Z_2 s'_i$ and indeed $(R' \circ R')(v', s'_i)$. Now, suppose there is $I \subseteq \{1, \ldots, k\}$ and u' with $R'(v', u')$ such that no s'_i with $i \in I$ is R'-accessible from u'. Then there is u with $S(v, u)$ such that $u Z_2 u'$ and no t_i with $i \in I$ is S-accessible from u. Namely, for contradiction suppose there is u with $u Z_2 u'$ and $S(u, s_j)$ for some $j \in I$. Then there is x with $s_j Z_2 x$ and $R'(u', x)$. But there is exactly one state in $\mathfrak{N}$ to which s_j is related via Z_2, so $x = s'_j$. This is a contradiction and we may conclude that $Z_2 : (\mathfrak{N}, w) \looparrowright (\mathfrak{M}', w)$. ■

Inspecting the proof of Lemma 9.4, we obtain the following:[5]

COROLLARY 9.5. *Suppose Z is a bisimulation from $(\mathfrak{N}, v)$ to $(\mathfrak{M}, w)$. If Z is a surjective function, then in fact Z is an asymmetric* MsL_0 *bisimulation from $(\mathfrak{N}, v)$ to $(\mathfrak{M}, w)$.*

We proceed to prove that there is no proper MsL_0 formula whose contradictory negation is expressible in MsL_0.

THEOREM 9.6 (Strong inexpressibility of contradictory negation). *If ϕ and ψ are* MsL_0 *formulas such that for all pointed models $(\mathfrak{M}, w)$ we have $\mathfrak{M}, w \models \psi$ iff $\mathfrak{M}, w \not\models \phi$, then ϕ is not proper.*

Proof. Assume ϕ is a proper formula of MsL_0. By Lemma 4.5 the class $Mod(\phi)$ is either not closed under $md(\phi)$-unraveling or else it is

[5] The author is grateful to Antti Kuusisto who observed that Lemma 9.4 has this corollary.

not closed under global duplication. That is, there is a pointed model $(\mathfrak{M}, w)$ of height at most $md(\phi)$ such that $\mathfrak{M}, w \models \phi$ while $\mathfrak{M}', w \not\models \phi$, where the pointed model $(\mathfrak{M}', w)$ is either (i) the unraveling of $(\mathfrak{M}, w)$, or else (ii) the global duplication of $(\mathfrak{M}, w)$. Now, suppose for contradiction that there is an MsL_0 formula $neg(\phi)$ which is true in a pointed model iff ϕ is not true therein. We note that $\mathfrak{M}', w \models neg(\phi)$. Let us consider the two cases separately. *Case (i):* In passing note that in this case $(\mathfrak{M}, w)$ cannot be tree-like, for otherwise $(\mathfrak{M}, w)$ would be isomorphic to its own $md(\phi)$-unraveling. Now, by Lemma 9.4 (setting $h := 0$), we have $(\mathfrak{M}', w) \looparrowright (\mathfrak{M}, w)$. Therefore, by Theorem 9.3, we may conclude that $\mathfrak{M}, w \models neg(\phi)$. This is a contradiction. *Case (ii):* First note that $(\mathfrak{M}, w)$ must be tree-like, for otherwise it would be its own global duplication. But if $(\mathfrak{M}, w)$ is tree-like, again Lemma 9.4 applies (this time setting $h := 1$) and we have $(\mathfrak{M}', w) \looparrowright (\mathfrak{M}, w)$. This yields a contradiction, since by Theorem 9.3 we have again $\mathfrak{M}, w \models neg(\phi)$. ■

Theorem 9.6 states a very strong inexpressibility result for the fragment MsL_0 of MsL. Now, by any reasonable criterion $\Box(\Diamond/1)\top$ is the syntactically simplest proper MsL formula. As a consequence of what just proven, we may observe the following fact concerning the larger language MsL, thereby vindicating the claim we made in Section 3 to the effect that Corollary 3.5 (according to which MsL is not closed under contradictory negation) can be improved.

COROLLARY 9.7. *Any fragment of* MsL *which contains the formula* $\Box(\Diamond/1)\top$ *fails to be closed under contradictory negation.*

Proof. Consider the pointed models $(\mathfrak{M}, a)$ and $(\mathfrak{M}', a')$ discussed in the proof of Theorem 7.5. The formula $\Box(\Diamond/1)\top$ is true in the former but not in the latter. On the other hand, these pointed models are evidently bisimilar. Thus, $\Box(\Diamond/1)\top$ is a proper MsL_0 formula. By Theorem 9.6, then, there is no MsL_0 formula that can express the contradictory negation of $\Box(\Diamond/1)\top$. Observe that any MsL formula of modal depth at most 2 is by syntactic criteria a formula of MsL_0. It follows that no MsL formula of modal depth at most 2 can express the contradictory negation of $\Box(\Diamond/1)\top$. Suppose, then, for contradiction that there is $\phi \in \mathsf{MsL}$ with $md(\phi) \geq 3$ such that ϕ is true in a pointed model iff $\Box(\Diamond/1)\top$ is not true therein. Now, clearly any MsL formula of modal depth at most 2 is true at v in $\mathfrak{N}$ iff it is true at v in the restriction of $\mathfrak{N}$ to the height 2. It is also clear that any MsL formula χ of modal depth at least 3 is truth-equivalent over models of height at most 2 to an MsL formula θ_χ of modal depth at most 2. It follows that for all pointed models $(\mathfrak{N}, v)$ we have: $\mathfrak{N}, v \models \phi$ iff $(\mathfrak{N} \restriction 2), v \models \phi$ iff $(\mathfrak{N} \restriction 2), v \models \theta_\phi$ iff $\mathfrak{N}, v \models \theta_\phi$, where $md(\theta_\phi) \leq 2$. But then θ_ϕ is a formula of MsL_0 expressing the contradictory negation of $\Box(\Diamond/1)\top$. This is a contradiction. We conclude that there is no MsL formula which would express the contradictory negation of $\Box(\Diamond/1)\top$. ■

As a by-product of what just established, we obtain a proof of the rather obvious fact that FOL cannot be translated into MsL; in Sections 3 and 7 we promised to exhibit such a proof.

COROLLARY 9.8. $\mathsf{FOL} \not\leq \mathsf{MsL}$.

Proof. The formula $\Box(\Diamond/1)\top$ can be translated into FOL. Therefore also its contradictory negation can. So if we had $\mathsf{FOL} \leq \mathsf{MsL}$, the contradictory negation of $\Box(\Diamond/1)\top$ could be translated into MsL, contrary to Corollary 9.7. ■

10. The robust decidability of modal logic revisited

We proceed to draw some general morals from the properties of MsL_0 that we have detected. Let us first note a couple of relevant facts about this logic.

10.1. Computational properties of MsL_0

Let us say that logic L is *strongly closed under complementation* if there exists a **PTIME**-computable function f such that for every $\phi \in L$ we have that $f(\phi)$ is true in a model iff ϕ is not true therein. If L is strongly closed under complementation, the satisfiability problem of L is decidable iff its validity problem is decidable. Moreover, if one of the problems is **C**-complete for some complexity class **C**, the other is **coC**-complete. If in particular **C** happens to be a deterministic complexity class, both problems are **C**-complete if one of them is: in this case **coC** = **C**. Now, many commonly encountered logics are strongly closed under complementation. Among them are the following: the logic $\mathsf{S5}$ (i.e., ML evaluated over models whose accessibility relation is an equivalence relation), the (poly)modal logics $\mathsf{K4}_n$ (i.e., the extension of ML with n modalities, each associated with a transitive accessibility relation), the hybrid logic $\mathcal{H}(@)$, the logic of until and since (US), the finite-variable fragments FOL^n of first-order logic, the guarded fragment (GF) and the loosely guarded fragment (LGF) of first-order logic, the fixed point extensions $\mu\mathsf{ML}$ (μ-calculus), $\mu\mathsf{GF}$ and $\mu\mathsf{LGF}$ of ML, GF and LGF, respectively. All these logics have available a syntactic sign for the contradictory negation, so in each of these cases there is a *constant time* computable function yielding the contradictory negation of a given input formula. Observe that in logics like MsL which are *not* closed under complementation to begin with (cf. Theorem 9.6), the satisfiability and validity problems need to be discussed separately — or at least there is no obvious way to reduce one to the other.

The following results are immediate consequences of results proven in (Tulenheimo and Sevenster, 2006) about the logic $\mathcal{L}_{\mathbf{SD}}$ (see *ibid.*, Theorems 15, 17 and 19). Logic $\mathcal{L}_{\mathbf{SD}}$ is a notational variant of MsL; more specifically there exist obvious linear-time computable truth-preserving translations from $\mathcal{L}_{\mathbf{SD}}$ to MsL and from MsL to $\mathcal{L}_{\mathbf{SD}}$.

PROPOSITION 10.1. (**a**) *The model-checking problem for* MsL_0 *is* **NP***-complete in combined complexity (i.e., measured in terms of the combined length of both inputs);* (**b**) *the validity problem of* MsL_0 *is* **PSPACE***-complete;* (**c**) MsL_0 *has the strong finite model property (i.e., there is a recursive upper bound to the size of every satisfiable* MsL_0 *formula); and* (**d**) *the satisfiability problem of* MsL_0 *is* **PSPACE***-complete.*

10.2. ML and MsL_0 as fragments of FOL

The fact that basic modal logic has many 'good' properties that first-order logic lacks has generated a remarkable literature. The model-checking problem and the validity and satisfiability problems of ML are decidable. More specifically, the combined complexity of the model-checking problem for ML is **PTIME**-complete (Grädel and Otto, 1999), while its satisfiability problem is **PSPACE**-complete (Ladner, 1977). Now, ML is semantically speaking a fragment of first-order logic. Via its translation, then, ML determines a decidable fragment of FOL. This has been considered interesting because in the relevant fragment no limit is imposed on the allowed quantifier alternations — while historically, known decidable fragments of FOL had been obtained by restricting patterns of allowed quantifier alternations; see, e.g., (Börger *et al.*, 1997). We may note that MsL_0 is not less interesting in this respect.

OBSERVATION 10.2. *Also* MsL_0 *is semantically a decidable fragment of* FOL *with unbounded quantifier alternation (Theorem 7.2). In particular, it is strictly more expressive than* ML *(Fact 3.1).*

Basic modal logic has been hailed not only because it is decidable (and decidable using an algorithm of a relatively reasonable complexity), but because it is 'robustly decidable': it has a host of independently interesting extensions or variants which are likewise (elementarily) decidable. Examples abound. The satisfiability problem of S5 is **NP**-complete (Ladner, 1977), while that of $\mathsf{K4}_n$ is **PSPACE**-complete for all $n \geq 1$; for hardness, see (Ladner, 1977) and for inclusion cf., e.g., (Horrocks *et al.*, 2007, Thm. 21). The satisfiability of the hybrid logic $\mathcal{H}(@)$ is **PSPACE**-complete (Areces *et al.*, 1999). It was noted in (Burgess, 1982) that the decidability of the logic of until and since (US) evaluated over arbitrary linear orders immediately follows from the results of (Rabin, 1969). Actually the satisfiability problem of US over arbitrary linear orders is **PSPACE**-complete: inclusion has been proven in (Reynolds, 2010; Rabinovich, 2010); hardness follows from the main result of (Reynolds, 2003). The satisfiability problem of the μ-calculus is **EXPTIME**-complete (Emerson and Jutla, 1988).

OBSERVATION 10.3. *Modal slash logic provides yet another confirmation of the robust decidability of* ML *— in a strikingly strong way.*

- **(a)** *Extending* ML *by performing slashing in the way permitted by the syntax of* MsL_0 *yields a logic with a decidable satisfiable problem.*[6] *Even better, insofar as the satisfiability problem is concerned, nothing is lost in terms of computational complexity in the transition from* ML *to* MsL_0*: the satisfiability problem of both logics is* **PSPACE***-complete. Unless* **PTIME** = **NP***, the difficulty of solving the model-checking problem is increased, though (cf. Proposition 10.1).*
- **(b)** *Consider the full logic* MsL *— which, as we recall, is not translatable into* FOL*. It was proven in (*Sevenster, 2010*, Thm. 4.9) that the satisfiability problem of a certain polymodal extension of (a notational variant of)* MsL *is decidable, actually in* **2NEXPTIME***. Therefore, in particular,* MsL *singles out a decidable fragment of* ESO *lying beyond* FOL*. And we obtain a host of new decidable fragments of* FOL*:* a fortiori *every first-order translatable fragment of* MsL *has a decidable satisfiability problem.*

10.3. Explaining the robust decidability of ML

It is natural to ask — and has indeed been asked by several people — what it is in ML that is responsible for its good computational behavior, and especially for its robust decidability. Looking at different ways of translating ML into FOL, one may reflect upon syntactic properties of the translations, and endeavor to identify, from the viewpoint of FOL,

[6] To employ an imprecise turn of phrase of a kind common in the literature, ML remains decidable even when it is subjected to some slashing. Of course in reality ML is not the subject of change here — there is no subject of change — rather one logic gives rise to another and a certain property is preserved.

syntactic features that would throw light on the relevant properties of ML. One may expect this to lead to a generalization: a fragment of FOL more expressive than ML, or at least syntactically an extension of the translation of ML, yet exhibiting appropriate 'nice' properties. The authors discussing the issue have phrased it as a quest for an *explanation*; cf. (Vardi, 1998, pp. 8, 22), (Andréka *et al.*, 1998, pp. 243, 261), (Grädel, 2001, pp. 8, 12). If any sort of explanation is sought for, it cannot simply consist of a description applying to ML and to nothing else — any more than the question 'Why does Moon revolve around the Earth?' could be answered by saying 'Because it is Moon.' At the very least, we need to pinpoint a more general phenomenon which subsumes the case of ML.

Concretely, at least three features have been proposed to have an explanatory role: the translatability of ML into the 2-variable fragment of first-order logic (FOL^2), its translatability into the guarded fragment of first-order logic (GF), as well as the tree-model property that ML enjoys. The first two involve syntactic conditions imposed on the relevant fragment of FOL, whereas the third one is a purely semantic condition. The relevance of these proposals comes from the good computational behavior of the logics FOL^2 and GF as well as from the fact that the decidability of logics having the tree model property can be approached using very widely applicable automata-theoretic tools. The combined complexity of the model-checking problem of FOL^2 is **PTIME**-complete (Grädel and Otto, 1999). The logic FOL^2 *without* equality is decidable (Scott, 1962). This is enough to conclude that ML is decidable, since equality is not needed to translate (in polynomial time) ML into FOL^2. The complexity of the satisfiability problem of FOL^2 is **NEXPTIME**-complete; hardness follows from a result of Fürer (1981) while inclusion was proven in (Grädel *et al.*, 1997). What about the logic GF? Like with FOL^2, also the combined complexity of the model-checking problem of GF is **PTIME**-complete (Berwanger and Grädel, 2001). The decidability of GF was proven in (Andréka *et al.*, 1998). The satisfiability problem of GF is **2EXPTIME**-complete (Grädel, 1999a).

Let us restrict attention to syntactically phrased explanations — such as FOL^2 and GF. Even when an automata-theoretic account of the decidability of a modal logic is available, we would still wish to understand of what its decidability consists from the syntactic viewpoint. Now, it makes sense to ask which syntactic features of the translation of ML into FOL are responsible for its decidability — except that of course there is no unique translation, instead there are syntactically innumerable translations which might lead to different assessments of what it is that is responsible for the properties considered. Relative to a given translation, what does it, then, mean to figure out which syntactic features are responsible for the decidability of the fragment. Well, it can only mean detecting irrelevant syntactic features and showing that generalizing with respect to those features still results to a fragment with sufficiently nice features. Of course even relative to a fixed translation there might be several ways of effecting such a generalization. When one poses questions such as 'Why is ML so robustly decidable?' or looks for an explanation of the decidability of ML, what one is doing is at least searching for such generalizations. One's search is typically guided by some preexisting idea as to what modal logic is. Then one aims to single out *modal* fragments of first-order logic and perhaps to understand in first-order terms of what the modal nature of a fragment consists.

How should we evaluate whether a proposed syntax-based explanation of the robust decidability of ML is correct or illuminating or credible? We may distinguish a strong and a weak criterion. In what follows, the options are, for simplicity, expressed in terms of the satisfiability problem, even though we might wish to consider also further computational problems. From now on, when we say without qualification that logic L' is *translatable* into L, we mean that there exists a polynomial time computable truth-preserving translation of L' into L (thereby deviating from the terminology introduced in Subsect. 1.1).

(1) We say that the decidability of an extension or variant of ML is *strongly explained* by the translatability of ML into logic L, if L is decidable (using an algorithm of such-and-such complexity). Logic L *strongly explains the robust decidability* of ML if it strongly explains the decidability of all relevant extensions or variants of ML. For example, one manifestation of the robust decidability of ML is that the satisfiability problem of the modal logic T (or the result of restricting the evaluation of ML to models with a reflexive accessibility relation) is also decidable, in fact **PSPACE**-complete (Ladner, 1977). This fact is strongly explained by the translatability of ML into FOL^2, since T can be translated into FOL^2: to find out whether a formula ϕ of T is satisfiable, apply a decision algorithm of FOL^2 to the formula $(\psi_\phi \wedge \forall x Rxx)$, where ψ_ϕ is a translation of ϕ into FOL^2 using only the variables x and y. Also basic tense logic is decidable and indeed translatable into FOL^2. For explanations in this sense, cf. (Vardi, 1998, p. 11). Note that it might be tempting to liberalize the definition of the notion of strong explanation by saying that it suffices that the satisfiability problem of M can be *reduced* to that of L. Here one should be careful, though. Since for example the satisfiability problem of FOL^2 is **NEXPTIME**-

complete, we could according to such a definition trivially 'explain' the decidability of any extension or variant M of ML whose satisfiability problem happens to lie in **NEXPTIME**. For, for such a logic M, there would trivially exist a **PTIME**-reduction $t : \mathsf{M} \to L$ such that for any formula ϕ of M, we would have that ϕ is satisfiable iff $t(\phi)$ is satisfiable. The more complex the satisfiability of the *explanans*, the more we would purportedly get explanatory power! We are well-advised to stay with reductions phrased in terms of truth-preserving (**PTIME**-computable) translations.

(2) If logic L is decidable (using an algorithm of such-and-such complexity), the translatability of ML into L *weakly explains* the decidability of an *extension* M of ML, if the result L' of extending L *in the same way* as M extends ML has a decidable satisfiability problem. For example, the translatability of ML into GF offers a weak explanation of the fact that the satisfiability problem of μML is decidable; actually this problem is **EXPTIME**-complete (Emerson and Jutla, 1988). Namely, the satisfiability problem of the fixed point extension μGF of GF is decidable, in fact **2EXPTIME**-complete (Grädel and Walukiewicz, 2006). Incidentally, the same fact is *not* weakly explained by the translatability of ML into FOL^2, because the fixed point extension of FOL^2 is highly undecidable (Grädel *et al.*, 1997). If M is a variant of ML obtained by imposing a restriction on its models, the fact of translatability weakly explains the decidability of the *variant* if the result of restricting the models of L in an analogous way is decidable as well. In connection with such model restrictions, several 'analogous' restrictions on the side of L may present themselves. If, say, ML is considered relative to models whose accessibility relation is required to be transitive (i.e., when considering the logic K4), according to the generalization chosen we must see how L behaves over structures which interpret exactly one/at least one/every binary predicate by a transitive relation; even further options are conceivable as will be noted below, cf. (Szwast and Tendera, 2004). Logic L *weakly explains the robust decidability* of ML, then, if it weakly explains the decidability of all relevant extensions and variants of ML.

The above terminology proves useful for systematic purposes, but one should not be misled by it. The existence of a strong explanation does *not* guarantee the existence of a weak explanation. To see this, write X for the result of translating μML into monadic second-order logic. Let $Y := \mathsf{FOL}^2 \cup X$. We assume for simplicity that FOL^2 utilizes the variables x, y but no formula of X uses these individual variables. Then we can check in **PTIME** whether a formula of Y is from X or not, and Y is decidable because FOL^2 and X are. Now, by the above criterion the translatability of ML into Y provides a strong explanation for the decidability of $\mu\mathsf{ML}$, since (trivially) $\mu\mathsf{ML} \leq Y$ and Y is decidable. Yet this fact does not provide a weak explanation for the decidability of $\mu\mathsf{ML}$: the fixed point extension of Y contains the fixed point extension of FOL^2 which is highly undecidable. (It does not matter in which particular way we define the fixed point extension of X.) Conversely, the existence of a weak explanation does *not* guarantee the existence of a strong one. For example the translatability of ML into GF weakly explains the decidability of $\mu\mathsf{ML}$ because $\mu\mathsf{GF}$ is decidable, but it does not strongly explain the decidability of $\mu\mathsf{ML}$, because $\mu\mathsf{ML}$ cannot be translated into FOL to begin with and so *a fortiori* not into GF.

10.4. How do the proposed explanations fare?

Let us look at the proposed explanations of the robust decidability of ML in some more detail. In the literature, the observation that ML is translatable into FOL^2 is often attributed to Gabbay (1981). So is the general proposal that there is an intimate relation between the finite-variable fragments of FOL and modal logics; cf., e.g., (Vardi, 1998; van Benthem, 1995; Andréka *et al.*, 1998). In reality Gabbay did not even mention the fact that basic modal or basic tense logic can be translated into FOL^2; neither did he 'identify' modal fragments of FOL with its finite-variable fragments. What he did, though, is to establish a lemma, namely Lemma 57 in (Gabbay, 1981), stating that a k-dimensional modal/temporal logic with a finite number of connectives $\sharp_i$ — the semantics of each connective being defined by some first-order formula of k free variables, called the *table* of the connective — can be translated into $\mathsf{FOL}^{\max_i n(i)+k}$, using at most $\max_i n(i)$ bound and k free variables, where for every relevant i, the number of bound variables used in the table of the connective $\sharp_i$ equals $n(i)$. The proof does not make use of properties of any specific class of models, so the result applies generally. Now, if one takes k-dimensional modal logics in the sense of Gabbay as exhausting the possible modal fragments of FOL *and* if one decides to restrict attention modal logics with only finitely many connectives, then indeed Gabbay's lemma implies that any such modal fragment can be translated into FOL^n for some $n < \omega$. Pointing to this doubly conditional consequence of the lemma is the best that can be said by way of justifying the attributions mentioned in the beginning of this paragraph. Inspecting, then, the appropriate tables, it is seen, e.g., that ML is translatable into FOL^2 and US into FOL^3. (Both logics are one-dimensional; the tables of the former use 1 and the tables of the latter 2 bound variables.) Returning to FOL^2 and the robust decidability of

ML, the explanatory value of the embeddability of ML into FOL^2 is limited. This fact does not strongly explain the decidability of μML, simply because $\mu\mathsf{ML} \not\leq \mathsf{FOL}^2$. For a further example, the logic US evaluated over arbitrary linear orders is a decidable variant of ML whose decidability is not strongly explained by the translatability of ML into FOL^2. For, 3 variables are actually needed to express the tables of *until* and *since* and therefore US cannot be translated into FOL^2 (relative to the class of all linear orders). In passing, note that in connection with a logic such as US there is no clear meaning to the notion of weak explanation: *how* should we extend FOL^2, say, in order to extend it as US extends ML?

In (Vardi, 1998), the logic S5 is discussed as a counterexample to the explanatory role of FOL^2. It is a fact that in order to express transitivity in FOL, 3 variables are needed. Because any translation of an S5 formula, evaluated over arbitrary models, must express that the interpretation of its binary relation symbol is transitive, the translation cannot be in FOL^2. No strong explanation of the decidability of S5 is offered by the fact that $\mathsf{ML} \leq \mathsf{FOL}^2$. Note that translatability into FOL^3 cannot be used to formulate an alternative strong explanation, since FOL^n with $n \geq 3$ (even without equality) is undecidable; this follows from the simple fact that already FOL^3 harbors the prefix class AEA which is undecidable; see, e.g., (Börger *et al.*, 1997). It could be argued, however, that it is unfair to require that transitivity should be expressible in a logic L in order for L to be able to explain the decidability of S5. After all, S5 itself is nothing but ML evaluated on a specific class of models, so we should rather ask whether $\mathsf{FOL}^2[\{R, P_1, \ldots\}]$ is decidable when evaluated over structures interpreting the binary relation symbol 'R' by an equivalence relation. That is, we should rather ask whether we can have a *weak* explanation here. As a matter of fact, it was proven in (Kieroński and Otto, 2005) that the satisfiability problem of FOL^2 with a single equivalence relation is decidable, in fact **NEXPTIME**-complete. So arguably S5 does not provide a counterexample to the explanatory role of FOL^2 in the weak sense. Yet, other examples suffice for showing that FOL^2 cannot function even as a weak explanation. It was already noted in Subsection 10.3 that the translatability of ML into FOL^2 does not weakly explain the decidability of the μ-calculus: the fixed point extension of FOL^2 is highly undecidable.

At the end of Subsection 10.3 we mentioned that the translatability of ML into GF explains weakly but not strongly the decidability of the μ-calculus. Another case in which strong explanation fails — a case in which, as noted, the question of weak explanation cannot even be formulated in a natural way — is that of the logic US evaluated over the arbitrary linear orders. Take for example the formula $until(p, q)$ which can be translated into FOL by the formula $\exists y(x < y \wedge Py \wedge \forall z([x < z \wedge z < y] \to Qz))$. This formula is not in GF: the quantifier $\forall z$ is not guarded in a suitable way. Actually it can be shown that this formula does not even have a translation into GF. While US over arbitrary linear orders is decidable, this cannot, then, be strongly explained with reference to the translatability of ML into GF. This observation need not yet discourage those who see in something like the guarded fragment the key to the explanation of robust modal decidability. For, we may reformulate the strong explanation thesis in terms of LGF instead of GF. Why should this be relevant? As an aside it can be mentioned that LGF over *arbitrary models* is decidable — a proof sketch was presented in (van Benthem, 1997) — in fact its satisfiability problem is **2EXPTIME**-complete (Grädel, 1999a). But this fact is irrelevant here. Write $\tau = \{<, P_1, \ldots, P_n\}$, with $<$ binary and the P_j unary. We can only proceed if the logic $\mathsf{LGF}[\tau]$ is decidable over structures interpreting '$<$' as a linear order. But this is the case: actually the full $\mathsf{FOL}[\tau]$ is decidable over such structures (Ehrenfeucht, 1959). Therefore, indeed we may take LGF as strongly explaining the decidability of US over the class of all linear orders. What about the logics $\mathsf{K4}_n$? Transitivity is not expressible in LGF (Grädel, 1999a), so the translatability of ML into LGF cannot strongly explain the decidability of $\mathsf{K4}_n$ for any $n \geq 1$. Can the decidability of $\mathsf{K4}_n$ be weakly explained with reference to GF or LGF? Let τ_k be a vocabulary containing (among other items) the relation symbols $R_1, \ldots, R_k$ and write $\mathcal{T}_k$ for the class of all τ-structures that interpret the relation symbols $R_1, \ldots, R_k$ by transitive relations. If $S_1, \ldots, S_k$ are binary relations, let $\mathit{Trans}[S_1, \ldots, S_k]$ be the statement (not expressible in LGF) that the interpretations of 'S_1', 'S_2', etc. are transitive. If we take the weak explanation to mean that the satisfiability problem of $\mathsf{GF}[\tau_k]$ (or that of $\mathsf{LGF}[\tau_k]$) is decidable over $\mathcal{T}_k$, the answer is in the negative for all $k \geq 2$. For, a $\mathsf{GF}[\tau_k]$ formula φ is satisfiable over $\mathcal{T}_k$ if and only if the statement $\varphi \wedge \mathit{Trans}[R_1, R_2]$ is satisfiable *simpliciter*. But Grädel has shown that the latter question is undecidable for formulas $\varphi \in \mathsf{GF}[\tau_2]$; see (Grädel, 1999a, proof of Theorem 5.10). In the case of LGF, the answer is known to be negative already for $k = 1$ (even when there is no equality in the language); see (Ganzinger *et al.*, 1999). Let us finally consider giving the following meaning to the notion of weak explanation. Define the *guarded fragment with transitive guards*, denoted $[\mathsf{GF} + \mathsf{TG}]$, to be the extension of GF (of a vocabulary without constant symbols) which allows even formulas of the form $\varphi \wedge \mathit{Trans}[R_1, \ldots, R_k]$, provided that in φ the binary relation symbols $R_1, \ldots, R_k$ occur in guards. Clearly there is a truth-preserving translation of type $\mathsf{K4} \to [\mathsf{GF} + \mathsf{TG}]$. So it might be suggested that if the satisfiability problem of $[\mathsf{GF} + \mathsf{TG}]$ is decidable, this would provide a weak explanation of the decidability of $\mathsf{K4}$. And actually it was proven in (Szwast and Tendera, 2004) that the former problem is indeed decidable, in fact **2EXPTIME**-complete. The discussion above shows that it is not obvious how to choose the specific first-order decision problem that would function as a weak explanation of a given modal decision problem. Though at first sight $[\mathsf{GF} + \mathsf{TG}]$ looks promising, it would be more in line with the modal case to find a direct model restriction to be imposed on models of GF than forcing such a restriction with expressions like $\mathit{Trans}[R_1, \ldots, R_k]$ not inherently belonging to our language. Unfortunately for those who would like to see GF as a universal explanation in modal decidability matters, no such model restriction is forthcoming in the case of $[\mathsf{GF} + \mathsf{TG}]$. The best we could do is the following 'hybrid' formulation. Every formula φ of $\mathsf{GF}[\tau]$ has some finite number $n(\varphi)$ of binary guards, so from φ we can select a set of binary guards in $2^{n(\varphi)}$ ways. For every $0 \leq i \leq 2^{n(\varphi)}$, define $\mathcal{C}^i_\varphi := \{(\varphi, \mathcal{M}) : \mathcal{M}$ is a τ-structure interpreting by a transitive relation every each binary guard belonging to the selection $i\}$. We stipulate that the case $i := 0$ corresponds to the selection of the empty set of guards. Note that $\mathcal{C}^0_\varphi = \{(\varphi, \mathcal{M}) : \mathcal{M}$ is a τ-structure$\}$. Finally, put $\mathfrak{C} := \{\mathcal{C}^i_\varphi : \varphi \in \mathsf{GF}[\tau]$ and $0 \leq i \leq m(\varphi)\}$. Now, the satisfiability problem of $[\mathsf{GF} + \mathsf{TG}]$ can be reduced to the following problem: given a class $\mathcal{C}^i_\varphi \in \mathfrak{C}$, find out whether there is a model $\mathcal{M}$ with $(\varphi, \mathcal{M}) \in \mathcal{C}^i_\varphi$ such that $\mathcal{M} \models \varphi$. This formulation makes use of model restrictions, but they are relative to the syntax of a given formula. That is, if we say that a weak explanation of the decidability of $\mathsf{K4}$ consists of the decidability of the latter decision problem, also this suggestion appears rather artificial.

10.5. Lessons from MsL_0

By the results discussed earlier in this paper, the logic MsL_0 is an extension of ML that appears to provide a rather strong counterexample to *each* of the explanations offered in the literature as being relevant to the robust decidability of ML. First, as noted in Observation 10.3, on the one hand MsL_0, and even the full MsL, are decidable extensions of ML, but on the other hand, insofar as the satisfiability and validity problems are concerned, the logic MsL_0 behaves algorithmically exactly as well as ML itself. Second, however, we have seen that the following negative facts hold:

1. The logic MsL_0 is able to state that certain paths of the same length in a pointed model converge (have a common successor), so this logic lacks the tree-model property (Fact 3.2).

2. Not only is MsL_0 not translatable into FOL^2, but it cannot be translated into any finite-variable fragment of FOL (Theorem 7.4).

3. There is no truth-preserving translation of MsL_0 into LGF (Theorem 7.5). That is, the translatability of ML into LGF does not strongly explain the decidability of MsL_0.

On the basis of these three negative properties, it can be said that MsL_0 rather convincingly resists features commonly associated with decidable extensions of ML. This is worth noting, even if it remains quite possible that the translatability of ML into LGF *weakly* explains the decidability of MsL_0: the relevance of LGF for the decidability of MsL_0 could still be saved by showing the decidability of the result of extending LGF in the same way as MsL_0 extends ML; cf. the discussion in (Tulenheimo and Sevenster, 2006). Different formulations of such a '*slashed* LGF' (or sLGF) could be proposed, but the most immediate formulation would be as follows. All LGF formulas are formulas of sLGF. Further, suppose a formula ϕ of sLGF has a subformula ϕ' of the form $\forall\vec{x}(G \to \psi)$. If here ψ contains a subformula $\psi' := \exists\vec{y}(G' \wedge \chi)$ such that in ψ no quantifier syntactically precedes $\exists\vec{y}$, the result of replacing in ϕ the subformula ψ' by the string $(\exists\vec{y}/\vec{x})(G' \wedge \chi)$ is likewise a formula of sLGF. The semantics of sLGF is obtained from that of FOsL by taking $Qx_1 \ldots x_n$ to be a shorthand notation for $Qx_1 \ldots Qx_n$ when $Q \in \{\forall, \exists\}$, and $(\exists y_1 \ldots y_m/x_1 \ldots x_n)$ to be an abbreviation of $(\exists y_1/x_1, \ldots, x_n) \ldots (\exists y_m/x_1, \ldots, x_n)$. We leave it as an open question whether the satisfiability problem of sLGF is decidable.

Given that semantically MsL_0 liberalizes certain restrictions inherent in ML, and that it is even most conveniently defined precisely as a result of such liberalization, it does not appear far-fetched to view it as a modal logic, or if more general term is needed, as a modal-like logic. In any case it is obviously an extension of ML. Since it furthermore behaves computationally very well indeed, any proposal purporting to clarify the computational properties of ML and its extensions should also prove to have explanatory power here. However, as just reiterated, none of the suggestions routinely resorted to in the literature in connection with modal logics is applicable (with the possible exception of LGF in the sense of weak explanation).

We conclude by some comments about the negative results (2) and (3) listed above. As to finite-variable fragments, we have mentioned Gabbay's lemma to the effect that any finite-dimensional modal logic with finitely many connectives, each with a first-order table, can be translated into FOL^n for some $n < \omega$. How does MsL_0 look like from this perspective? Does it have finitely many connectives? Which *are* its connectives? From the viewpoint of slash logic, the answer to the last question is clear: its connectives are the two modal operators $\Box$, $\Diamond$ and the propositional connectives. The slash symbol '/' is not an operator — but, as (Hintikka, 1997, p. 523) puts it, a punctuation device, which helps to complete the job of parentheses as scope-indicators. However, it is plain that the tables of these connectives only yield the semantics of ML, not that of MsL_0. If we wish to attempt forcing MsL_0 into the format in which the semantics of a modal language is defined via tables associated with its connectives, it appears advisable to treat the expression $(\Diamond/1)$ as if it was an operator of its own. Then we can assign the following table to this pretended operator: $(\exists y/x)(Rxy \wedge Py)$. Supposing the remaining connectives $\Box, \Diamond, \wedge, \vee, \neg$ have their usual tables, we end up having a way to assign to every MsL_0 formula ϕ a formula $T[\psi]$ of FOsL expressing the truth condition of ϕ. Such a formula $T[\psi]$ is simply computed from atomic subformulas of ϕ by making use of the tables; incidentally, we may arrange that the formulas $T[\psi]$ only use the variables x and y. Construed in this way, MsL_0 is a one-dimensional modal logic with finitely many connectives. The tables of these connectives are strings of FOsL; the truth conditions of MsL_0 formulas are defined by formulas of FOsL obtained via the tables. Yet MsL_0 is not translatable into any finite-variable fragment of FOL. That is, Gabbay's lemma cannot be generalized by allowing the tables of connectives to be even very simple FOsL formulas. The interest of this observation is reinforced by the fact that still the truth condition of each single MsL_0 formula *can* be expressed even in FOL. However, if we wanted to use tables written in FOL to express the semantics of MsL_0, we would need to introduce *infinitely* many connectives: essentially the connectives $C_{n,S}$ of grammar C. While the table of $(\Diamond/1)$ can be translated into FOL, there is no n such that every 'substitution instance' of the table can be translated into FOL^n.

Let us then note how MsL as a semantic fragment of FOL behaves from the viewpoint of guards. In the guarded fragment, quantifiers are both (syntactically) relativized and (semantically) guarded: quantifiers only appear relativized, i.e., in contexts of the form $Q\vec{x}(G \star \psi)$ with $(Q, \star) \in \{(\forall, \rightarrow), (\exists, \wedge)\}$, and furthermore the relativizer G is a guard in the sense that it imposes an explicit condition relating the variables in the tuple $\vec{x}$ to the free variables of ψ. We might conceive of separating the two functions of guards. A formula might serve to impose the relevant sort of semantic condition without being syntactically a relativizer of the quantifier. For example in the formula $\exists z \forall y(Rxy \rightarrow (Ryz \wedge Pz))$ the atomic formula Ryz in this sense guards the quantifier $\exists z$ without being, syntactically speaking, its relativizer.

Indeed, in the standard first-order translation of MsL_0 the two functions of guards *are* separated. Let us state in a precise way how MsL_0 in this respect differs from ML on the one hand and GF on the other. Let us say that an atomic formula (or a quantifier) appearing in a formula ψ is *ungoverned* if in ψ this atom (respectively this quantifier) does not itself appear in the scope of any quantifier in ψ. First note that one way to express the *locality* of ML is to say that in its standard first-order translation, (1) any token of a quantifier Qx is associated with a unique formula $(G \star \psi)$ where G is the guard of Qx, (2) the guard is the relativizer of the quantifier, i.e., Qx is actually prefixed to the formula $(G \star \psi)$, and there are strict syntactic restrictions on where the variable x introduced by the quantifier Qx may appear in the formula ψ: (3) x is the only free variable of ψ, and (4) in ψ the variable x may only appear in an ungoverned atomic formula or in the guard of an ungoverned quantifier. Now, if we allow the use of binary syntactic structures in the syntax of first-order logic, then by inspecting the translation explained in the proof of Theorem 7.2, it is readily seen that MsL_0 deviates from ML only in giving up restriction (2) for existential quantifiers. For, in such a translation existential quantifiers may occur in contexts $\exists x_{i_1} \ldots \exists x_{i_{|S|}} \forall x(\mathsf{R}x_0 x \to C_{n,S}(\xi_1, \ldots, \xi_{n+1}))$, where each ξ_{i_j} is of the form $(G_{i_j} \wedge \psi_{i_j})$ and the context $(G_{i_j} \wedge \psi_{i_j})$ satisfies restrictions (1), (3) and (4).

The case of MsL_0 suggests, then, that being a relativizer may not be essential for being a guard, insofar as one wishes to consider being guarded as a feature of quantifiers which would explain decidability. On the other hand, MsL_0 retains also the locality properties (3) and (4) of ML, so it is not clear that its decidability can be predicated on its being guarded even in this weaker sense. We may note that from the *first-order* perspective, formulas of MsL_0, like those of ML, serve to impose relational tests on variables introduced by quantifiers, but unlike in the case of ML, in MsL_0 such a test is carried out so to say with a delay: not immediately when a value corresponding to an existentially quantified variable is introduced, but only later it is checked whether this value stands in a suitable relation to a value introduced by the subsequent evaluation of a universal quantifier. For comparison we note that in GF requirements (1) and (2) are retained but constraints (3) and (4) are given up in favor of allowing the variable x to appear anywhere in the scope of Qx, as long as G serves to impose an explicit relational condition on x and the free variables of ψ; indeed instead of a single variable x a tuple $\vec{x}$ of variables is allowed. Logic sLGF, again, results from GF by allowing limited deviations of restriction (2). Should sLGF turn out to be decidable, suitable non-relativizing guards might gain some credibility as features weakly explaining decidability in the realm of modal-like logics.

References

Andréka, H., Németi, I., and van Benthem, J., 1998, "Modal languages and bounded fragments of predicate logic," *J. Phil. Logic*, 27:217–74.

Areces, C., Blackburn, P. & Marx, M., 1999, "A road-map on complexity for hybrid logics." In *Lecture Notes in Computer Science*, (Vol. 1683, pp. 307–321), Berlin: Springer.

Barwise, J., 1979, "On branching quantifiers in English," *Journal of Philosophical Logic*, 8: 4780.

Berwanger, D. & and Grädel, E., 2001, "Games and model checking for guarded logics." In *Logic for Programming, Artificial Intelligence, and Reasoning 8th International Conference, LPAR 2001 Havana, Cuba, December 37, 2001 Proceedings*, LNCS Vol. 2250, R. Nieuwenhuis and A. Voronkov(Eds.), Springer, 70–84.

Blackburn, P., de Rijke, M., & Venema, Y., 2002, *Modal Logic*, Cambridge: Cambridge University Press.

Bradfield, J., 2000, "Independence: logics and concurrency." In *Lecture Notes in Computer Science*, (Vol. 1862, pp. 247–261), London: Springer.

Bradfield, J., & Fröschle, S., 2002, "Independence-friendly modal logic and true concurrency," *Nordic Journal of Computing*, 9, 102–117.

Burgess, J. B., 1982, "Axioms for tense logic I. 'Since' and 'Until'," *Notre Dame Journal of Formal Logic* 23(4), 367–374.

Börger, E., Grädel, E., & Gurevich, Y., 1997, *The Classical Decision Problem*, Berlin: Springer.

Cameron, P., & Hodges, W., 2001, "Some combinatorics of imperfect information," *Journal of Symbolic Logic*, 66, 673–684.

Ebbinghaus, H.-D., & Flum, J., 1999, *Finite Model Theory*, Berlin: Springer.

Ehrenfeucht, A., 1959, "Decidability of the theory of the linear ordering relation," *Notices of the American Mathematical Society* Vol, 6 No. 3, 538–556.

Emerson, E.A., & Jutla, C., 1988, "The complexity of tree automata and logics of programs." In *Proceedings of the 29th IEEE Symposium on Foundations of Computer Science*, 368–377.

Fürer, M., 1981, "Alternation and the Ackermann case of the decision problem," *L'Enseignment Mathematique*, 27, 137–162.

Gabbay, D., 1981, "Expressive functional completeness in tense logic." In U. Mönnich (Ed.), *Aspects of Philosophical Logic*, Reidel, Dordrecht, 91–117.

Ganzinger, H., Meyer, C., Veanes, M., 1999, "The two-variable guarded fragment with transitive relations." In *Fourteenth Annual IEEE Symposium on Logic in Computer Science, Trento, 1999*, 24–34.

Grädel, E., 1999a, "On the restraining power of guards," *Journal of Symbolic Logic* 64(4), 1719–1742.

Grädel, E., 1999b, "The decidability of guarded fixed point logic." In J. Gerbrandy, M. Marx, M. de Rijke, and Y. Venema (Eds.), *JFAK. Essays Dedicated to Johan van Benthem on the Occasion of his 50th Birthday*, Amsterdam UP, 1999.

Grädel, E., 2001, "Why are modal logics so robustly decidable?" In *Current Trends in Theoretical Computer Science. Entering the 21st Century* (G. Paun, G. Rozenberg, and A. Salomaa, Eds.), pp. 393–408, Singapore: World Scientific. References to the file retrieved December 3, 2011, from `http://www.logic.rwth-aachen.de/pub/graedel/Gr-trends01.ps`.

Grädel, E., Kolaitis, P., and Vardi, M., 1999, "On the decision problem for two-variable first-order logic," *Bulletin of Symbolic Logic* 3(1), 53–69.

Grädel, E., and Otto, M., 1999, "On logics with two variables," *Theor. Comp. Science* 224, 73–113.

Grädel, E., Otto, M., and Rosen, E., 1997, "Undecidability results on two-variable logics." In *Proceedings of 14th Symposium on Theoretical Aspects of Computer Science STACS '97 Lecture Notes in Computer Science No. 1200*, Springer, pp. 249-260.

Grädel, E., & Walukiewicz, I., 2006, "Positional determinacy of games with infinitely many priorities," *Logical Methods in Computer Science*, 2, 1–22.

Hintikka, J., 1996, *Principles of Mathematics Revisited*, Cambridge: Cambridge University Press.

Hintikka, J., 1997. "No Scope for Scope?," *Linguistics and Philosophy*, 20, 515–544.

Hintikka, J., and Kulas, J., 1983, *The Game of Language: Studies in Game-Theoretical Semantics and Its Applications*, Dordrecht: Reidel.

Hintikka, J., and Kulas, J., 1985, *Anaphora and Definite Descriptions: Two Applications of Game-Theoretical Semantics*, Dordrecht: Reidel.

Hintikka, J., & Sandu, G., 1989, "Informational independence as a semantical phenomenon." In J. Fenstad, I. Frolov, & R. Hilpinen (Eds.), *Logic, Methodology and Philosophy of Science*, (Vol. 8, pp. 571–589), Amsterdam: Elsevier.

Hodges, W., 1983, "Elementary predicate logic." In Handbook of Philosophical Logic, Vol. 1, D. M. Gabbay, and F. Guenthner (Eds.), Dordrecht: Reidel, pp. 1131.

Hodges, W., 1997a, "Compositional semantics for a language of imperfect information," *Logic Journal of the IGPL*, 5, 539–563.

Hodges, W., 1997b, "Some strange quantifiers." In *Lecture Notes in Computer Science*, (Vol. 1261, pp. 51–65), London: Springer.

Hodges, W., 2006, "Logic and Games." In *The Stanford Encyclopedia of Philosophy (Summer 2006 Edition)*, E. Zalta (Ed.), URL = <http://plato.stanford.edu/archives/sum2006/entries/logic-games/>.

Hodges, W., 2007, "Logics of Imperfect Information: Why Sets of Assignments?" In *Interactive Logic*, J. van Benthem, D. M. Gabbay, and B. Löwe (Eds.), Amsterdam: Amsterdam University Press, pp. 117–133.

Horrocks, I., Hustadt, U., Sattler, U. and Schmidt, R., 2007, "Computational modal logic." In *Handbook of Modal Logic*, P. Blackburn, J. van Benthem and F. Wolter (Eds.), Elsevier, 181–245.

Hyttinen, T., & Tulenheimo, T., 2005, "Decidability of IF modal logic of perfect recall." In R. Schmidt, I. Pratt-Hartmann, M. Reynolds, & H. Wansing (Eds.), *Advances in Modal Logic*, (Vol. 5, pp. 111–131), London: KCL Publications.

Immerman, N., 1999, *Descriptive Complexity*, New York: Springer.

Kieronski, E. & Otto, M., 2005, "Small substructures and decidability issues for two-variable first-order logic," *Proceedings of 20th IEEE Symposium on Logic in Computer Science LICS 2005*, 448–457.

Ladner, R.E., 1977, "The computational complexity of provability in systems of modal propositional logic," *SIAM Journal on Computing*, 6(3), 467–480.

Mann, A.L., Sandu, G., and Sevenster, M., 2011, *Independence-Friendly Logic: A Game-Theoretic Approach*, Cambridge: Cambridge University Press.

Rabin, M.O., 1969, "Decidability of second-order theories and automata on infinite trees," *Transactions of the American Mathematical Society*, vol. 141, 1–35.

Rabinovich, A., 2010, "Temporal logics over linear time domains are in PSPACE." In *RP'10, volume 6227 of Lecture Notes in Computer Science*, Springer, 29–50.

Reynolds, M., 2003, "The complexity of the temporal logic with 'until' over general linear time," *Journal of Computer and System Sciences*, 66(2), 393–426.

Reynolds, M., 2010, "The complexity of decision problems for linear temporal logics," *Journal of Studies in Logic*, 3(1), 19–50.

Sandu, G., & Hintikka, J., 2001, "Aspects of compositionality," *Journal of Logic, Language and Information*, 10, 49–61.

Scott, D., 1962, "A decision method for validity of sentences in two variables" (abstract), *Journal of Symbolic Logic* 27, 477.

Sevenster, M., 2006, *Branches of Imperfect Information*. Ph.D. thesis, University of Amsterdam.

Sevenster, M., 2009, "Model-Theoretic and Computational Properties of Modal Dependence Logic," *Journal of Logic and Computation* 19(6), 1157–1173.

Sevenster, M., 2010, "Decidability of Independence Friendly Modal Logic," *Review of Symbolic Logic* 3(3), 415–441.

Szwast, W. & Tendera, L., 2004, "The guarded fragment with transitive guards," *Annals of Pure and Applied Logic*, 128, 227–276.

Tulenheimo, T., 2003, "On IF modal logic and its expressive power." In P. Balbiani, N.-Y. Suzuki, F. Wolter, & M. Zakharyaschev (Eds.), *Advances in Modal Logic*, (Vol. 4, pp. 475–498), London: KCL Publications.

Tulenheimo, T., 2004, *Independence-Friendly Modal Logic: Studies in its Expressive Power and Theoretical Relevance*, Ph.D. thesis, University of Helsinki.

Tulenheimo, T., 2009, "Hybrid logic meets IF modal logic," *Journal of Logic, Language and Information* 18(4), 559–591.

Tulenheimo, T., & Rebuschi, M., 2009, "Equivalence criteria for compositional IF modal logics," *Electronic Notes in Theoretical Computer Science* 231, 153–173.

Tulenheimo, T., & Sevenster, M., 2006, "On modal logic, IF logic and IF modal logic." In G. Governatori, I. Hodkinson, & Y. Venema (Eds.), *Advances in Modal Logic*, (Vol. 6, pp. 481–501), London: College Publications.

Tulenheimo, T., & Sevenster, M., 2007, "Approaches to independence friendly modal logic." In J. van Benthem, D. Gabbay, & B. Löwe (Eds.), *Interactive Logic*, (Texts in Logic and Games Vol. 1, pp. 247–280), Amsterdam: Amsterdam University Press.

Vardi, M., 1998, "Why is modal logic so robustly decidable?" In *Descriptive complexity and finite models*, *Series in Discr. Math. & Theor. Comp. Science* vol. 31, 149-84, Amer. Math. Society. References to the file retrieved December 3, 2011, from `http://www.cs.rice.edu/~vardi/papers/dimacs96.ps.gz`.

van Benthem, J.F.A.K., 1995, "Temporal logic." In *Handbook of Logic in Artificial Intelligence and Logic Programming*, D. Gabbay, C. Hoggar and J. Robinson (Eds.), vol. 4. Oxford University Press, 241–350.

van Benthem, J.F.A.K., 1997, "Dynamic bits and pieces," Technical Report LP-97-01, Institute for Logic, Language and Computation, University of Amsterdam.

Väänänen, J., 1999, *A Short Course on Finite Model Theory*, Department of Mathematics, University of Helsinki, `http://www.math.helsinki.fi/logic/people/jouko.vaananen/shortcourse.pdf`.

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

Väänänen, J., 2008, "Modal Dependence Logic." In K. Apt & R. van Rooij (eds.): *New Perspectives on Games and Interaction*, (Texts in Logic and Games Vol. 4, pp. 237–254), Amsterdam: Amsterdam University Press.