SEMANTIC PERSPECTIVES IN LOGIC

Johan van Benthem, Stanford & Tsinghua

http://staff.fnwi.uva.nl/j.vanbenthem

1 Introduction

The pleasantly erudite and still highly readable paper Beth 1963 saw logic as consisting of three basic strands, historically entangled and complementary: linguistic definability (semantics, if you will), proof, and algorithm (Beth's tag for the theme of computation). With this perspective in mind, in this light paper, I weave a story connecting some basic notions and issues on the semantic side of logic. My starting point will be invariance, but gradually, other major themes come in, such as consequence relations, preservation theorems, and the semantic role of games and agency. There will be no new technical results, my emphasis is rather on broad integrating themes in logic, hopefully along a path that the reader finds unusual, with occasionally different vistas from what one sees in standard texts.

2 Semantic invariance and definability

A common view, reflecting the standard textbook order of presentation, is that a logical language is uninterpreted syntax that stands in need of a semantic interpretation in order to create meaningful assertions. In that sense, syntax is prior to semantics. But historical reality may well have been the other way around. What came first in evolution was meaningful communication and description of reality, and human languages evolved to provide a vehicle for this. And logical languages arose out of reflection on human language,

2.1 Invariants, languages, and logics

The world-to-language perspective may be traced back to the influential view of Helmholtz 1878. Reality has structure in stable patterns, and these patterns reveal themselves to the human observer as *invariants* under suitable *transformations*. For instance, Helmholtz thought that the basic geometric notions we employ in daily life and in mathematics, reveal themselves as invariants for the natural transformations in visual perspective that correspond to ubiquitous human movements, namely those of walking in a straight line and turning around. Taken to mathematics (moving from observer movements to translations and rotations of space in geometry, or more abstract transformations in algebra), this is reflected in Klein's "Erlanger Programm": any mathematical theory needs to start from structures plus transformations setting its 'invariance level'.

The step to the emergence of *language* is then this: invariant structure is important, so a language will emerge to define and communicate information about this structure. A perceptive early exposition of this theme is in Weyl 1963, who observes that mathematical languages define invariants for associated sets of transformations over the relevant structures, while he also draws attention to what he sees as the much harder converse question whether, given those transformations, the language is rich enough to define all invariants. By now, invariance thinking has penetrated everywhere, from physics to computer science, ecology, psychology (Suppes 2002) and philosophy (Barwise & Perry 1983).

Alternatives There are also other views of how language may have arisen, more in terms of human abilities to pick up information from the world, or in terms of the basic process structure of acts of communication. We will encounter such views later on in this paper.

Invariance in logic Invariance has been implicit in logic for a long time, with invariance for isomorphisms acting as a general constraint on properly logical notions (Mostowski 1957, Lindström 1966). It became more of a conscious philosophical concern only as late as the 1980s, when Tarski 1983 and others independently started emphasizing the general power of this idea. Some logicians and philosophers even think that invariance, suitably conceived, is all there is to 'logical constants' (Sher 1991). It is not my aim here to tell the story of this foundational debate. Van Benthem 2001 has extensive discussion of strengths and weaknesses of invariance as a criterion for logicality, a wide array of relevant literature references, and a brief survey of alternative approaches to logicality. Instead, I start by pointing out some technical features of the connection between logical languages and invariance, broadening the canvas to include other themes as we go along.

Permutation invariance Much of the philosophical and linguistic literature chooses for its transformations *permutations* π of some fixed domain D of objects. For instance, the identity relation $x = y$ is invariant for such transformations, since it holds iff $\pi(x) = \pi(y)$. Together with non-identity, the universal relation and the empty relation, we get the only four permutation-invariant relations between objects. In other linguistic categories, say, the quantifier "some" is permutation-invariant, since $some(A)$ holds for some set of objects A (that is, A is non-empty) iff $some(\pi[A])$ holds for any permutation image $\pi[A]$ of the set A. Essentially, the only thing that matters here is the cardinality of the set A, and this is also true for many other quantifiers, such as "all", "one", or "most": after all, these are expression of 'quantity' only. We will return to the behavior of quantifiers in a moment.

Permutation invariants occur in many different kinds of expression in natural language, where we can match linguistic categories with a simple type system in Montague's style, with base domains of objects and truth values, and operations of product and implication. Laüchli 1970 proved the still intriguing result that the types constructed in this way whose domains always contain permutation-invariant objects are exactly those that correspond, when read as formulas in a propositional language with conjunction and implication, to validities of intuitionistic logic. For a systematic discussion of permutation invariance in finite type theory, we refer to van Benthem 1989, a contribution to a seminal issue of the *Notre Dame Journal of Formal Logic* that collected many approaches to the nature of logical constants co-existing at the time: semantic, but also proof-theoretic. [1]

Isomorphisms and automorphisms A more general view, in line with what we said earlier, is one of notions being invariant for relations between different structures. Permutation invariance is a special case of invariance for *bijections*. And one step further, bijections are an extreme case of *isomorphisms* between structures ('automorphisms', when we are inside the same structure), being bijections that preserve all relevant atomic predicates and operations in a natural sense. [2] This is how we mostly will phrase things henceforth.

[1] Further results, including connections to notions of invariance appropriate to the lambda calculus due to Statman 1982 and Plotkin 1980, can be found in van Benthem 1991.

[2] Bijections abstract away from all content, and only preserve identity and non-identity of objects.

We now continue with the general theme of isomorphism invariance in logic. For a start, the basic system of first-order logic satisfies the following invariance property:

Fact If F is an isomorphism from a model $\boldsymbol{M}$ to model $\boldsymbol{N}$, then, for any first-order formula φ and assignment of objects $\boldsymbol{d}$ to the free variables $\boldsymbol{x}$ of φ, $\boldsymbol{M}, d \models \varphi$ iff $\boldsymbol{N}, F(\boldsymbol{d}) \models \varphi$.

Many basic results about definability in first-order logic involve invariances of this sort. For instance, Beth's Definability Theorem can be viewed in this light – but for present purposes, we rather state Svenonius' Theorem, in a version proved in van Benthem 1982. The following result gives a precise sense, in first-order logic, to a pervasive invariance-related intuition that being able to fix denotations up to isomorphism leads to definability.

Theorem A predicate P is fixed up to isomorphism by theory $T(P, \boldsymbol{Q})$, meaning that, in all models for $T(P, \boldsymbol{Q})$, any $\boldsymbol{Q}$-automorphism is automatically a P-automorphism, iff P is explicitly definable in T up to finite disjunction: $\vee_i \forall \boldsymbol{x}(P\boldsymbol{x} \leftrightarrow \delta_i(\boldsymbol{Q}, \boldsymbol{x}))$ is a semantic consequence of T for some finite set of formulas $\delta_i(\boldsymbol{Q}, \boldsymbol{x})$.

We have concentrated on first-order logic here, since it is still the hothouse for developing logical ideas, but isomorphism invariance is so widespread in logic that it is a defining characteristic of logical systems in Abstract Model Theory (Barwise & Feferman, eds., 1985).

2.2 Combining invariance and inference

Isomorphism invariance is a somewhat abstract criterion, but it often acquires more bite when combined with other logical phenomena. As a simple example in a linguistic setting, let us go back to bijection invariance, considering binary quantifiers with their usual pattern of occurrence in natural language, the type

$Q\ A\ B$ standing for binary relations between sets of objects.

This realm was studied in van Benthem 1984, restricting attention to finite domains of objects. On any given domain of size n, isomorphism invariance tells us that the denotation of the quantifier is fixed by the set of 4-tuples (n_1, n_2, n_3, n_4) with $n_1 + n_2 + n_3 + n_4 = n$, where the four numbers refer to the sizes of the zones $A \cap B$, $A \cap -B$, $B \cap -A$, $-B \cap -A$. [3] Now we combine this with one more plausible restriction on quantifiers, satisfied by all the usual quantifiers studied in logic and linguistics, making the A-domain paramount:

Conservativity $Q\ A\ B$ iff $Q\ A\ (B \cap A)$

This leads to a view of quantifiers as sets of pairs of numbers (adding up to the size of A), visualized geometrically as a set of points $(|A \cap B|, |A-B|)$ in the so-called 'Tree of Numbers'

$$
\begin{array}{ccccc}
 & & (0, 0) & & \\
 & (1, 0) & & (0, 1) & \\
(2, 0) & & (1, 1) & & (0, 2) \\
 & & \ldots & &
\end{array}
$$

Now conservativity may be viewed as a basic Boolean property of quantifiers, and what it shows is how quantifiers display basic inferential behavior from the start. More generally, most quantifiers in a daily use satisfy Boolean inference properties such as the following:

[3] This reduces the size of the set of all generalized quantifiers on n objects by an exponential.

Monotonicity $Q\,AB,\ B \subseteq B'$ *imply* $Q\,AB'$

This property may be called ‘upward right’ in a natural sense. Thus, there are four basic monotonicity properties in all: right or left, upward or downward. Typical examples for each occur in the Square of Opposition: “all”, “some”, “no”, “not all”.

In the Tree of Numbers, monotonicity properties acquire a direct geometrical meaning. For instance, upward right monotonicity means that on horizontal rows in the tree, once a position is accepted, so is everything to its right, while upward left monotonicity says that once a point is accepted, so is the whole downward subtree generated from it.

Now we can classify all possible quantifiers that are permutation-invariant, conservative, and that support a rich set of inferences. Here is one sample result:

Theorem All doubly monotone isomorphism invariant quantifiers are first-order definable.

These first-order quantifiers exhibit a geometrical pattern of finite unions of convex sets, as can be seen from the above facts, which typically differs from the tree pattern for “most AB” $(|A \cap B| > |A - B|)$ whose boundary follows a zigzag line through the middle of the tree.

The original motivation for results like this in the 1980s was an investigation to which extent natural language manages to define all ‘natural’ and ‘useful’ counting expressions. Part of this utility is richness in available inferences, where monotonicity represents a sort of stability: the quantifier still holds when we encounter changes in one or both of its set arguments. This highlights the widespread entanglement of invariance with inference, an important theme to which we shall return several times in what follows.

2.3 Potential isomorphism

However, isomorphism is not the only game in town, and there are alternatives, even in the heartland of logic. First-order logic is also invariant for a much less-demanding invariance called *potential isomorphism*. This is a family $\boldsymbol{F}$ of finite partial isomorphisms F between two models satisfying the Back and Forth properties, saying that given any object a in one of the models, there is an object b in the other model such that the extended function $F \cup \{(a, b)\}$ also belong to the family $\boldsymbol{F}$. Isomorphisms induce potential isomorphisms by taking all their finite submaps, but the converse does not hold. To show the attraction of this notion we cite one result from van Benthem & Bonnay 2008. These authors identify the abstract form of the Back and Forth properties as a natural diagrammatic form of ‘commutation of an invariance relation with object expansion’.

Theorem Potential isomorphism is the smallest similarity relation between models that commutes with object expansions.

They then prove a very general result that, for any binary relation E and equivalence relation S over some class of objects, S commutes with E iff the inverse E^{-1} preserves S-invariance. As a consequence, potential isomorphism is the smallest similarity relation S between models that respects truth values of atoms while object projection is S-invariant. [4]

However, the fit is still not precise: invariance for potential isomorphism also holds for several strong extensions of the first-order language. It only becomes a precise match

[4] This abstract analysis also applies to modal logic and bisimulation, a topic to be introduced below.

when we move to *infinitary first-order logic* $L_{\infty\omega}$, which allows conjunctions and disjunctions over arbitrary sets of formulas. Then we can show the following equivalence: [5]

Theorem Two models have a potential isomorphism between them
iff they satisfy the same sentences of $L_{\infty\omega}$.

For more on the conceptual importance of $L_{\infty\omega}$ and potential isomorphism, cf. the earlier-mentioned Barwise & Feferman, eds., 1985, van Benthem & Bonnay 2008.

2.4 Fixed-point logic and computation

Another crucial extension of first-order logic, incomparable in that it can define well-foundedness of binary orders (which $L_{\infty\omega}$ cannot), is *first-order fixed-point logic LFP(FO)* (Ebbinghaus & Flum 2005) having operators for defining smallest fixed points for formulas

$$\mu P, \mathbf{x}\bullet \varphi(P, \mathbf{Q}, \mathbf{x}) \qquad \text{where } P \text{ occurs only positively in } \varphi,$$

– and likewise, there are definitions $\nu P, \mathbf{x}\bullet \varphi(P, \mathbf{Q}, \mathbf{x})$ of greatest fixed-points. This system is a natural extension of first-order logic, with still countable syntax, that can encode the fundamental theory of induction and recursion, the staples of computation. [6] Thus, it represents a major further motive across logic that ties in with our earlier concerns.

One might zoom in even more closely on fixed-point logic by tightening invariance for potential isomorphism. But one can also see this as a point where other notions come into their own. Induction and recursion are basic structures of computation, the third ingredient in Beth's view of logic, and fixed-point formulas may be seen as recipes for computational algorithms. Perhaps, then, what is characteristic of systems like *LFP(FO)* is the interplay of two factors: semantic invariance, and algorithmic structure. [7]

2.5 Discussion

More permissive views of logicality In the above setting, we thought of the logical constants as invariants for bijections: rough isomorphisms respecting only identity of objects. However, invariance for *L*-isomorphisms suggests a more liberal view that is not 'all or nothing'. We allow the atomic predicates in the language *L* as parameters (invariant by definition), and we ask which complex predicates are then also invariant, or perhaps more pointedly, which constructions (sometimes called 'logical glue') maintain invariance. Now even in the realm of logical expressions, some parameters in *L* may have a special status.

Consider 'mass quantifiers' in natural language, such as "all the wine", "some wine", "most wine". Now there is no discrete base domain, we are rather in a mereological setting of continuous objects (say, bits of water) ordered by inclusion. In this case, the definition of the quantifiers may refer essentially to this inclusion structure, not just to identity of objects. Thus, mass quantifiers fail the above test of invariance under bijections. But they

---

[5] The proof uses some basic model-theoretic definability techniques such as 'Scott sentences' describing types of objects occurring in models up to any given ordinal depth of recursion.

[6] Unlike Recursion Theory, fixed-point logic makes no assumptions about the data one computes over, disentangling the mix of recursion and coding that is characteristic of the latter field.

[7] The lack of a purely semantic fit also shows in the lack of an abstract model-theoretic Lindström Theorem characterizing *LFP(FO)*, cf. Van Benthem, ten Cate & Väänänen 2009.

are still natural, they support Boolean inferences just like the standard ‘count quantifiers’ (Peters & Westerståhl 2006), and they are invariant under inclusion isomorphisms. In other words, the appropriate invariance level for even logical expressions may differ. [8]

Whence the base structures? The case of mass versus count quantification also shows how invariance analysis depends on a prior choice of structures to work on, a choice that involves other considerations. In particular, working in standard mode in a set-theoretic universe suggests an underlying base domain of primitive objects out of which everything is constructed ‘upward’ by set-forming operations. Mereology, on the other hand, suggests a very different view: a universe where we can only analyze things ‘downward’ into smaller components, without a guarantee that we will hit smallest objects. Building up versus analyzing represent very different views of what logical semantics is about, and the choice between these two perspectives is not made for us by invariance thinking.

Still, an irreducible role for one’s conceptual choice of base objects and their patterns is not a problem for a story of the sort we are telling: it rather makes our discussion all the more semantic in the general sense of that word in logic.

Function words in natural language The general picture emerging here fits the functioning of natural language. In addition to logical expressions like Booleans and quantifiers, there are also many ‘function words’ with a somewhat logic-like behavior, such as modals (“can”, “may”, “must”), prepositions (“in”, “out””, “of”, “with”, “to”), or other abstract functional items such as comparatives or indexicals. All these linguistic expressions come with their own inferential behavior, and their own notions of invariance referring to, often ubiquitous, semantic structure that is appropriate to them. Thus, from an invariance perspective, logicality is a widespread phenomenon that can come in degrees, or levels.

2.6 Conclusion

Invariance under structure transformations is a pervasive aspect of semantical analysis for logical languages, and it supports sophisticated notions and technical results in the field. Also, it naturally leads us to acknowledge a wide variety of ‘logicality’ across natural language, depending on the invariance level, making the standard logical constants less isolated. Finally, and quite significantly, invariance is naturally entangled with other basic logical notions. One of these, as we saw in the pilot case of monotonicity, is inference and proof. A second entangled notion is computation, the third strand in Beth’s view of logic, which emerged when we looked at fixed-point logics for induction and recursion.

Syntax What is interesting to the latter points is this. Both proof and computation depend crucially on syntax, as the code we are working with. Thus, while invariance might suggest a priority for structure over syntax, once the language is there, other logical intuitions become naturally available, and can operate freely, sometimes even without immediate semantic counterparts. This entanglement of themes will return in the sections to follow.

---

[8] A similar point holds for logical operators in intuitionistic, rather than classical logic on models of information stages ordered by inclusion. Permutation invariance is not the right notion there, invariance for inclusion-isomorphisms is (or even stronger criteria to be discussed below).

3 Plurality of zoom levels

3.1 Plurality and zoom

Plurality While invariance has been proposed as a unique underpinning for a core of logic, scientific practice seems very different. When analyzing reality, many different levels of structure can play a role. Consider mathematics: it has a wide array of legitimate theories of space, ranging from affine and metric geometry to topology. These different theories represent different 'zoom levels' for looking at reality. Each of these levels comes with its own logical language: richer if the structural similarity relation is more detailed (making invariance under this relation less demanding), poorer if the similarity relation is coarser.

Looking down, or up This observation fits with a fact about logical analysis. Many people see the task of logic as providing ever more detail, formalizing each small step in reasoning and each feature of structure. The paradigm for this would be total formalization of informal mathematics into machine-checkable languages, with nothing taken for granted. Steps of reasoning at this most detailed ground level can be considered a true base logic.

But proofs in this style may be as unreadable as machine code, and an equally legitimate form of logical analysis does the precise opposite. One looks at a reasoning practice, abstracts away from details, and looks for global patterns representing some high level of reasoning that may even bring to light patterns undiscovered so far. For a sample of the contrast, compare the extremely detailed first-order language of Tarski's 'elementary geometry' with the highly general modal logic of the interior operation in topology (van Benthem & Bezhanishvili 2008). We will discuss two ways in which this plurality may arise, and point at some open problems that emerge when we look at this generally.

3.2 Logics for graphs

Let us fix one particular similarity type, annotated directed graphs *(W, R, **P**)* with a set of points *W,* a binary relation *R*, and a bunch of unary predicates ***P*** of points. If we choose isomorphism as our invariance, then languages appropriate to studying graphs are first-order logic and its extensions – or even second-order logics – and these systems are indeed a good fit with much reasoning found in Graph Theory. However, there are also other natural ways of looking at graphs, where we identify more structure.

One such invariance is *bisimulation*, where we are only interested in local properties of points as well as the structure of accessibility at each point, i.e., its arrows to other points. [9] A bisimulation *E* between two annotated graphs ***M***, ***N*** is a binary relation between points in ***M*** and ***N*** matching only points with the same local properties, and satisfying the following Back and Forth properties: (a) if *s E t*, and *s R s'* in ***M***, then there exists a point *t'* with *t R t'* in ***N*** such that *s' E t'*, (b) the same clause in the direction from ***N*** to ***M***.

It is easy to find non-isomorphic bisimilar graphs, and indeed, bisimulation is a rougher invariance level, proposed independently in modal logic, set theory, and computer science.

---

[9] Notice how points in graphs derive their identity from two sources: their local properties, plus: their connections with other points. This recursive nature, which can be made precise in terms of our earlier fixed-point logics, provides a much more sophisticated view of these structures that what might be suggested by standard discussions of 'possible worlds' as isolated entities.

Finding natural structural similarity relations is often a matter of having some intuition in mind – and for bisimulation, a powerful intuition is one of equivalence of processes, where we see points in graphs as states of a process, and arrows as possible state transitions.

The modal landscape The language that fits bisimulation is *modal logic*, and a rich theory has sprung up at this level that we cannot survey here: cf. Blackburn, de Rijke & Venema 2002, van Benthem 2010 on the mathematics of modal logic. Our main point is just that modal logic is a good illustration of how invariance thinking ties up with language design.

For a start, there is a *translation* from the modal language into first-order logic, sending universal modalities $\Box p$ to guarded quantifier formulas $\forall y\ (Rxy \to Py)$, and existential $\Diamond p$ to $\exists y\ (Rxy \wedge Py)$. So, we can view the logic of a coarser level as a fragment of the logic of a richer level in a precise manner. Moreover, there is a certain dynamics of design. By now, there is a whole landscape of languages in between modal logic and first-order logic, that can be viewed as arising in two ways. Either we extend the syntax of modal logic by certain 'hybrid' expressive devices available in first-order logic, or we devise new notions of simulation in between bisimulation and isomorphism, and create matching languages.

Zoom levels This setting also explains why working at different zoom levels can be useful. The first-order language, though richer than modal logic, and natural in its own way, comes at a cost. Its theory is undecidable by Church's Theorem – and at a more domestic level, its syntax is more complicated. Modal logic has a variable-free notation that makes checking for truth provably easier, and allows for perspicuous notation of basic proof patterns without variable management. This demonstrates a much more general point that is often under-appreciated. In logic, as in acting, saying less is sometimes saying more.

Zooming out: from actions to powers Invariance thinking also suggests coarser levels with weaker logics than the basic modal one. Here is an illustration about games, viewed as processes of interactive computation (van Benthem 2014 has details). Consider a finite extensive game tree with transitions as moves, while local properties of nodes record turns for players and pay-offs at end nodes. Now we may want to abstract away from local moves, as often done in Game Theory, being interested only in players' *powers* for controlling the outcomes of the game. Then we can identify game trees where players have the same powers using a natural notion of 'power bisimulation'. For instance, in the games

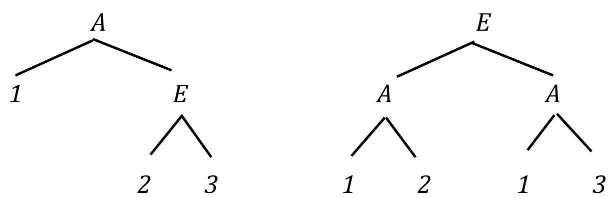

player E has the same powers $\{1, 2\}$, $\{1, 3\}$, though her local moves are quite different.

Naturally, there exists a language matching this invariance, less expressive than the modal logic of game trees. To be aligned with the power structure, it has 'forcing modalities' $\{i\}\varphi$ saying that player i has a strategy for playing the game such that, with any counterplay by the other players, only end nodes result that satisfy φ. The logical theory of forcing has many similarities with modal logic – be it that $\{i\}\varphi$ only has upward monotonicity $\Box\varphi \to \Box(\varphi \vee \psi)$ as a base law (plus laws linking powers of different players). The aggregation law $(\Box\varphi \wedge \Box\psi) \to \Box(\varphi \wedge \psi)$ of basic modal logic is invalid for powers, as is easy to see.

This reflects the fact that powers suggest a transition from graph models to 'neighborhood models' for modal logic, relating points to sets of end points reachable from them. [10]

The search for natural invariances and logical languages that fit games is still ongoing. [11] We have mentioned this illustration to show how invariance thinking is a live topic today, while also, games will make a brief appearance in this paper later on.

3.3 Interlevel connections

Translation While we are advocating variety of simulations and matching languages, logic abhors chaos. To keep things in hand, one must study *connections* between levels. We already mentioned the existence of translations between languages, and modal logic provides examples of how these can be related to invariances. For instance, the Modal Invariance Theorem (van Benthem 1977) says the following:

> Any first-order formula φ in the signature $(R, \mathbf{P})$ is equivalent to the translation of a modal formula iff φ is invariant for bisimulations.

Many such characterization theorems exist for modal logics in a broad sense, but there are also other ways in which we can relate different levels and their languages, for instance, by mapping similarity relations directly. [12] A good setting for this is Category Theory, and much work has been done for modal languages in co-algebra (Venema 2012).

Digression: weaker or stronger? We have suggested that invariance themes are rewarding, we did not offer a final theory. For instance, we pointed out how logical languages at different invariance levels can be related. But we have not offered a definitive view on precisely how. One of the vexing (but also intriguing and wonderful) things about logic is that one can look at the same situation in different ways in tandem. For instance, is the modal language really more specialized and weaker than that of first-order logic, or is it more general? The latter view is developed in generalized modal semantics for first-order logic (Andréka, van Benthem & Németi 1998), and the strong analogies between modal logic and first-order logic are analyzed further at an abstract level, via a correspondence between potential isomorphism and bisimulation, in van Benthem & Bonnay 2008. [13]

3.4 Neighborhood structures

Generalizing a similarity type Now consider an issue discussed earlier. Searching for new invariances may also make us change the similarity type of the models we work with. For example, let's generalize annotated graphs to *neighborhood models (W, N, P)* with $N\,s\,X$ a relation between points s in W and sets of points X (cf. the textbook Pacuit 2017). [14]

---

[10] We do not elaborate this (but cf. van Benthem 2014) as it may confuse the reader with neighborhood models as a richer, rather than a poorer, level beyond modal logic to be introduced below.
[11] For a new 'instantial bisimulation', in between modal and power bisimulation, that fits with game-theoretic equilibria involving all players, see van Benthem, Bezhanishvili & Enqvist 2016.
[12] To be added: direct match of modal bisimulation and power bisimulation. Maybe also: display the translation from the poorer power modalities into the richer standard modal logic or PDL.
[13] Generalized semantics (cf. Andréka, van Benthem, Bezhanishvili & Németi 2014) is related to our themes here. But we still do not see exactly how things fit, and therefore omit further discussion.
[14] This move is natural, since neighborhood models are parametrized 'hypergraphs'.

Neighborhoods, evidence and plausibility For a concrete case, consider the treatment in van Benthem & Pacuit 2011 of agents' evidence for their beliefs. Starting at the ordered graph level, belief is a modality using a reflexive transitive plausibility ordering $\leq$ of worlds (points in a graph), with $B\varphi$ saying that φ is true in all the most plausible worlds – while conditional belief $B^{\psi}\varphi$ refers to truth in the most plausible ψ–worlds. Indeed, plausibility models support a standard modal language with even further doxastic notions.

Now, what if we add semantic structure explaining how the plausibility order came about? We can do this by giving models a family $\boldsymbol{E}$ of subsets encoding the evidence received at the current stage. And we stipulate that $s \leq t$ iff for each set E in $\boldsymbol{E}$, if $s \in E$, then $t \in E$. Thus, each evidence model $\boldsymbol{M}$ induces a plausibility graph *ord(**M**).* A natural language matching this richer level of structure contains not just analogues for the modalities at the derived plausibility level, but also new operators, such as $\Box\varphi$ saying there exists an evidence set 'supporting' φ in the sense that each point $s \in E$ makes φ true. In fact, several languages make sense here, depending on which similarity relation, less or more strict, one imposes on evidence structures (cf. van Benthem, Bezhanishvili, Enqvist & Yu 2016). Now, comparison between levels is not as direct as before: we have not just different invariances on the same structures, but maps transforming structures at one level into those at another.

We will not pursue details of this ongoing work, but summarize the point we want to make. Generalization of similarity types is another active force in semantics, going beyond invariance inside one type. Having said this, type change will be a side theme in this paper.

3.5 Conclusion

Variety of zoom levels in semantics is a natural ongoing process. We gave illustrations from the realm of modal logic, driven by interests in spatial logics, process theories, or varieties of information. We showed that this variety is not just a free-for-all. There are connections between different zoom levels and their invariances – though open problems abound, and maintaining the unity of logic lies in getting clear on these.

4 Preservation, generalized consequence, and model change

4.1 Standard consequence

Logical semantics, with its emphasis on meaning and truth also offers an account of valid reasoning. The standard semantic notion of consequence says that $\Sigma \models \varphi$ iff for all models making all formulas in Σ true, φ is also true. Thus, the notion of consequence gets reduced to that of truth in a model, which can then be analyzed in its own recursive manner. This is not a arbitrary choice. The usual completeness theorems for logical systems then say that there is an extensional equivalence (qua admissible transitions from sets of formulas to formulas) with syntactic notions of derivability defined in some appropriate manner, perhaps even according to independent proof-theoretic intuitions.

4.2 Interpolation and invariance

Consequence mixes well with our earlier topic of invariance. We give one illustration. Consider Craig's Interpolation Theorem for first-order logic.

Theorem For all first-order formulas φ, ψ with $\varphi \models \psi$, there exists a formula α whose non-logical vocabulary is contained in that both φ and ψ such that $\varphi \models \alpha \models \psi$.

Now, intuitively, there is a semantic surplus to the existence of an interpolant involving only part of the vocabulary of the antecedent and the consequent of an inference. What the existence of an α as above guarantees is the following 'transfer property':

Let $\boldsymbol{M} \models \varphi$ and let there be an $L_\varphi \cap L_\psi$–potential isomorphism between $\boldsymbol{M}$ and any other model $\boldsymbol{N}$: then $\boldsymbol{N} \models \psi$.

Let us say that, in this case, 'φ entails ψ along $L_\varphi \cap L_\psi$–potential isomorphism'. The following analysis comes from Barwise & van Benthem 1999. Here is a new version of the first-order Interpolation Theorem highlighting this special behavior:

Theorem The following are equivalent for all first-order formulas φ, ψ: (a) there is an L-interpolant for φ, ψ, (b) 'φ entails ψ along L–potential isomorphism'. [15]

Meta-logic Here is an interesting fact. Standard Interpolation fails for natural extensions of first-order logic: in particular, the earlier infinitary logic $L_{\infty\omega}$. However, the preceding invariance version of Interpolation can be shown to hold for $L_{\infty\omega}$. Thus, invariance versions of well-known meta-properties may have better prospects of extending to other logics. What becomes clear at the same time is this. What a meta-property of a logical system is may depend on its formulation, and hence received views of their 'holding' or 'failing' across systems in the landscape of logics may need reconsideration. [16]

4.3 Preservation and generalized consequence

Preservation theorems The same thinking extends to model-theoretic preservation theorems, to some the most attractive results that started Model Theory in the 1950s. A key example is the Los–Tarski Theorem saying that a first-order formula is preserved under submodels iff it is definable in a purely universal syntax starting from literals (atoms and their negations) using only conjunction, disjunction, and universal quantifiers. Another such result (relevant to our earlier discussion of monotonicity inference) is Lyndon's Theorem: a first-order formula is upward monotonic in the predicate P iff it is equivalent to a formula in which P has only positive syntactic occurrences. These results embody the essence of what logicians like: syntactic form determines semantic behavior. [17] [18]

Entailment along a relation From our current perspective, preservation theorems are really about a generalized notion of consequence, allowing for transfer in that, when the premises hold in one model, the conclusion holds in some other model. Standard consequence is the special case where we stay, a bit timidly, inside the same model.

---

[15] One can also rework most standard preservation theorems in this interpolation style.

[16] This reanalysis also carries over to modal logic. For instance, a valid consequence between two modal formulas φ, ψ has a modal L-interpolant iff φ entails ψ along L-bisimulation.

[17] In line with our earlier discussion of invariance, we could also reverse this, and say that useful semantic transfer behavior across models will lead to the emergence of matching special syntax.

[18] We forego other features, such as the fact that special syntax, say, universal or positive, may make inferences perspicuous, and sometimes even recursively computable from mere syntax (cf. van Benthem 1986), in ways that arbitrary first-order formulas do not.

A general way of reasoning about entailment along any relation R is in a modal format

$$\varphi \to [R]\psi$$

The general properties of this style of reasoning come out in a calculus having different relations. For instance, with *;* as relational composition, and ^{v} as converse, we have:

(a) $\varphi \to [R]\psi$ and $\psi \to [S]\alpha$ imply $\varphi \to [R; S]\alpha$,

(b) $\varphi \to [R]\psi$ implies $\neg\psi \to [R^v]\neg\varphi$

Van Benthem 1996, 1998 give complete logics for the 'universal Horn fragment' of propositional dynamic logic PDL, which describes the structural properties of such generalized reasoning in their most abstract form. We submit that real inference usually involves jumps across situations, a theme already emphasized in Barwise & Perry 1983.

4.4 Excursion: 'alternative logics'?

By making the relational jumps in transfer inferences explicit as modalities, our implication stayed classical, supporting all the usual properties of classical inclusion, except for Reflexivity (since $\varphi \to [R]\varphi$ clearly fails in general). Still, our setting supports variations that break laws of classical consequence. For instance, new consequence notions arise if we endow models with an order of relevance or importance. If we then say that a conclusion holds if it is true in all most important models for the premises (so premises can influence which of their models are relevant), we get non-classical *non-monotonic logics* of a sort widely studied in the 1980s. But again, classical logic lies close by. There is a clear analogy between this way of thinking and the earlier plausibility models for belief. Instead of insisting on non-monotonic consequence $\varphi \Rightarrow \psi$, we might just as well add a formula $B^{\varphi}\psi$ with an explicit modality for belief to classical logic (van Benthem 2011).

4.5 Dynamic logics of model change

Transfer across models occurs more often in the recent literature. In particular, current logics of information change analyze how formulas expressing knowledge, belief, or yet other attitudes of agents change as new information comes in. We cannot go into the motivations or the further structure of such 'dynamic epistemic logics' here (see Baltag & Smets 2006, van Benthem, van Eijck & Kooi 2006, van Ditmarsch, van der Hoek & Kooi 2007, van Benthem 2011). But of relevance to us is this key technical feature: update with new information is treated as definable model change.

In particular, an event $!\varphi$ of getting the 'hard information' that proposition φ is the case restricts a current model $\boldsymbol{M}$ to a definable submodel $\boldsymbol{M}/\varphi$. This is a semantic 'relativization' of a model to a definable subdomain. If we now try to axiomatize the corresponding logic, with explicit dynamic modalities $[!\varphi]$ for informational actions, and static modalities K and B for agents' knowledge and beliefs, we are no longer interested in transfer implications, but in equivalences telling us how facts in the updated model relate to facts in the original one. A good example of such an equivalence are these two valid laws:

$$[!\varphi]K\psi \leftrightarrow (\varphi \to [!\varphi]\psi)$$

$$[!\varphi]B\psi \leftrightarrow (\varphi \to B^{\varphi}[!\varphi]\psi)$$

that give basic properties of semantic relativization. More general dynamic logics of update also axiomatize the behavior of operations that modify plausibility order (say, $\Uparrow\varphi$ puts all φ–points on top of all $\neg\varphi$–points in the current model, while retaining the old order inside these zones), or transform models in yet other definable ways. More complex counterparts to the above laws generate complete logics of such model transformations.

Technically, one can see these logical systems as formalizing the theory of cross-model relations matching definable transformations of models. This is more special than the scenarios that we considered before, where entailment could be along any relation, say, that of being an arbitrary submodel, definable or not. For definable transformations, often the static base logic has enough expressive power to supply all needed recursion laws of the above sort (cf. van Benthem & Ikegami 2008 on 'product closure' of logics). Indeed, dynamic-epistemic logics tend to be decidable if their static base logic is, whereas logics of consequence along arbitrary relations can be much more complex. [19]

Still, this simplicity of logics for definable model change is fragile. Löding & Rohde 2003, a study of 'sabotage games' where one player can delete arrows from a graph, while the other tries to reach some goal point (cf. van Benthem 2014), showed how the modal logic of removing arbitrary links from relations is undecidable. Aucher, van Benthem & Grossi 2016 show how even simpler 'stepwise' versions of dynamic-epistemic logics, where just some point lacking the property φ (or some link failing to pass the relevant update recipe) gets removed, may well become undecidable – even when the base logic is quite simple.

4.6 Zoom levels and tracking

Our general point here is simple, but sweeping. The universe of models for semantics is criss-crossed by links of various sorts. In inference, we are often interested in what can be said about one model in terms of what we know about another model linked to it. And to put this in its proper perspective, logics of model change and transfer seem an appropriate medium. Moreover, there are strong connections here with our earlier semantic themes. In particular, the update operations introduced in this section should respect whatever invariance relation we have chosen for our static models. But there is more. Updates can take place at various zoom levels, and in that case, there is a significant issue of when updates at a coarser level faithfully 'track' updates at a finer level. For more on this connection between model change and zoom levels, cf. van Benthem 2016, Cina 2017.

4.7 Conclusion

We have shown how a concern with consequence is a natural companion to our earlier semantic considerations of invariance and definability. We also showed how generalized notions of consequence make sense then, that merge with model-theoretic preservation results and with model transformations as studied in dynamic-epistemic logics. [20]

5 Games and agents

While our topics so far were well within the realm of standard model theory as description of the world, in this final section, we explore a slightly more radical perspective.

---

[19] To be added: Why FOL with a model extension modality is not RE, open problem of complexity.

[20] All this is not to deny that there are also independent more syntactic proof-theoretic intuitions concerning consequence – but these are not the topic of this paper.

5.1 Logic games

Model comparison games Our starting point is the fine-structure of similarity relations and invariance. A well-known technique for analyzing this are 'model comparison games' played between two given models $\boldsymbol{M}$ and $\boldsymbol{N}$, due to Ehrenfeucht and Fraïssé. Such games, and in particular, their strategies, form a natural bridge between logical syntax and invariant semantic structure. The game works as follows. A player $\boldsymbol{S}$ ('Spoiler') who claims dissimilarity of the models chooses one of the models, and picks an object d in its domain. The counter-player $\boldsymbol{D}$ ('Duplicator') then chooses an object e in the other model, and the pair (d, e) is added to the current list of matched objects. After k rounds, the object matching is inspected. If it is a partial isomorphism, $\boldsymbol{D}$ wins; otherwise, $\boldsymbol{S}$ does. For concrete examples of this game and its uses, we refer to Doets 1996, van Benthem 2014.

It can be shown that this game is adequate in the following sense.

Theorem For all models $\boldsymbol{M}, \boldsymbol{N}$, and $k \in N$, the following two assertions are equivalent:

(a) $\boldsymbol{D}$ has a winning strategy in the k-round game,

(b) $\boldsymbol{M}, \boldsymbol{N}$ agree on all first-order sentences up to quantifier depth k.

In fact, the correlation with syntax is much tighter than it might seem from this version.

Theorem There exists an explicit correspondence between

(a) winning strategies for $\boldsymbol{S}$ in the k-round comparison game for $\boldsymbol{M}, \boldsymbol{N}$,

(b) first-order sentences φ of quantifier depth k with $\boldsymbol{M} \models \varphi$, *not* $\boldsymbol{N} \models \varphi$.

Likewise, $\boldsymbol{D}$'s winning strategies can be made more explicit in terms of 'towers' of partial isomorphisms. Perhaps most satisfyingly, in the game over infinitely many rounds, the winning strategies for $\boldsymbol{D}$ are the potential isomorphisms between the models (if any).

Other logic games The same sort of strategic analysis provides fine-structure to other logical notions such as truth, proof, or model construction (Väänänen 2011). All these games show interesting connections. For instance, the above match of first-order formulas and strategies for the Spoiler $\boldsymbol{S}$ can be reworked into a precise correspondence of modal comparison games and evaluation games for first-order formulas in single models.

5.2 Introducing agents

But to us, the use of games in logic signals something of much greater import. Games are played by *agents*, and what really comes to the fore here is the role of agents dealing with truth, similarity, or consequence. [21] A focus on agents emphasizes what might be called 'the other face of logic': not as description of the world, in either physical or mathematical structure, but as a form of *structured activity*, highlighted in communication, argumentation, and other pervasive themes in the history of logic – often having to do with the information that agents have, and with the interaction of multiple agents.

Games are a concrete focus for studying agency, and what makes them even more attractive is that they are also a model for the modern face of *computation* (one of our basic themes) as interactive agency between computers (and humans) in large social networks. Indeed, van Benthem 2014 develops various links between game solution procedures, game-theoretic equilibria, and the fixed-point logics that we discussed earlier.

[21] This agency theme also underlies the dynamic-epistemic logics in the preceding section.

Logical constants once more This is not just a tool, a game-theoretic stance can affect our very understanding of logic itself. From a game-theoretic perspective, logical constants are not so much most general invariants of reality, the line of our earlier discussion, but rather the most general structures one can find in creating games. Conjunction and disjunction reflect choices different players can make in games, negation correlates with role switch, and quantifiers involve sequential composition of games. Even further logical constants, beyond the classical repertoire, arise when we consider further natural operations on games such as various forms of parallel composition, or infinite iteration. We merely mention this more radical switch here, but the consequences of such a more radical shift for semantics as traditionally understood remain to be fully investigated.

5.3 Agent diversity

Our last illustration of the agency perspective concerns what it might do to topics that we discussed earlier, and that seemed settled. For instance, consider the notion of similarity between models and invariance. We might think that whether two models are similar is not an objective given, but something that depends on the view of the agents doing the comparison. But then what becomes relevant are two new features.

Bounded agents and automata One issue is what idealized capacities we assign to agents in performing the tasks demanded by standard logical notions. Clearly, real agents operate under tight bounds on what they can infer, observe, or remember. Some awareness of this theme exists in computational logic (and some parts of game theory), where agents are modeled by *automata* of various sorts. For instance, in 'pebble versions' of model comparison games, the agents have finite memories given by a fixed number of pebbles at their disposal that they can use to mark objects when drawing samples from the models and making comparisons as to partial isomorphism. This leads to even more fine structure than we displayed above, in terms of winning strategies matching the syntax of finite-variable fragments of first-order logic. Other uses of automata abound (cf. Graedel, Thomas & Wilke, eds., 2002), but these ideas have not percolated yet into core logic.

From fragments to agents What might agents do in the heartland of logic? Normally, we think in terms of logical systems, i.e., complete machineries for definition and inference. Or, if we think of users, we do so implicitly, considering low-complexity fragments of complex systems, presumably better usable in practice in terms of complexity of model-checking or inference. Well-known examples of decidable fragments inside the undecidable system of first-order logic are monadic predicate logic, or modal logic.

But here is an alternative perspective. If a bounded agent is handed a complex logical system, that agent will only be able to use part of that system, or at least, she will only use part of the full system correctly. But then we can rethink what are usually seen as 'fragments' of complete logical systems in terms of what is available to agents, modeled, say, as automata with various computational limitations. One instance of this type are the earlier pebble games: the usual adequacy results in terms of finite-variable fragments determine for which parts of complete predicate logic such agents can perform model comparisons correctly. A more proof-theoretic example would be simple agents (say, finite automata) doing a parity count to determine positive or negative occurrence, and then working correctly with the monotonicity subcalculus of full predicate logic.

Invariance and agent types A second way in which agents enter semantic considerations goes to the notions of structural similarity and invariance that were our running threads. When we ask whether two given structures are 'the same', there is a hidden parameter: 'the same for *whom'*? Agents with restricted powers of inspection or memory see fewer differences than idealized ones. Consider our theme of games, It is attractive to view two extensive games as equivalent when they have the same Backward Induction solution (van Benthem 2014). [22] However, this is a strong assumption about the type of player, viz. as following classical rationality in choosing best actions for oneself given one's preferences and beliefs. Such assumptions are under discussion in modern game theory, where players can also behave quite differently, making 'agent types' a parameter to be chosen explicitly before we can analyze a game– and in the same vein, our discussion of invariance might also need an additional parameter. When are two structures the same for whom?

Diversity once more But here is one more step. The usual implicit assumption in logic seems to be that agents are the same qua powers and styles of behavior. This is true in logical games: although players are allowed to have different information, roles and pay-offs, their styles of observation and reasoning run in exactly the same way. But in reality, agents differ qua powers in all these respects (Liu 2009), and normal scenarios all around us pit very different players against each other, often even humans versus machines. We can of course leave this in place as an idealization. But we may be missing an interesting object of logical study then. What about taking this diversity seriously, and introducing more social scenarios? What are solutions to logic games when players are different, say games of model comparison, or games of dialogue? Or what happens to 'resource logics' (Restall 2000) when we include interplay of agents with different reasoning resources? Bits and pieces of this sort of thinking occur in the literature (cf. the modeling of players with different 'sights' in Grossi & Turrini 2012), but nothing like a theory exists. [23] [24]

Relocating logic Does logic disappear in this concert, perhaps a cacophony, of different agent voices? I do not think so. What we have to give up is one kind of unique logical rationality for all agents. But the role of logic merely shifts to something more subtle. Taking out agent types as a parameter, much structure remains, and logic then regulates the rational interaction of different kinds of agents.

5.4 Conclusion

An agent perspective using games as its vehicle is a natural supplement to our semantic considerations, bridging between structure and syntax in new ways. [25] But if we take this viewpoint seriously, it may also be much more radical then it looks at first sight, and it may start affecting our very understanding of the basic logical notions.

---

[22] To be added: Display the BI-failure for the earlier valid propositional distribution law.

[23] Uniformity assumptions about agents are also pervasive in philosophy, and van Benthem 2014 discusses, e.g., moral agents that are different in how much moral reasoning they are capable of.

[24] Or could we define logic the first abstraction level where we have taken out all agent variety?

[25] Games can also be studied from proof-theoretic and category-theoretic viewpoints, so we are not claiming an exclusive tie with semantics here.

6 Conclusion

This short piece has given a broad picture of the semantic strand of logic, following the main themes of similarity, invariance, consequence, and a bit of games and agency to spice things up. Admittedly, our treatment was sketchy – and in particular, all that we have said about games and agency was meant as an appetizer only, not as a serious introduction. Even so, we hope to have given an impression of the richness of our themes, and of the way they integrate across logic, crossing between different systems and subfields. [26]

7 References

H. Andréka, J. van Benthem & I. Németi, 1998, 'Modal Languages and Bounded Fragments of Predicate Logic', Journal of Philosophical Logic 27:3, 217–274.

H. Andréka, J. van Benthem, N. Bezhanishvili & I. Németi, 2014, 'Changing a Semantics: Opportunism or Courage? In M. Manzano, I. Sain, & E. Alonso, eds., *The life and Work of Leon Henkin*, Birkhaeuser, Zürich, 307–337.

G. Aucher, J. van Benthem & D. Grossi, 2016, 'Modal Logics of Sabotage Revisited', *Journal of Applied Non-Classical Logics*, to appear.

A. Baltag & S. Smets, 2006, 'Dynamic Belief Revision over Multi-Agent Plausibility Models', in G. Bonanno, W. van der Hoek, M. Wooldridge, eds., *Proceedings LOFT'06*, Department of Computing, University of Liverpool, 11–24.

J. Barwise & J. van Benthem, 1999, 'Interpolation, Preservation, and Pebble Games', *Journal of Symbolic Logic* 64:2, 881–903.

J. Barwise & S. Feferman, eds., 1985, *Model-Theoretic Logics*, Springer, New York.

J. Barwise & J. Perry, 1983, *Situations and Attitudes*, the MIT Press, Cambridge MA.

J. van Benthem, 1977, *Modal Correspondence Theory*, Ph.D. dissertation, Mathematical Institute, University of Amsterdam.

J. van Benthem, 1982, 'The Logical Structure of Science', *Synthese* 51:3, 431–472.

J. van Benthem, 1984, 'Questions about Quantifiers', *Journal of Symbolic Logic* 49:2, 443–466.

J. van Benthem, 1986, *Essays in Logical Semantics*, Kluwer, Dordrecht.

J. van Benthem, 1989, 'Logical Constants across Varying Types', *Notre Dame Journal of Formal Logic* 30:3, 315-342.

J. van Benthem, 1991, *Language in Action,* North-Holland, Amsterdam.

J. van Benthem, 1996, *Exploring Logical Dynamics*, CSLI Publications, Stanford University.

J. van Benthem, 1998, 'Dynamic Bits and Pieces', Institute for Logic, Language and Computation, University of Amsterdam

J. van Benthem, 2010, *Modal Logic for Open Minds*, CSLI Publications, Stanford University.

J. van Benthem, 2014, *Logic in Games*, The MIT Press, Cambridge MA.

J. van Benthem, 2016, 'Tracking Information', in K. Bimbó, ed., *Michael Dunn on Information-Based Logics*, Springer, Dordrecht, 363–389.

J. van Benthem & G. Bezhanishvili, 2008, 'Modal Logics of Space', in M. Aiello, I. Pratt-Hartman & J. van Benthem, eds., *Handbook of Spatial Logics*, Springer, Heidelberg, 217–298.

[26] Evidently, we had to leave out many topics. Thus, we have underplayed the role of alternative logics arising on an agent conception of invariance or inference. One reason for this de-emphasis is the 'voracity' of classical logic that, so far, has been able to translate alternative logics faithfully.

J. van Benthem, N. Bezhanishvili & S. Enqvist, 2016, 'A New Game Equivalence and its Logic', Institute for Logic, Language and Computation, University of Amsterdam and Philosophical Institute, Stockholm University.

J. van Benthem, N. Bezhanishvili, S. Enqvist & J. Yu, 2016, 'Instantial Neighborhood Logic', *Review of Symbolic Logic* online.

J. van Benthem, J. van Eijck & B. Kooi, 2006, 'Logics of Communication and Change', *Information and Computation* 204, 1620–1662.

J. van Benthem & D. Bonnay, 2008, 'Modal Logic and Invariance', *Journal of Applied Non-Classical Logics* 18: 2-3, 153–173.

J. van Benthem, B. ten Cate & J. Väänänen, 2009, 'Lindström Theorems for Fragments of First-Order Logic', *Proceedings LICS* 5:3, 1–27.

J. van Benthem & D. Ikegami, 2008, 'Modal Fixed-Point Logic and Changing Models', in A. Avron, N. Dershowitz & A. Rabinovich, eds., *Pillars of Computer Science: Essays Dedicated to Boris (Boaz) Trakhtenbrot on the Occasion of his 85th Birthday*, Springer, Berlin, 146–165.

J. van Benthem & E. Pacuit, 2011, 'Dynamic Logics of Evidence-Based Beliefs', *Studia Logica* 99, 61–92.

E-W Beth, 1963, 'Konstanten van het Wiskundige Denken', *Mededelingen van de Koninklijke Nederlandse Akademie van Wetenschappen,* 26, 231–255.

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

G. Cina, 2017, *Interweaving Modal Logic and Category Theory*, Dissertation, Institute for Logic, Language and Computation, University of Amsterdam.

H. van Ditmarsch, W. van der Hoek & B. Kooi, 2007, *Dynamic-Epistemic Logic*, Cambridge University Press, Cambridge UK.

H. C. Doets, 1996, *Basic Model Theory*, CSLI Publications, Stanford University.

H-D Ebbinghaus & J. Flum, 2005, *Finite Model Theory*, Springer, Berlin.

E. Graedel, W. Thomas & Th. Wilke, eds., 2002, *Automata, Logics, and Infinite Games*, Lecture Notes in Computer Science, Springer Verlag, Heidelberg.

D. Grossi & P. Turrini, 2012, 'Short-Sight in Extensive Games', in V. Conitzer & M. Winikoff, eds., *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multi-Agent Systems* (AAMAS), 805–812.

H. Helmholtz, 1878, *The Facts of Perception*, Wesleyan University Press, Middletown CT.

A. Läuchli, 1970, 'An Abstract Notion of Realizability for which the Predicate Calculus is Complete', in J. Myhill, A. Kino & A. Vesley, eds., *Intuitionism and Proof Theory*, North-Holland, Amsterdam, 227-234.

P. Lindström, 1966, First-Order Predicate Logic with Generalized Quantifiers', *Theoria* 32:3, 186–195.

F. Liu, 2009, 'Diversity of Agents and Their Interaction', *Journal of Logic, Language and Information* 18:1, 23–53.

C. Löding & P. Rohde, 2003, 'Model Checking and Satisfiability for Sabotage Modal Logic', in P. K. Pandya & J. Radhakrishnan, editors, FSTTCS 2003, Springer Lecture Notes in Computer Science 2914, 302–313.

A. Mostowski, 1957, 'On a Generalization of Quantifiers', *Fundamenta Mathematicae* 44, 12–36.

E. Pacuit, 2017, *Neighborhood Semantics for Modal Logic*, Department of Philosophy, University of Maryland. To appear in Springer textbook series *Logic in Asia*.

S. Peters & D. Westerståhl, 2006, *Quantifiers in Language and Logic*, Oxford University Press, Oxford.

G. Plotkin, 1980, 'Lambda Definability in the Full Type Hierarchy', in J. Seldin & J. Hindley, eds., *To H.B. Curry. Essays on Combinatory Logic, Lambda Calculus and Formalism*, 363-373.

G. Restall, 2000, *An Introduction to Substructural Logics*, Routledge, London.

G. Sher, 1991, *The Bounds of Logic*, MIT Press, Cambridge MA.

R. Statman, 1982, 'Completeness, Invariance, and l–Definability', *Journal of Symbolic Logic* 47, 17–26.

P. Suppes, 2002, *Representation and Invariance of Scientific Structures*, CSLI Publications, Stanford University.

A. Tarski, 1986, 'What are Logical Notions?', In J. Corcoran, ed*., History and Philosophy of Logic* 7, 143–154.

J. Väänänen, 2011, *Models and Games*, Cambridge University Press, Cambridge UK.

Y. Venema, 2012, *Lectures on the Modal Mu-Calculus*, Institute for Logic, Language and Computation, University of Amsterdam.

H. Weyl, 1963, *Philosophy of Mathematics and Natural Science,* Atheneum, New York, reprint from the 1930s.