

The Dynamic World of Martin Stokhof

Johan van Benthem, Amsterdam & Stanford

The editors of this tribute have clearly found a natural nexus where duty coincides with inclination. It is a great pleasure to write something to honour a colleague who is both well respected and well liked by people of the most diverse interests and tempers – and I am happy to count myself among them.

Formal semantics Martin Stokhof has a continuous record of major ideas that set the agenda for semantics over the last decades. He helped introduce Montague Grammar, but then went on to work with Jeroen Groenendijk and Frank Veltman on feeding major new themes into, and beyond, that framework. Just think of his work on the pervasive role of information, questions, and dynamics of meaning. These paradigms have not only enriched the post-Montagovian interface of logic and language, but they have also been hooks for new contacts with computer science and cognitive science. I could talk here about my personal take on these trend-setting themes, involving dynamic logics of information and evaluation, games and strategies, the evolutionary grindstone of long-term learning procedures, and the emancipation of the interactive agents that use language and perform other social tasks. But I won't.

Organization Best known to insiders only, Martin Stokhof also has a long record of organizational creativity. His lasting impact can be seen in the still vigorous Amsterdam Colloquia, the ILLC in Amsterdam, and the European umbrella organization FoLLI. I could talk about that perhaps more rare talent, too, since I have been a fellow architect (and conspirator) on many of these institutional fronts. But I won't.

Philosophy Over the recent years, Martin has become much more of a philosopher than a formal semanticist. It was in fact the combination of these two dimensions, importantly including his original new work on Wittgenstein, that carried him into the Dutch Royal Academy of Arts and Sciences. I could talk about that, too, since philosophical reflections on language use have exercised me ever since I became exposed to classical Chinese views. These recommend minimal use of text, with maximal

freedom of interpretation, to capture one's audience ('a net needs holes to catch fish') rather than maximal explicitness enforcing one's own interpretation. But still, I won't.

Dissidence and doom But the two souls living in one breast are not in harmony yet. In recent years, Martin has become a dissident in formal semantics, shaking the very ladder that brought him to his current scientific and institutional status. I am not sure that I agree with all of Martin's recent worries and criticisms, but they do resonate. Indeed, I have often voiced methodological worries about what formal semantics is, starting from my paper 'Why is Semantics What?' in the mid 1980s. Personally, I have always managed to push these ghosts back into the bottle – but sometimes, they refuse. This need not be due to an increase in critical philosophical acumen, but rather a change in the *Zeitgeist*. Nowadays, no one seems happy with the old lofty appeal to expert 'intuitions', usually of the very same professors whose theories were at stake. Through the window, one feels an invigorating breeze of corpus statistics and psychological data, some calling us directly from the brain. What is the status of formal semantics today, our cosy little island of logic in between the probabilistic behavior of brain function and the statistics of large-scale long-term public opinion?

Teaching 'first-order translation' Frankly, I have no deep thoughts to offer on where semantics stands. But talking with Martin recently, we discussed a small example of *Zeitgeist* blues. When introducing logic, I teach the usual skill of 'translating' natural language into the formal language of first-order logic, and in my didactical prose, I unthinkingly copy what generations of colleagues have said about it. But recently, working with the team of a new internet course "Logic in Action", it suddenly occurred to me that I no longer believe in these formulations. Or maybe worse, I no longer understand what we are claiming about this skill, and hence what we are teaching. Clearly, we are not translating in any standard sense of the word. First-order logic is not a language like natural language, and also, in what sense is the 'translation' that we teach a useful improvement? Let me take this as my simple theme for this piece.

Is first-order logic a language at all? Probably it is if one looks just at grammatical structure – the way Montague's Thesis did way back in the 1970s, proclaiming the fundamental unity of natural and formal languages. (But then, I do not understand that

once so convincing thesis so well any more either: another Zeitgeist phenomenon.) But first-order logic is definitely not a language that is used for communication the way other languages are. Still, this does not mean that it has no useful purpose. It is not a ‘living language’, but it is not a ‘dead tongue’ either. So, what is it?

Logics as models for natural language One popular view in the literature is this. Logical languages are abstract *models* for natural languages, and logical systems are models for natural practices of reasoning. I think there is some truth to this, and I also think that the above doubts about the precise status of the system should not overdo it. There is the simple fact that generations of students have found it illuminating to learn how to think in first-order terms, and be sensitized to phenomena such as scope, inference patterns, and more generally, high standards of precision. First-order logic naturally commands allegiance: it is not imposed by force. And such a success for an abstract theory is also a concrete empirical fact! We need to understand what that allegiance means. At least, it seems to be like learning a new mathematical way of thinking, something that we appreciate for the richer view of the world it provides.

The danger of system imprisonment But how good is the model of natural language provided by first-order logic? There is always a danger of substituting a model for the original reality, because of the former’s neatness and simplicity. I have written several papers over the years pointing at the insidious attractions and mind-forming habits of logical systems. Let me just mention one. The standard emphasis in formal logical systems is ‘bottom up’. We need to design a fully specified vocabulary and set of construction rules, and then produce complete constructions of formulas, their evaluation, and inferential behavior. This feature makes for explicitness and rigor, but it also leads to *system imprisonment*. The notions that we define are relative to formal systems. This is one of the reasons why outsiders have so much difficulty grasping logical results: there is usually some parameter relativizing the statement to some formal system, whether first-order logic or some other system. But mathematicians want results about ‘arithmetic’, not about the first-order Peano system for arithmetic, and linguists want results about ‘language’, not about formal systems that model language. Even inside logic, this restriction can be annoying. In studying natural

reasoning, we want to talk about monotonicity inferences, not relativized to first-order logic, or any system at all: just monotonicity reasoning. But how to say it?

And there is still one more reason why system imprisonment is annoying. One argument in favor of logical methods is that they bring a degree of reflection that is unique: meta-theorems about the properties and scope of these methods themselves. But again, meta-theorems are relative to formal systems (first-order logic obeys the Compactness Theorem, other logics do not). Where are the general insights in the meta-theory of logic, and at what level of generality should we state them?

From bottom-up to top-down I used the term ‘bottom up’, and the natural counterpart to that would be ‘top down’. Many real uses of formal methods seem to be top-down. We look at some phenomenon in reality, but we do not translate it into some fully specified formal counterpart. What we rather do is analyze a little bit of relevant structure from the outside, like when we introduce a few equations in the description of some physical scenario. I think that is the better way of viewing what first-order logic does: it does not replace sentences by formulas, but it offers us a way of stylizing some features of a sentence, such as its quantifier scope structure, while perhaps leaving large chunks of the sentence untouched. Viewing things in this way will also remove some misunderstandings of the role that logic has played in the past. It is often said, taking the systems view, that the traditional syllogistic was a ‘small fragment of monadic first-order logic’. This is true in a very technical sense, but it is also very misleading. On the top down view, syllogistic patterns can be discerned in any sort of natural language discourse, where the unanalyzed predicates inside might have arbitrary complexity, potentially going far beyond first-order logic.

Contents versus wrappings Here is one more point where I have long dissented from the received wisdom concerning logical systems in understanding natural language and natural reasoning. We are told that, by Church’s Theorem, complete reasoning in natural language is ‘undecidable’, and hence that it has a much greater complexity than what famous traditional logicians may have thought. But once again, we must ask: exactly what is so complex here? Natural reasoning itself, or the formal system of first-order logic that we use to model this reasoning with? In my papers in the 1990s,

I have shown how one should, and can, make a distinction here between complexity of *contents* versus that of *wrappings*. In particular, it can then be proved that the core validities of quantifier reasoning encoded in first-order logic are still decidable, while it is a number of additional design decisions for the set-theoretic Tarski semantics in terms of full function spaces of available assignments that generate the undecidability. Once more, the model should not be identified with the reality.

Virtues of the first-order model And yet I will not end this part of my musings on a negative note. First, the deconstructionism advocated in the preceding paragraph can also be taken to show a virtue of using formal models. We might not even have noticed the choice points clearly without having the formal system of first-order logic. Moreover, that very system suggests parameters for variation, once we step back and look at its standard built-in assumptions. And this is just one instance of what a well-chosen model can do for you. Almost a century after it was first proposed as an important formal system inside much richer higher-order logics, first-order logic is still a highly successful conceptual laboratory for playing with new ideas. Most recently, for instance, experimentation with first-order logic has led to new views on strategic interaction in games of perfect and imperfect information, and on a rich world of different dependency structures between agents involved in social tasks. And the same youthful vigor will be even clearer when you look at the role of first-order logic toward *programming languages*, the third corner of the triangle, where it has sparked such different paradigms as logic programming and program semantics.

From semantics to conceptual analysis Let me go on a bit in this line, leading up to the other perspective on formal versus natural languages that I want to raise. The idea that logical languages may, or may not, be good ‘models’ for a given natural language is still in line with the idea that some fixed reality is given, which logical systems are then supposed to model. In that sense, the very term ‘semantics’ is conservative, because the given natural language becomes the yardstick for judging the quality of the logical system. But a logical system is *not just* a model for natural language, or some reasoning practice. It is also a tool for independent conceptual analysis of some cognitive activity, which can bring to light important features that are not encoded in natural language as we have it. In that sense, logical systems can also be much more

radical in their thrust, helping us, as Marx advocated, to move away from interpreting the world to *changing* it. I find nothing repugnant to the idea that we might want to change natural language because of logical considerations. In fact, I often worry about the adjective “natural”, which often hides a sort of unthinking admiration for what is supposed to be pristine ‘nature’. In reality, ‘natural language’ is about as natural as a Dutch polder or a Chinese lake: constantly tinkered with by humans. What also helps in seeing this broader conceptual horizon is the fact that a language is first and foremost a *practice*, not the algebra with operators that Montague saw as its core.

Enriching natural language In this light, my own current interest is no longer the whole-sale opposition between formal and natural languages. I have become much more intrigued by the more local phenomenon of mutual influences and *hybridization*. Consider an expert language like that of mathematics. On the one hand, despite a century of mathematical logic, even mathematics research seminars resound happily with natural language as their substratum. But on the other hand, that language has all sorts of formal insertions: formal notation, equations, even the occasional quantifier symbol, when this is needed for the purposes of precision and communication. One should take that mixed language very seriously. It is not a problematic bastard of two pure systems, but a lively and crucial medium of expression and communication. And this is not an exceptional situation. Many influences from computer science have this feature: algorithmic thinking about action is absolutely crucial, and quite natural to the human mind, but it does not lie encoded ideally in natural language. And so, we see how diagrams, and features of program constructions make their way into how we write and communicate instructions. Many more parts of natural language have this mixed character, where technical features come in freely as required. In other words then, pure formalisms do not replace natural language, and they do not ‘lecture it’ on how to improve, but they enrich it as needed for some serious purpose that we humans have. On a small scale then, looking at a student’s notepad, and seeing a sentence with a few logical symbols stuck in is not a peep at some dark intermediate recess of the mind, but at a viable and intriguing new medium of expression in its own right.

Dynamics of language choice Now let me talk about the Amsterdam trademark of dynamics after all, which may soon outperform tulips and windmills. Mixed notations

on your notepad are not a sign of sloppy thinking, but a creative new medium. But the process of using them is dynamic. In one direction, there is the activity of *precisation* or *abstraction*, needed, e.g., when you want to use some algebraic module for quick symbolic computation, say pushing a negation inside a formula via De Morgan laws. Actually, this process of stepping up formalization can also be useful for many other human purposes: responding to a criticism in argumentation, or just communicating very clearly what you want to say. Formalization may sometimes be the greatest form of courtesy, and greater objectivity the best guarantee of inter-subjectivity. But even though many logicians and philosophers will tell you with Messianic fervor that this process of precisation is the inevitable course of history, there is an equally important process that runs in the opposite direction. Often, we want to look at a mass of correct formal detail, and reduce that information to a well-chosen paraphrase identifying the main idea, and communicating that. Indeed, we would often judge whether someone really understands a given text by seeing how well she would be able to paraphrase that text at higher non-technical language level. The two processes, of increasing and decreasing formality, of course work in tandem. In my view, it is that dynamic world of language use and indeed, of language creation, which should be our main focus of investigation. And as a dynamic logician, I think that this world has a lot of technical structure, far richer than the quiet algebraic homomorphisms of Montague's world reflecting structures without interaction, the way a mountain lake reflects some peaks.

Immune system instead of panacea In particular, our skills at language creation defy 'fundamentalism' about using formal languages as a guaranteed medicine for thinking clearly. There is no such guarantee: we can always make errors of formulation, and errors of thinking. Here is an earlier point I have made on many occasions about our dynamic abilities at revising beliefs in the face of inconsistencies that have come to light. The most fascinating feature of human intelligence is not some foundationalist wonder drug for correctness, but rather the working of logic as the *immune system of the mind*, that can learn from mistakes, revise beliefs, change formulations, and do all the things needed to recover whenever something unsatisfactory has come to light. Echoing this, I would say that it is this creative tension between natural and formal languages which leads to innovation, and which should command our respect.

A lesson? So, with all this said, what should I tell my students nowadays about translation into first-order logic? I do not know yet. Certainly not the above musings, because they will think I have gone over the top. The language of textbooks is among the most conservative parts of natural language, and I still need to find a gentle way to push it a bit away from the traditional mantras. But let that moral be my problem.

Conclusion No matter how all this rethinking of old certainties may end, one can draw a much more definite conclusion right now. It is this. Martin Stokhof's intellectual trajectory in life is always worth watching. The 60 years we are celebrating is of course just one milestone in this respect. Martin's thoughts will no doubt have many further twists that will surprise and inspire us. But perhaps more important than intellectual insight is ethical example. Martin's gentle diplomatic, but at the same time direct human presence will continue to make life *better* for all of us.