

GUARDS, BOUNDS, AND GENERALIZED SEMANTICS

Johan van Benthem, Amsterdam & Stanford

March 2005

Abstract

Some initial motivations for the Guarded Fragment still seem of interest in carrying its program further. First, we stress the equivalence between two perspectives: (a) satisfiability on standard models for guarded first-order formulas, and (b) satisfiability on general assignment models for arbitrary first-order formulas. In particular, we give a new straightforward reduction from the former notion to the latter. We also show how a perspective shift to general assignment models provides a new look at the fixed-point extension $LFP(FO)$ of first-order logic, making it decidable. Next, we relate guarded syntax to earlier quantifier restriction strategies for the purpose of achieving effective axiomatizability in second-order logic – pointing at analogies with 'persistent' formulas, which are essentially in the Bounded Fragment of many-sorted first-order logic. Finally, we look at some further unexplored directions, including the systematic use of 'quasi-models' as a semantics by itself.

1 Basics of the Guarded Fragment

1.1 *Guarded syntax*

The Guarded Fragment of Andr eka, van Benthem & N emeti 1998 is a decidable part of first-order syntax with a semantic philosophy: quantifiers only access the total domain of individual objects 'locally' by means of predicates over objects. But there is more to the motivation and ambitions of guarding, as will be shown in this paper. But first, we make a quick tour of some known results and proof methods.

Here are some syntactic preliminaries. In what follows, mostly for convenience, we consider only languages with predicate symbols and variables: no function symbols or identity predicates occur. But we do allow so-called polyadic first-order quantifiers $\exists \mathbf{x}\phi$, $\forall \mathbf{x}\phi$ over tuples of variables \mathbf{x} , with their obvious interpretation. Finally, we also use polyadic notations $[\mathbf{u}/\mathbf{y}]\phi$ for simultaneous substitutions. These are taken in the standard syntactic sense that the substitution is performed provided the \mathbf{u} are free for the \mathbf{y} . If not, some suitable alphabetic variant is taken first for ϕ .

Our key idea is that objects \mathbf{y} can only be introduced relative to given objects \mathbf{x} , as expressed by a 'guard atom' $G(\mathbf{x}, \mathbf{y})$ where objects can occur in any order and multiplicity – and that the subsequent statement refers only to those guarded \mathbf{x}, \mathbf{y} .

Definition 1 Guarded Formulas.

Guarded formulas are all those constructed according to the syntax rules

$$\text{atoms } Px \mid \neg \mid \vee \mid \exists \mathbf{y} (G(\mathbf{x}, \mathbf{y}) \ \& \ \phi(\mathbf{x}, \mathbf{y}))$$

Here, bold-face \mathbf{x}, \mathbf{y} indicate finite tuples of variables, and G is a predicate letter. *Loosely guarded formulas* allow a conjunction of atoms $\gamma(\mathbf{x}, \mathbf{y})$ instead of $G(\mathbf{x}, \mathbf{y})$ in the quantifier clause, provided each variable from \mathbf{y} co-occurs with each variable from \mathbf{x} , \mathbf{y} in at least one atom of $\gamma(\mathbf{x}, \mathbf{y})$. The set of all guarded first-order formulas is the *Guarded Fragment GF*. There is also a *Loosely Guarded Fragment LGF*. ♣

1.2 Decidability via quasi-models

The initial motivating result was that guarding quantifiers leads to decidability.

Theorem 1 GF and LGF are decidable.

In what follows, we stick mostly to GF . We shall have occasion to refer to the proof of Theorem 1, and hence we reproduce its outline here.

Proof The first observation is that truth of first-order formulas in any model is witnessed in some finite syntactic object, called a 'quasi-model'. Let formula ϕ be true in standard model \mathbf{M} . Let V be the finite set of variables occurring in ϕ – free or bound. In effect, we are inside a finite-variable fragment of first-order logic here. Next, we restrict attention to the finite set Sub_ϕ consisting of ϕ and its subformulas, while also closing under simultaneous substitutions $[u/y]$ using only variables in V , that do not change syntactic forms. This is feasible because of the following simple observation, provable by some syntactic manipulation:

Lemma 1 Finite-variable fragments are closed under simultaneous substitutions.

Now each variable assignment s on \mathbf{M} verifies a set Δ_s of formulas from Sub_ϕ with special properties, that we call a *type*. Note that any model realizes at most finitely many types. A 'quasi-model' is a finite set of types with some properties and mutual relations that obviously hold if the source is indeed some model \mathbf{M} .

Definition 2 Quasi-models.

Let F be the finite set of all formulas of length $\leq |\phi|$ that use only variables from V . Note that $\phi \in F$ and F is closed under taking subformulas and the above 'alphabetic variants' used with substitutions. An F -type is a subset Δ of F which satisfies

- | | | | | | |
|-----|----------------------------|---------|--|----------|------------------------|
| (a) | $\neg\psi \in \Delta$ | iff | $\text{not } \psi \in \Delta$ | whenever | $\neg\psi \in F$ |
| (b) | $\psi \vee \xi \in \Delta$ | iff | $\psi \in \Delta \text{ or } \xi \in \Delta$ | whenever | $\psi \vee \xi \in F$ |
| (c) | $[u/y]\psi \in \Delta$ | only if | $\exists y \psi \in \Delta$ | whenever | $\exists y \psi \in F$ |

Next, we write $\Delta =_y \Delta'$ if Δ, Δ' share the same formulas with all their free variables disjoint from y . A *quasi-model* is a set of F -types S such that

- (d) for each $\Delta \in S$ and each formula $\exists y \psi \in \Delta$,
there is a type $\Delta' \in S$ with $\psi \in \Delta'$ and $\Delta =_y \Delta'$.

We say that ϕ holds in a quasi-model if $\phi \in \Delta$ for some Δ in this quasi-model. ♣

Clearly, this definition justifies the following assertion:

Lemma 2 If a first-order formula has a model, it is true in some quasi-model.

The converse is not true for all first-order formulas, but it does hold for GF .

Lemma 3 If a guarded formula has a quasi-model, then it has a standard model.

The key fact is that quasi-models can be 'unraveled' to tree-like standard models without affecting truth values of guarded formulas in their set F . Details can be found in Andréka, van Benthem & Némethi 1998 – but they will not matter here. The same tree-model construction also leaves truth values of loosely guarded formulas invariant. Decidability of GF or LGF now follows because we can test satisfiability for arbitrary (loosely) guarded formulas ϕ by testing for the existence of a quasi-model for ϕ whose size is effectively bounded by the length of ϕ . ♣

This decision procedure can be adapted easily to give an optimal complexity result (Grädel 1999B). Satisfiability is $2EXPTIME$ -complete for guarded formulas, and it is $EXPTIME$ -complete for GF with a fixed bound on the arities of predicates.

1.3 Other metaproperties

The Guarded Fragment was meant to serve several purposes at once. On the one hand its complexity is low enough to be decidable, while it is expressive enough to generalize most common modal languages. This demonstrates the balance sought in all good modal-like languages. Another desirable feature concerns its meta-theory. Basic modal logic resembles first-order logic in all its meta-properties, even 'existential' ones that do not follow from just being a sublanguage, such as Craig Interpolation, Beth Definability, and the standard model-theoretic preservation theorems. GF shares this good behaviour to some extent, witness the Los-style preservation theorem for

submodels given in Andréka, van Benthem & Németi 1998. Cf. also van Benthem 2001 on GF as an instrument for finding out 'what makes modal logic tick'. But subsequent work has shown that the picture is somewhat mixed. Beth Definability holds (Hoogland, Marx & Otto 1999), but Craig Interpolation fails in its strong general form – though it remains valid when we view guard predicates as part of the logical vocabulary (Hoogland & Marx 2002). On the positive side again, GF has shown logical resilience in other ways, not foreseen when it first appeared. A striking example is the result in Grädel 1999C, that the extension $LFP(GF)$ of GF with fixed-point operators μ, ν remains decidable, whereas validity for the version $LFP(FO)$ for all of first-order logic is non-axiomatizable – indeed non-arithmetical. Incidentally, Grädel 1999A also determines the complexity for $LFP(GF)$.

In the remainder of this paper, we pursue some other, less generally appreciated aspects of guarded syntax, that also played a role at the time of its invention.

2 Guards and General Assignment Models

2.1 *Restricted syntax versus generalized semantics*

Giving each quantifier a guard may be viewed as a syntactic restriction, refraining from all unbounded quantification. In this sense, GF is indeed a fragment of FOL . But there is also another perspective, where this move rather represents a semantic generalization. We now assume that quantification will normally take place in 'structured domains', where access from one group of objects to another must go via some connecting relation R of some appropriate arity. Binary modal accessibility is a typical example. Standard models are the special case with R the universal relation. Informally, then, there seems to be an analogy between

- (a) using guarded formulas over standard models, and
- (b) using arbitrary first-order formulas over suitably generalized models.

This can be made more precise using the following semantics, going back to the general relativization technique for algebraic models in Németi 1985.

2.2 *Modal models and general assignment models*

For a start, by simple inspection of the standard truth definition, it is easy to interpret a complete first-order language on abstract *modal models*

$$\mathbf{M} = (S, \{R_x\}_{x \in \text{VAR}}, I)$$

with S a set of 'states', R_x a binary accessibility relation between states for each variable x , and I an interpretation function giving a truth value to each atomic formula in each state s . This is a huge extension of standard first-order semantics, where no domain of 'individual objects' need be present underpinning the states. Indeed, it has often been observed that quantifiers are modalities, with

$$\mathbf{M}, s \models \exists x \phi \quad \text{iff} \quad \text{for some } t: R_x s t \text{ and } \mathbf{M}, t \models \phi.$$

Thus, first-order logic is a poly-modal logic with an existential modality $\exists x$. This gives us a broad space for generalized semantics of many sorts.

More concrete is the following semantics which merely takes away the existential assumption of 'fullness' from standard Tarski models for *FOL*.

Definition 3 General assignment models.

A *general assignment model* is an ordered pair (\mathbf{M}, \mathbf{V}) with \mathbf{M} a standard first-order model with domain D and interpretation function I , and \mathbf{V} any non-empty set of assignments on \mathbf{M} , i.e., a subset of D^{VAR} . The first-order language is interpreted as usual, now at triples $\mathbf{M}, \mathbf{V}, s$ with $s \in \mathbf{V}$ – with the following clause for quantifiers:

$$\mathbf{M}, \mathbf{V}, s \models \exists x \phi \quad \text{iff} \quad \text{for some } t \in \mathbf{V}: s =_x t \text{ and } \mathbf{M}, \mathbf{V}, t \models \phi$$

Here $=_x$ is the standard relation between assignments of identity up to x -values. ♣

Assignment gaps in these models model the natural phenomenon of dependencies between variables: changes in value for one variable x may induce, or at least be correlated with, changes in value for another variable y . General assignment models also support *new vocabulary*, reflecting distinctions beyond standard first-order logic. Examples are irreducibly *polyadic quantifiers* $\exists \mathbf{x}$ binding tuples of variables \mathbf{x} , with the following truth condition:

$$\mathbf{M}, \mathbf{V}, s \models \exists \mathbf{x} \phi \quad \text{iff} \quad \text{for some } t \in \mathbf{V}: s =_{\mathbf{x}} t \text{ and } \mathbf{M}, \mathbf{V}, t \models \phi$$

This time, $=_{\mathbf{x}}$ is identity between assignments up to values for all the variables in \mathbf{x} . E.g., in standard first-order logic, the notation $\exists x y \phi$ is just short-hand for $\exists x \exists y \phi$ or $\exists y \exists x \phi$ in any order. But in *GAM*-semantics, these two expressions are no longer equivalent, as not all 'intermediate assignments' for x - or y -shifts need be present – and they are both non-equivalent to $\exists xy$, as defined just now. Moreover, one can also interpret single or polyadic *substitution* operators directly in this style:

$$\mathbf{M}, \mathbf{V}, s \models [y/x] \phi \text{ iff } s[x:=s(y)] \in \mathbf{V} \ \& \ \mathbf{M}, \mathbf{V}, s[x:=s(y)] \models \phi$$

The resulting logic 'CRS' over general assignment models has been much studied (Marx & Venema 1997). Its valid principles consist of the standard axioms for poly-modal *S5* plus all atomic 'locality principles' $(\neg)Px \rightarrow \forall y (\neg)Py$ with $x \frown y = \emptyset$. Not universally valid in the above models, however, are the following two principles:

- (i) $[u/y]\psi \rightarrow \exists y \psi$ with u free for y in ψ *Existential Generalization*
(ii) $\phi(x) \rightarrow \forall y \phi(x)$ with no y free in $\phi(x)$ *Full Locality*

These failures reflect the special handling of variables in models where not all assignments need be available. All of x, y, z, \dots then acquire a sort of 'individuality', due to their possibly different interactions with other variables. We omit technical details here (an excellent source is Némethi 1996). Instead, we turn to the connection between *GAM*-semantics for *FOL* and standard semantics for *GF*.

2.3 Reducing *GAM* logic to *GF*

The following result is proved in Andréka, van Benthem & Némethi 1998, Section 5.

Definition 4 Guarded translation.

Consider any k -variable language $L\{x_1, \dots, x_k\}$. Let R be a new k -ary predicate. The translation *guard* takes k -variable first-order formulas ϕ to guarded first-order formulas $guard(\phi)$ by relativizing all quantifiers to the same atom $Rx_1\dots x_k$. This translation works for polyadic first-order quantifiers just as well as single ones – and it even extends to the above substitution operators, if desired. There is also a matching semantic operation of model expansion. Let (\mathbf{M}, \mathbf{V}) be any general assignment model for $L\{x_1, \dots, x_k\}$ – without the new predicate R . The standard model $GUARD(\mathbf{M}, \mathbf{V})$ is \mathbf{M} viewed as a standard model, and expanded with the following interpretation:

$$R(d_1, \dots, d_k) \text{ iff the assignment } x_i := d_i \ (1 \leq i \leq k) \text{ is in } \mathbf{V}. \quad \clubsuit$$

The following is easy to prove by induction on first-order formulas:

Lemma 4 For all available assignments s in \mathbf{V} , and all k -variable formulas ϕ ,
 $\mathbf{M}, \mathbf{V}, s \models \phi$ iff $GUARD(\mathbf{M}, \mathbf{V}), s \models guard(\phi)$

Here is a reduction of *GAM*-semantics to the Guarded Fragment.

Theorem 2 For all first-order k -variable formulas ϕ , the following are equivalent:

- (a) ϕ is satisfiable in general assignment models,
(b) $Rx_1\dots x_k \wedge guard(\phi)$ is satisfiable in standard models.

Proof From (a) to (b), Lemma 4 supplies the reason. For the converse, suppose that $Rx_1 \dots x_k \wedge \text{guard}(\phi)$ has a standard model \mathbf{M} under some assignment s . Now define a general assignment model (N, \mathbf{V}) by retaining only those assignments on \mathbf{M} whose values for x_1, \dots, x_k stand in the relation $R_{\mathbf{M}}$. These include the assignment s itself. Then it is easy to see that $N, \mathbf{V}, s \models \phi$ as with Lemma 4. ♣

The translation also works directly for the full first-order language without the k -restriction, by a slightly modified translation. The converse direction was left as an open question in Andréka, van Benthem & Némethi 1998, but we tackle it here. Remark 1 below briefly discusses an earlier solution by Marx.

2.4 Reducing GF to GAM logic

We need a translation again. This time, it is not compositional in the earlier sense. The reason is the earlier failure of Existential Generalization (i) and Full Locality (ii) in general assignment models. We need these principles for some finite set of relevant formulas in the proof to follow, and hence we put them into the translation.

Definition 5 GAM translation.

Let ϕ be any guarded first-order formula with a total set of variables $\mathbf{x} = x_1, \dots, x_k$. Let $\text{set-up}(\phi)$ be the finite conjunction of all formulas of the following form

- (i)' $\forall \mathbf{x} ([\mathbf{u}/\mathbf{y}]\psi \rightarrow \exists \mathbf{y} \psi)$ where $\mathbf{u}, \mathbf{y} \subseteq \mathbf{x}$ and $\psi(\mathbf{z})$ is a subformula of ϕ
- (ii)' $\forall \mathbf{x} (\psi(\mathbf{z}) \rightarrow \forall \mathbf{y} \psi(\mathbf{z}))$ where $\mathbf{z}, \mathbf{y} \subseteq \mathbf{x}$ with \mathbf{z} disjoint from \mathbf{y} ,
and $\psi(\mathbf{z})$ is a subformula of ϕ

The, not necessarily guarded, formula $\text{gam}(\phi)$ is the conjunction $\phi \wedge \text{set-up}(\phi)$. ♣

In particular, the prefixed polyadic universal quantifier $\forall \mathbf{x}$ running over all relevant variables makes sure that the implications (i)', (ii)' hold throughout any general assignment model which has $\text{set-up}(\phi)$ true at any assignment at all.

Theorem 3 For all guarded formulas ϕ , the following are equivalent:

- (a) ϕ is satisfiable in standard models,
- (b) $\text{gam}(\phi)$ is satisfiable in general assignment models.

Proof From (a) to (b), it suffices to note that any standard model for ϕ also satisfies $\text{gam}(\phi)$, since the formulas in the second conjunct are universally valid. And standard models are general assignment models with a full set of assignments.

Next, from (b) to (a), let $\mathbf{M}, \mathbf{V}, s \models \text{gam}(\phi)$. As in Section 1.2, this situation induces a quasi-model for ϕ . Recall that the relevant formulas are all subformulas of ϕ plus their alphabetic variants with variables from \mathbf{x} . The types of the quasi-model are now all sets of relevant formulas true at the assignments in \mathbf{V} . We must check the four clauses of Definition 2. Here, the first two follow directly by the truth definition for Boolean operations. Next, the third existential generalization clause holds for all types because of the truth of conjunct (i)' in $\text{set-up}(\phi)$. And finally, the special 'witness clause' (d) for existential quantifiers in suitably related types holds because of the truth condition for the existential quantifier in general assignment models plus the true transfer condition (ii)' in $\text{set-up}(\phi)$. Thus, the given guarded formula ϕ has a quasi-model – and hence it also has a standard model by Lemma 3. \clubsuit

The same reasoning extends to the loosely-guarded fragment *LGF*. One can also recast Theorem 4 in other interesting ways, as is done in Andr eka & N emeti 2005. The latter uses conjunctions of the special formulas $\forall x ([u/y]\psi \rightarrow \exists y \psi)$, $\forall x (\psi(z) \rightarrow \forall y \psi(z))$ in Definition 5 to raise further issues comparing standard first-order validity and *GAM*-validity. We do not pursue these here.

Remark 1 Marx 2001, which came to our attention after completing this paper, also addresses the issue of semantic relativization versus syntactic guarding. Following earlier results from 1997, Marx gives a two-way reduction like our Theorems 2, 3 – but instead of 'loading' the translation *gam* as we have done, he works with special general assignment models satisfying Full Locality plus a 'Local Cube' property guaranteeing the validity of existential generalization. Another illuminating result in the paper is a model-theoretic preservation theorem for the 'packed fragment' of first-order logic (roughly, a slight extension of *LGF*). This is shown to be the largest fragment of first-order logic that is insensitive between evaluation in standard models and models relativized to some 'tolerance relation'.

2.5 Consequences: general semantics for fixed-point languages

Sections 2.3, 2.4 show that we can think in two modes about first-order logic: either via syntactic quantifier guards, or via generalized semantics. This duality transports insights from one area to another. Here is one example that may be of interest. Consider the *fixed-point version* *LFP(FO)* of full first-order logic. This language extends the usual inductive formation rules for first-order syntax with an operator

$$\mu P, \mathbf{x} \bullet \phi(P, Q, \mathbf{x})$$

where P may occur only positively in $\phi(P, \mathbf{Q}, \mathbf{x})$, and \mathbf{x} is a tuple of variables of the right arity for P . The relevant predicates are the smallest fixed-points of the following monotone set operation on predicates in any given model \mathbf{M} :

$$F_{\phi}^{\mathbf{M}} = \lambda P \bullet \{d \text{ in } M \mid (\mathbf{M}, P), d \models \phi(P, \mathbf{Q})\}$$

We saw that Grädel 1999B showed the fixed-point version $LFP(GF)$ of the Guarded Fragment to be decidable, unlike the much stronger $LFP(FO)$ itself. But then, this phenomenon can also be understood in dual mode, on general assignment models. In algebraic terms, the 'relativized' version of $LFP(FO)$ is well-behaved!

Formulas ϕ in the language $LFP(FO)$ now need a bit more care, since variables are less 'anonymous' in general assignment models, as we noted before. In particular, when defining a predicate $\mu P, \mathbf{x} \bullet \phi(P, \mathbf{Q}, \mathbf{x})$, the particular variables \mathbf{x} matter. This suggests that we are only defining values for the specific atom $P\mathbf{x}$, whereas variants such as $P\mathbf{y}$ must be viewed as substitution instances $[y/x]P\mathbf{x}$. With this understanding, we can give a definition of semantic evaluation as before.

Definition 6 GAM fixed-point evaluation.

Formulas ϕ in the above language induce the following map in general assignment models (\mathbf{M}, \mathbf{V}) with some given assignment s for the free variables in ϕ :

$$F^{M, s}_{\phi} = \lambda P \bullet \{d \text{ in } M \mid s[x:=d] \in \mathbf{V} \ \& \ (\mathbf{M}, P), s[x:=d] \models \phi(P, \mathbf{Q})\}$$

Smallest and greatest fixed-points are then defined as usual. ♣

Example 1 Transitive closure of guard predicates.

Consider the fixed-point formula $\phi = \mu P, \mathbf{x} \bullet Q\mathbf{x} \vee \exists y [y/x]P\mathbf{x}$. Its approximation sequence as defined above starts with the empty set for P , and it ends by stage ω where iteration of the map $F^{M, s}_{\phi}$ produces nothing new. Here are some stages:

$$\begin{aligned} P^0 &= \emptyset \\ P^1 &= \{d \mid s[x:=d] \in \mathbf{V} \ \& \ (\mathbf{M}, P^0), s[x:=d] \models Q\mathbf{x} \vee \exists y [y/x]P\mathbf{x}\} \\ &= \{d \mid s[x:=d] \in \mathbf{V} \ \& \ Q(d)\} \\ P^2 &= \{d \mid s[x:=d] \in \mathbf{V} \ \& \ (\mathbf{M}, P^1), s[x:=d] \models Q\mathbf{x} \vee \exists y [y/x]P\mathbf{x}\} \\ &= \{d \mid s[x:=d] \in \mathbf{V} \ \& \ (Q(d) \vee \text{for some object } e: \\ &\quad s[x:=d][y:=e] \in \mathbf{V} \ \& \ s[x:=e][y:=e] \in \mathbf{V} \ \& \ Q(e))\} \end{aligned}$$

Iteratively, one computes the set of all objects d for which there is an object e satisfying Q reachable from d in the transitive closure of the following relation:

$$R_s ab \quad \text{iff} \quad s[x:=a][y:=b], s[x:=b][y:=b] \in \mathbb{V}. \quad \clubsuit$$

We forego further details here – but note how these fixed-point computations bring to light the hidden dependency structure of the relevant general model.

Theorem 4 $LFP(FO)$ is decidable over general assignment models.

Proof First, the translation *guard* in Section 2.3 from arbitrary formulas to guarded ones can easily be extended to the language with added fixed-point operators. Next, translations are inside the language of $LFP(GF)$, and fixed-point evaluation must stay inside the set of tuples satisfying the guard relation R . Lemma 4 will still go through then – and Grädel's earlier-mentioned result supplies the decidability. \clubsuit

The more general direction suggested by these results is a systematic use of general assignment models in abstract model theory for extensions of first-order logic.

3 Guards in second-order logic

3.1 Lowering complexity in higher-order logic

Guarding was introduced for making first-order logic decidable – but strategies for lowering complexity of logics have been around at least since Henkin introduced his general models for second-order logic. These make the latter system effectively axiomatizable, and indeed equivalent to a two-sorted first-order logic over objects and sets with a primitive relation \in between them. On these models, one can then impose any family of Comprehension axioms for set existence as additional first-order axioms. Van Benthem 1996A asks whether the guarding and general model strategies are related. What follows is a partial answer, triggered by another precursor of GF , viz. the use of 'persistent formulas' in so-called 'extended pragmatic languages' in Montague 1970. These languages were inspired by earlier work of Orey in the 1950s on fragments of higher-order logics, and in practice, they can deal with much of what one would naturally want to say in natural language.

3.2 Persistent formulas

Montague observed (together with his student Kamp) that, inside the full language of second-order logic, the following set of formulas is particularly well-behaved.

Definition 7 Persistent second-order formulas.

Persistent second-order formulas have all their second-order quantifiers relativized in the format $\exists X(R(X, \mathbf{Y}, \mathbf{z}) \wedge \phi)$, where R is some third-order predicate, \mathbf{Y} a tuple of predicate variables, and \mathbf{z} a tuple of object variables. The *persistent fragment* of second-order logic with just these forms of quantification is called *PSOL*. ♣

Here, third-order predicates express properties of sets or relations. Examples of third-order type-theoretic objects are generalized quantifiers (see below), as well as other higher-order constructions in mathematics or natural language semantics.

Theorem 5 *PSOL* is effectively axiomatizable.

Proof The key observation is the following.

Lemma 5 *PSOL*-formulas have standard models iff they have general models.

The direction from left to right is obvious. Conversely, consider any general model $\mathbf{M}, \mathbf{P}, \mathbf{P}, s \models \phi$, where \mathbf{P} is the restricted range of predicates. Next, consider the full standard model $\mathbf{M}^+ = (\mathbf{M}, \mathbf{P}, s)$ which allows for the family *PRED* of all predicates. Moreover, just copy all third-order predicates R used in the bounded quantifiers of ϕ : so that they only relate predicates from the earlier family \mathbf{P} and possibly individual objects. Now it is easy to prove by induction that

$$\mathbf{M}, \mathbf{P}, \mathbf{P}, s \models \phi \quad \text{iff} \quad \mathbf{M}, \text{PRED}, \mathbf{P}, s \models \phi, \text{ for all } \text{PSOL}\text{-formulas } \phi.$$

The key inductive step is that for the existential quantifiers $\exists X(R(X, \mathbf{Y}, \mathbf{z}) \wedge \phi)$ over predicates. From left to right, this is obvious, since *PRED* extends \mathbf{P} . From right to left, we use the truth of the bounding predicate R in \mathbf{M}^+ as defined above to conclude that the predicate X must belong to \mathbf{P} .

Finally, since validity in general models is effectively axiomatizable by the usual Henkin-style proof, the same holds for standard validity of *PSOL*-formulas. ♣

The above proof does not check whether general models satisfy Comprehension Principles $\exists Y \forall \mathbf{x} (Y\mathbf{x} \leftrightarrow \psi(\mathbf{x}, \mathbf{z}, \mathbf{P}))$ for all second-order formulas ψ . The latter serve to make sure that the logic satisfies universal instantiation in its strongest form. For our purposes here, we omit this point. The above is really a reduction from satisfiability for *PSOL* to satisfiability in two-sorted first-order logic.

3.3 *Bounded Fragment versus guarding*

The key inductive step in the proof of Lemma 5 expresses a familiar fact from first-order logic concerning *bounded formulas*, having all their quantifiers relativized to some atomic predicate. For, the above general model \mathbf{M}, \mathbf{P} was a so-called 'generated submodel' of the full model \mathbf{M}^+ with respect to its predicate subdomain – and Feferman 1969 showed that the characteristic semantic feature of bounded formulas is their *invariance for generated submodels*. More precisely, Lemma 5 also allowed for unrestricted quantifiers over objects, as object domains are the same in \mathbf{M}, \mathbf{P} and \mathbf{M}^+ . Van Benthem 1983 contains some semantic preservation theorems for these mixed settings. We now focus on the *Bounded Fragment BF*. This differs from the Guarded Fragment in allowing the more general quantification

$$\exists y (G(x, y) \ \& \ \phi(x, y, z))$$

where the formula at the end may contain new free variables. This fragment is still undecidable, but its semantic invariance behaviour has applications in arithmetic and set theory, as a way of defining suitably 'absolute' properties not affected by the difference between standard models and generalized models. Ten Cate 2005 has a modern treatment with interesting new results, including the one that *BF* equals the first-order definable part of basic modal logic with added propositional quantifiers.

For the general perspective of this paper, the relevant point is this. Like *GF*, *BF* also represents a general strategy for lowering complexity of given logical systems. Within second-order logic, it reduces complexity of validity to *RE*. It does not take things further to decidable, as we just noted – but the latter would happen if we restrict the syntax of *PSOL* still further, to just guarded quantifiers over predicates and objects. It may be of interest to see which parts of higher-order logic or set theory that motivated the uses of *BF* are even guarded in this stronger fashion. Finally, as with *GF*, the *BF* strategy occurs in two guises: either syntax restriction, or semantic model generalization. Thus, it seems that *BF* and *GF* are a natural pair.

3.4 *Further modal second-order fragments*

The preceding ideas apply to existing modal languages with a second-order flavour.

Example 2 Neighbourhood semantics

In modal neighbourhood semantics, the key semantic clause is

$$\mathbf{M}, s \models []\phi \quad \text{iff} \quad \begin{array}{l} \text{there exists some set of worlds } X \\ \text{with } Ns, X \ \& \ \forall t \in X: \mathbf{M}, t \models \phi. \end{array}$$

Here N is a given third-order relation between worlds and sets of worlds. This form is guarded, and hence the basic modal logic of these models is decidable. Moreover, one can even extend the logic while staying guarded – by introducing a second sort of formulas expressing modal properties of sets X , with operators referring to their elements. Such decidable languages can be made still stronger by using truth conditions in the loosely guarded fragment *LGF*. ♣

Another line of investigation would be guarded or loosely guarded fragments of quantified modal logic, or monadic second-order logic *MSOL*. And of course, if only *RE*-ness is the target, truth conditions in a *persistent* format will suffice!

Example 3 Generalized quantifiers.

The most exciting challenge from our viewpoint is that of *generalized quantifiers* in a modal format. It is known that adding simple quantifiers to first-order logic, such as *Most A are B* or *At least as many A are B as C* makes the language non-axiomatizable, as the standard natural numbers can be defined. Can this complexity be tamed by some sort of guarding strategy? One way of doing this uses guard relations R among objects. One might interpret, e.g.,

Most $x \bullet (\phi(x, y), \psi(x, y))$ as saying that the majority
of objects x satisfying $\phi(x, y) \wedge Rxy$ also satisfy $\psi(x, y)$.

But it is not known if this move makes the logic decidable. E.g., van der Hoek & de Rijke 1993 do axiomatize a decidable basic modal logic defining an operator "most R -successors of the current world...", but this only works on equivalence relations R , where iterated operators collapse – and things can be proved by brute force. Eric Pacuit (p.c.) has claimed decidability for the general modality "in the majority of successors".

But there is also another relevant analysis. A generalized quantifier is a third-order predicate, as with persistent formulas. Now this does not help by itself, as we are using a fixed interpretation for this predicate, not a freely assignable one as needed in the above reduction to *RE*-complexity. Nevertheless, one can try to bound or guard the second-order quantifiers occurring *in that fixed interpretation* itself. E.g., the above statement about a 'majority' says that there is some injective map from the set $\{x / \phi(x, y) \wedge \neg Rxy\}$ to the set $\{x / \phi(x, y) \wedge Rxy\}$, while there is no such map in the opposite direction. But we can bound this as follows:

There is some injective map *satisfying some third-order property* P from

$\{x/ \phi(x, y) \wedge \neg Rxy\}$ to $\{x/ \phi(x, y) \wedge Rxy\}$, but not in the opposite direction.

The property P may be motivated as restricting attention to some sort of 'available' maps, say computable ones – or maps that are simply definable from the given predicates. At least in natural language semantics, this is a natural restriction. ♣

These examples are merely meant to illustrate that both bounding and guarding strategies may be of use in second-order logic, in ways unexplored so far.

4 Further ways to go

Sections 2, 3 have just addressed a few of the more general issues that can still be raised about guarding. In this final section, we list a few more, as a short agenda.

4.1 *Guarding lower down?*

Guards make sense, not just higher up from first-order logic, but also but lower down in much more restricted formalisms. For instance, Kerdiles 2001 considers the language CG of *conceptual graphs* which has only atoms, conjunction, and existential quantifiers. The complexity of the general consequence problem between such formulas is NP , but we have that *consequence between guarded CG -formulas is in P* . This result suggests that guarding can take the 'N' out of 'NP' sometimes, but the precise extent of this phenomenon is unknown.

4.2 *Replacing modal logic by GF ?*

The Guarded Fragment considerably extends the expressive resources of basic modal logic. Moreover, it lives inside the familiar formalism of first-order logic, and that with meta-properties roughly similar to those of basic modal logic. One could ask then for a systematic replacement of the latter formalism by the much more liberal GF in all sorts of contexts. For instance, the μ -calculus is a prime example of a decidable modal fixed-point language, but why not use the full $LFP(GF)$ instead (cf. also Grädel 1999A)? Likewise, whatever can be done with standard modal bisimulation seems also feasible with the 'guarded bisimulations' of Andr eka, van Benthem & N emeti 1998. This substitution has not yet been studied in its generality, and there may still be cases of logic combination where system properties with GF may have surprises compared with basic modal logic. All these are theoretical questions, though – and there remains a much more practical concern:

4.3 *Empirical coverage of guarding*

Basic modal logic captures a certain amount of reasoning about monotonicity and distribution, and part of its attraction is that modal patterns of expression without

explicit variable-binding devices are widespread in many areas. What about guarded syntax? If we do a real empirical survey of types of expression found in standard uses of first-order logic, are we going to find a high degree of guardedness? For a start, generalized quantifiers $Q(A, B)$ in natural language always occur restricted to definable subdomains A ("all *humans* are mortal", "most *gentlemen* prefer blondes"). Still, this is not yet full guarding (cf. Section 3.4) – but rather a case of *bounding*, as there may be new free variables in the follow-up predicate B .

Perhaps a better test case than natural language discourse are formal proofs. One good way of appreciating the power and limits of guards is by looking at *tiling arguments* for undecidability. One writes a first-order sentence describing a tiling on some rectangular grid which is satisfiable iff a tiling exists on $N \times N$ with the finite set of given tiles. Examining the relevant assertions, one finds that most of them are guarded – but one crucial property is not, viz. the confluence property of a grid: going north-east can be mimicked by going east-north. The latter assertion is not guarded, and not even loosely guarded. Thus, guards fail whenever some sort of confluence is required. This border line is very significant, e.g., when designing process logics in computer science. Another rich test area for the scope of *GF* and *BF*-style fragments are the theories of space-time investigated in Andr eka, Madarasz & N emeti, to appear.

5 Conclusion

The original motivations for the Guarded Fragment were diverse, ranging from the design of decidable well-behaved modal fragments to generalized semantics, and from complexity lowering of validity in first-order logic to similar concerns in other areas. These ideas still seem alive, and they can still be pushed a little further.

6 Acknowledgement

I thank Hajnal Andr eka, Balder ten Cate, Maarten Marx, Istvan N emeti, and the anonymous *JoLLI* referee for their very helpful comments and corrections.

7 References

- H. Andr eka, J. van Benthem & I. N emeti, 1998, 'Modal Languages and Bounded Fragments of Predicate Logic', *Journal of Philosophical Logic* 27, 217–274.
- H. Andr eka, I. Hodkinson & I. N emeti, 1999, 'Finite algebras of relations are representable on finite sets', *Journal of Symbolic Logic* 64, 243–267.
- H. Andr eka, J. Madarasz & I. N emeti, to appear, 'Logics of Relativistic Space-Time', in M. Aiello, J. van Benthem & I. Pratt, eds., *Handbook of Spatial Reasoning*, Kluwer-Springer Academic Publishers, Dordrecht.
- H. Andr eka & I. N emeti, 2005, private communication.

- J. van Benthem, 1983, *Modal Logic and Classical Logic*, Bibliopolis, Napoli.
- J. van Benthem, 1996A, 'Complexity of Contents versus Complexity of Wrappings', in M. Marx, M. Masuch & L. Pólos, eds., *Arrow Logic and Multimodal Logic*, Studies in Logic, Language and Information, CSLI Publications, Stanford and Cambridge University Press, 203–219.
- J. van Benthem, 1996B, *Exploring Logical Dynamics*, CSLI Publications, Stanford.
- J. van Benthem, 1997, 'Dynamic Bits and Pieces' and Pieces', Report LP-97-01, Institute for Logic, Language and Computation, University of Amsterdam.
- J. van Benthem, 1999, 'The Range of Modal Logic', *Journal of Applied Non-Classical Logics*, vol. 9:2/3, 407–442.
- J. van Benthem, 2001, 'Modal Logic in Two Gestalts', in M. de Rijke, H. Wansing & M. Zakharyashev, eds., *Advances in Modal logic, vol. II*, Uppsala 1998, CSLI Publications, Stanford, 73–100.
- B. ten Cate, 2005, *Model Theory for Extended Modal Languages*, Dissertation, Institute for Logic, Language and Computation, University of Amsterdam.
- S. Feferman, 1969, 'Persistent and Invariant Formulas for Outer Extensions', *Compositio Mathematica* 20, 29–52.
- E. Grädel, 1999A, 'Decision Procedures for Guarded Logics', in *Automated Deduction - Proceedings CADE 16*, Lecture Notes in Computer Science 1632, Springer Verlag, Berlin, 31 – 51.
- E. Grädel, 1999B, 'On the Restraining Power of Guards', *Journal of Symbolic Logic* 64, 1719–1742.
- E. Grädel, 1999C, 'The Decidability of Guarded Fixed Point Logic', in *JFAK. Essays Dedicated to Johan van Benthem on the Occasion of his 50th Birthday* (J. Gerbrandy, M. Marx, M. de Rijke, and Y. Venema, eds.), CD-ROM <http://turing.wins.uva.nl/~j50/cdrom/>, Amsterdam University Press.
- W. van der Hoek & M. de Rijke, 1993, 'Generalized Quantifiers and Modal Logic', *Journal of Logic, Language, and Information* 2, 19–50.
- E. Hoogland & M. Marx, 2002, 'Interpolation and Definability in Guarded Fragments', *Studia Logica* 70, 373–409.
- E. Hoogland, M. Marx & M. Otto, 1999, 'Beth Definability for the Guarded Fragment', in H. Ganzinger, D. McAllester & A. Voronkov, eds., *Logic for Programming and Automated Reasoning, LPAR 6*, Springer Lecture Notes in AI 1705, 273–285, Springer, Berlin.
- G. Kerdiles, 2001, *Saying it with Pictures: a logical landscape of conceptual graphs*, Dissertation DS-2001-09, Institute for Logic, Language and

Computation, University of Amsterdam.

M. Marx, 2001, 'Tolerance Logic', *Journal of Logic, Language and Information* 10, 353–373.

M. Marx & Y. Venema, 1997, *Multi-Dimensional Modal Logic*, Kluwer, Dordrecht.

R. Montague, 1970, 'Pragmatics and Intensional Logic', *Synthese* 22, 68-94.

I. Németi, 1985, 'Cylindric-Relativized Set Algebras Have Strong Amalgamation', *Journal of Symbolic Logic* 50:3, 689-700.

I. Németi, 1995, 'Decidability of Weakened Versions of First-Order Logic', in *Logic Colloquium 92*, CSLI Publications, Stanford, 177–241.

I. Németi, 1996, 'Fine-Structure Analysis of First-Order Logic', in M. Marx, M. Masuch & L. Pólos, eds., *Arrow Logic and Multi-Modal Logic*, CSLI Publications, Stanford, 221–247.

8 Appendix Quasi-models per se

The original working methods for GF may have a broader spin-off by themselves.

In particular, *quasi-models* are a mix of modal filtration and the 'mosaics' of algebraic logic – and to some extent even semantic tableaux for first-order logic. Right now, mosaics – introduced in Németi's 1986 dissertation, with Németi 1995 as a more up-to-date reference – seem the method of choice for proving decidability in modal and algebraic logics – with contributions by Marx, Mikulas, Reynolds, and many others. But quasi-models may also be appreciated on their own. First, a quasi-model for some initial formula ϕ is a modal model \mathbf{M}_ϕ for a first-order language as it stands. The types are the worlds, there are accessibility relations $=_x$ of agreeing on all formulas having no free variables in x , and for atoms, $V(\Delta, Px) = 1$ iff $Px \in \Delta$. The following Truth Lemma is then easily proved by induction:

Lemma 6 For all $\alpha \in SUB_\phi$, and all types Δ in \mathbf{M}_ϕ , $\mathbf{M}, \Delta \models \alpha$ iff $\alpha \in \Delta$.

Thus, quasi-models are models by themselves, and this may lead to new applications of modal model theory to first-order logic. The connection between modal models and standard ones has been studied in some detail van Benthem 1996, Sections 9.8 and 9.9, with results such as the following. Let x, y stand for finite sequences of variables. The notation R_x denotes the sequential composition of accessibility relations R_x as they occur in their given order in x .

Theorem 6 An abstract modal frame $(S, \{R_x\}_{x \in VAR})$ is isomorphic to the frame of some general assignment model if and only if

the R_x are equivalence relations satisfying all 'Path Principles'

- (a) if $s R_{z_1} t, \dots, s R_{z_k} t$, and the only variable occurring in all of z_1, \dots, z_k is x , then $s R_x t$.
- (b) if no variable occurs in all of z_1, \dots, z_k , then $s=t$.

Representing modal models with valuations in this same way is an open problem. At least, stronger Path principles are needed. But less, rather than more, is needed if we are content with a weaker structural equivalence than isomorphism. Then an unraveling argument like the earlier one for quasi-models shows the following.

Theorem 7 A finite modal model is bisimilar to a general assignment model if and only if its accessibilities are all equivalence relations.

For bisimulations with standard models, the situation is more complex. First, having a quasi-model – finite by definition – does not imply having a finite standard model. E.g., it is easy to find a quasi-model for the first-order formula $\forall xyz((Rxy \ \& \ Ryz) \rightarrow Rxz) \ \& \ \forall x \exists y Rxy \ \& \ \forall x \neg Rxx$ which only has infinite models. But in fact, having a quasi-model need not imply standard satisfiability at all. E.g., the predicate-logically inconsistent formula $\exists x \exists y Rxy \ \& \ \neg \exists y \exists x Ryx$ is clearly satisfiable in the general assignment model \mathbf{M} with domain $\{1, 2\}$, relation $\{<1, 2>\}$, and just one admissible assignment s , viz. $\{(x, 1), (y, 2)\}$. This model also satisfies the earlier Existential Generalization and Full Locality. The single type of \mathbf{M} induced by s is therefore a quasi-model for $\exists x \exists y Rxy \ \& \ \neg \exists y \exists x Rxy$. This 'inconsistency' in a set of types may seem strange – but it also shows that quasi-models are intriguing structures.