Philosophy of logical practice:

a case study in formal semantics

NIKHIL MADDIRALA

www.nikhilmaddirala.com

Philosophy of logical practice:
a case study in formal semantics

MSc Thesis *(Afstudeerscriptie)*

written by

Nikhil Maddirala
(born February 17th, 1989 in Hyderabad, India)

under the supervision of **Benedikt Löwe**, and submitted to the Board of Examiners in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:
July 28th, 2014

Members of the Thesis Committee:
Ronald de Wolf
Benedikt Löwe
Arianna Betti
Michael Franke
Gert-Jan Munneke

Institute for Logic, Language and Computation

Abstract

This thesis attempts to delineate the contours of a nascent domain of inquiry which shall be known as "philosophy of logical practice" and to make a modest contribution to this field, by way of a case study in formal semantics. Over the past few decades, logic has spawned a lively scientific community with its own social norms, rules of behavior and procedures for generating new results. Consequently, I believe that an adequate philosophy of logic needs to account for logical practice and provide an explanation for the practices and procedures of the logical community. Philosophy of logical practice seeks to do so by combining historical, philosophical and social scientific studies of logic. In this thesis I demonstrate one possible approach to philosophy of logical practice by way of a case study in formal semantics, which is a particular form of logical practice. The case study seeks to discuss the question "is formal semantics a failed discipline?" by drawing primarily on two methodological frameworks: (1) qualitative research in the social sciences — in particular, this case study is structured as an interview study featuring interviews with critics, insiders and outsiders of formal semantics — and (2) Thomas Kuhn's framework for the understanding the history and philosophy of science. Major themes that emerge in the case study are: (1) the tension between the scientific and philosophical aspirations of formal semantics as a discipline, (2) the putative rivalry between formal semantics (based on mathematical logic) and computational linguistics (based on machine learning and stochastic processes), and (3) the trend towards empirical, data-driven research in the larger field of linguistics.

Acknowledgements

First and foremost, I would like to acknowledge my thesis supervisor and academic mentor, Benedikt Löwe, whose guidance was a *sine qua non* to the conception and production of this thesis. I am greatly indebted to my conversations with him for much of the stimulation of my thoughts relating not only to this thesis, but also to logic, philosophy, mathematics and academia in general. For enthusiastically participating in my research, by allowing me to interview them and permitting me to publish the interview transcripts, I would like to thank the interviewees: Martin Stokhof, Michiel van Lambalgen, Jeroen Groenendijk, Floris Roelofsen, Frank Veltman, Katrin Schulz, Khalil Sima'an, Ivan Titov, Philip Schulz and Remko Scha; needless to say, such a thesis would not have been possible without their cooperation. For carefully reading my thesis and for providing me with valuable feedback, comments and questions, I am grateful to the members of my thesis committee: Ronald de Wolf, Benedikt Löwe, Arianna Betti, Michael Franke and Gert-Jan Munneke.

The two years that I have spent as a Master of Logic student at the Institute for Logic, Language and Computation (ILLC) in Amsterdam have undoubtedly been the most intellectually stimulating and inspiring years of my life. Although I am beholden to the entire ILLC community for this experience, I would like to acknowledge in particular Ulle Endriss, for being the meticulous curator and custodian of the ILLC's Master of Logic program, and Tanja Kassenaar, for the indefatigable administrative services that she provides for this community.

I would also like to acknowledge my undergraduate institution (or in American parlance, my *alma mater*), The University of Chicago, for initiating me into "the life of the mind" and for nurturing in me the intellectual curiosity and scholarly inclination without which I never would have embarked upon this academic trajectory. Lastly (but by no means least), I would like to express the deepest gratitude to my parents — Subrahmanyam Maddirala and Pranathi Reddy — for their unconditional love and support throughout all of my life's endeavors thus far.

"The aim of philosophy, abstractly formulated, is to understand how things in the broadest possible sense of the term hang together in the broadest possible sense of the term."

Wilfrid Sellars (1963)

Contents

Chapter 1

Executive summary

This thesis attempts to broadly delineate the contours of a nascent and almost altogether uncharted domain of inquiry which shall be known as "philosophy of logical practice" (PLP) and to make a modest contribution to this field, by way of a case study in formal semantics, which is a particular form of logical practice.

PLP is fundamentally motivated by the following consideration: over the past few decades, logic has spawned a lively scientific community with its own social norms, rules of behavior and procedures for generating new results; therefore an adequate philosophy of logic needs to account for logical practice and to provide an explanation for the procedures and practices of the logical community. PLP seeks to do so by combining historical, philosophical, and social scientific studies of logic. Although PLP is a nascent and almost altogether uncharted domain of inquiry, there is a closely related and reasonably well charted domain of inquiry known as "philosophy of mathematical practice" (PMP) from which I draw inspiration and guidance. PMP is an interdisciplinary approach to philosophy of mathematics that focuses on mathematical practice, i.e. the scientific practice of mathematics, rather than on abstract philosophical and mathematical models of mathematics that have been the focus of traditional philosophy of mathematics. Similarly PLP is an interdisciplinary approach to philosophy of logic that focuses on logical practice. PLP can be described by the following key characteristics:

1. *Critical attitude towards logic.* Traditional philosophy of logic usually assumes that logic is a special science (and a special human activity) and that logical knowledge enjoys a special kind of objectivity and certainty. PLP maintains a critical attitude towards logic by treating it as an ordinary (as opposed to special) scientific practice and as an ordinary human activity.

2. *Focus on logical practice.* Traditional philosophy of logic is inadequate for modern logical practice, because it typically pays attention to only a few key areas of logic and to the products or results of those areas, but it ignores the actual practice of logic. PLP on the other hand seeks to take into account various forms of logical practice and to provide an explanation for the practices and procedures of the logical community.

3. *Interdisciplinary perspective.* Traditional philosophy of logic is limited by the disciplinary boundaries of philosophy, logic and mathematics. PLP is committed to an interdisciplinary methodology that includes not only philosophical, logical and mathematical perspectives, but also historical, social scientific and other related perspectives.

In this thesis I demonstrate one possible approach to PLP by way of a case study in formal semantics, which is a particular form of logical practice. Formal semantics was chosen as the subject of this case study primarily due to the fact that there was a serious meta discussion about the successes and failures of this discipline within the formal semantics community. In particular, Martin Stokhof and Michiel van Lambalgen (two prominent (former) formal semanticists) initiated a significant debate in the journal *Theoretical Linguistics*, where they raised the question of whether formal semantics could be an example of a *failed discipline*. The debate centers around the way in which formal semantics conceptualizes its central objects of study so as to fit a particular methodology. They allege that the central objects of study of formal semantics (i.e. "language", "meaning", etc.) have been deliberately constructed through a process of idealization, which does not meet the standards of a rigorously scientific inquiry — in particular, they allege that formal semantics does not meet the criterion of empirical verifiability / falsifiability.

Taking this debate as a point of departure, the case study seeks to discuss the question "is formal semantics a failed discipline?" by drawing primarily on two methodological frameworks: (1) qualitative research in the social sciences — in particular, this case study is structured as an interview study featuring interviews with critics, insiders and outsiders of formal semantics — and (2) Thomas Kuhn's framework for the understanding the history and philosophy of science. I take this to be an exercise in PLP as opposed to philosophy of logic, because it adheres to the three characteristic features of PLP described in the previous chapter: (1) The case study adopts a critical attitude towards formal semantics, treating it as an ordinary scientific practice and an ordinary human activity. (2) The case study is focused on the practice of formal semantics as opposed to the formal models that it produces. (3) The case study adopts an interdisciplinary approach by drawing on the Kuhnian perspective in the philosophy of science together with the interview methodology of qualitative research in the social sciences.

Logical practice in general presupposes and depends crucially on the processes of modeling and formalization; in particular, the logician will produce a formal model of some real world phenomenon in order to represent it and to reason about it. Formal semantics, being a particular form of logical practice, is no exception to this general rule of thumb. In this case the relevant real world phenomenon to be modeled is natural language. The formal semanticist will produce a formal model of natural language (or at least a model of some particular region of natural language) in order to represent it and to reason about it. Given that formal semantics is concerned with modeling, and given that modeling is a process that replaces one thing with another — in particular, it replaces the real world phenomenon of natural language with its formal representation in a formal language — care must be taken to make sure that the model stays close to the modeled object. This requires some external benchmark against which the models can be measured for success or failure. Typically in philosophy and in the humanities expert intuition is used as the benchmark, whereas in the natural sciences experimental data is used as the benchmark. But what about formal semantics? Here formal semantics is torn between its philosophical and its scientific aspirations.

Traditionally formal semantics has been interested in a conceptual analysis of

abstract, structural features of natural language such as the concept of propositional or semantic meaning. This conceptual analysis is quite detached from any empirical or experimental motivation and is guided primarily by the intuitions of the researchers. The critics criticize such a methodology because it leads the researchers to idealize their object of study in a manner that renders it unsuitable for empirical investigation. Moreover, the use of intuitions and conceptual analysis in this manner is now seen to be in conflict with the discipline's own scientific aspirations, according to which "empirical ratification of analytical work is our main ambition and touchstone for success."[1] Thus, we are now beginning to see a shift in formal semantics away from intuitions towards empirical data, resulting an identity crisis for the discipline. The critics urge that formal semantics should concretely determine its own identity by choosing between one of its two aspirations: either become rigorously scientific by focusing on observable linguistic behavior, or give up the aspiration to be a rigorously scientific discipline. However, based on my interviews, it seems to me that formal semanticists are unable to decide on an identity, and they seem to suffer from the syndrome of "wanting to have their cake and eat it too."

The formal semanticists interviewed in this study all acknowledge the criticisms made by the critics. They recognize that there is some degree of idealization involved in their process of modeling natural language by means of formal representations, and they further acknowledge that this also poses problems for the empirical validity of their models. Yet they maintain that their discipline is rigorously scientific. Although they recognize that they are lacking an experimental methodology which would allow for empirical verification or falsification of their models, they explain this problem away by delegating the empirical work to some point in future or to some other group of researchers. By the critics' lights, however, this would count merely as an avoidance rather than solution to the problem, because formal semantics has not yet given a clear account of how such an experimental methodology could be developed, or even if it is in principle possible for their theoretical models to be related to empirical work.

Additionally, we see that there are sociological processes at work, and that they are equally important to the future of this debate. It seems clear (to the

[1] Institute for Logic, Language and Computation (2014a)

critics as well as the formal semanticists) that the success or failure of formal semantics will depend to a large degree on its ability to sustain a thriving community of students and researchers. One of the critics points out that despite some influx of young researchers into the field, it seems to lack the kind of intellectual coherence that is required to sustain the discipline. Moreover, the formal semanticists are themselves quite pessimistic about the future of the discipline. They point out that what was once their home community — linguistics — is now reluctant to accept their work, because the linguists want to move away from the stereotype of "armchair linguistics" towards a more rigorously empirical, data-driven approach. Therefore formal semanticists are finding it increasingly difficult to sustain a thriving community in what was once their home.

A related development is that formal semantics is also beginning to feel threatened by the emergence of competing disciplines such as computational linguistics (which is more closely aligned with the empirical demands of the linguistics community). In addition to its scientific and philosophical aspirations, formal semantics also voices some engineering aspirations to be involved in the building of practical applications such as a machine or computer program that understands natural language. With respect to these aspirations, computational linguistics has enjoyed far more success than formal semantics. The critics also point to this as an example of formal semantics' failure, and they note that an important reason for the success of computational linguistics is the fact that their models stay closer to the original phenomenon (of natural language) and hence are more amenable to empirical testing. In response to this, the formal semanticists allege that the models of computational linguists are *merely* engineering devices and that they do not contribute to a deeper scientific or philosophical understanding of natural language. My interviews with the computational linguists revealed that there is indeed some truth to this response — at least in the sense that the overtly theoretical and philosophical aspirations of formal semantics certainly have no place in a discipline such as computational linguistics. In the end, it remains unclear where such questions have a place.

In conclusion, I would like to note that the outcome of these various developments remains very much unclear at this stage, but I would like to suggest that we might look to the careers of the two critics as an example of what the future

might hold. Sharing some fundamental dissatisfactions and criticisms about the discipline, they have both come to believe that formal semantics as it is practiced today is able to live up to neither its philosophical nor its scientific aspirations. This belief has caused them both to give up on formal semantics, but has lead each of them in different directions. One critic has abandoned not only formal semantics, but also his aspiration to study language strictly scientifically. Having abandoned his scientific aspirations, he is now primarily occupied with philosophical and hermeneutical investigations into the nature of language. The other critic, however, retains his scientific aspirations with respect to language. Having abandoned formal semantics, he has gone in the direction of empirical research in cognitive science, by means of which he hopes to better understand the brain processes associated with language. In sum, they either gave up their philosophical aspirations and entered into a more rigorously scientific and empirical discipline, or they gave up their rigorously scientific aspirations and entered into a philosophical discipline. The very same trend is also observed in some of the other (current and former) formal semanticists interviewed for this study.

With this case study I have demonstrated one particular approach to PLP, but with this thesis I also hope to have sparked some interest in in PLP more generally. Besides making a contribution to the debate around formal semantics, I also hope to have encouraged formal semanticists and logicians more generally to engage in some philosophical reflection on the way in which they make use of formal models. Hopefully I have convinced the reader that it is a worthwhile and fruitful academic endeavor with the potential to reveal something new and interesting. It would be a great pleasure for me to see more logicians, philosophers, mathematicians, linguists, historians, social scientists and other researchers engage in such explorations. A good starting point for further exploration would be an expansion of this existing case study into something richer. The material covered in the interviews is far richer than the analysis provided in the main chapters of this thesis. My expertise on this subject matter and my capacity for analysis are limited, but perhaps someone with more expertise, time and other resources will be able to use this material as a data point for further investigation. Beyond this, there is always the possibility of additional case studies and additional methodologies (e.g. quantitative research).

Chapter 2

Introducing philosophy of logical practice

This thesis attempts to delineate the contours of a nascent and almost altogether uncharted domain of inquiry which shall be known as "philosophy of logical practice" (PLP) and to make a modest contribution to this field, by way of a case study in formal semantics, which is a determinate form of logical practice. PLP is fundamentally motivated by the following consideration: over the past few decades, logic has spawned a lively scientific community with its own social norms, rules of behavior and procedures for generating new results; therefore an adequate philosophy of logic needs to account for logical practice and to provide an explanation for the procedures and practices of the logical community. PLP seeks to do so by combining historical, philosophical, and social scientific studies of logic.

Although PLP is an almost altogether uncharted domain of inquiry, there is a closely related, partially charted domain of inquiry known as "philosophy of mathematical practice" (PMP) from which I draw inspiration and guidance.[1] In this chapter, I will first (in section 2.1) sketch an overview of the motivation and development of the PMP movement against the background of traditional

[1] As far as I know, there has only been one attempt made to chart the domain which I call PLP — by Catarina Dutilh Novaes. That attempt, which is discussed further in section 2.2, also used a different label. Instead, I use the label PLP to signal closer affinity to the PMP movement. At the time of this writing, the search query "philosophy of logical practice" returns zero results on Google Search.

philosophy of mathematics, and then (in section 2.2) I will argue that philosophy of logic is in need of an analogous PLP movement.

2.1 Philosophy of mathematical practice

Philosophy of mathematical practice (PMP) is an interdisciplinary approach to philosophy of mathematics that focuses on mathematical practice, i.e. the scientific practice of mathematics, rather than on abstract philosophical and mathematical models of mathematics that have been the focus of traditional philosophy of mathematics.

> The philosophy of mathematical practice movement arises from a dissatisfaction among philosophers and historians with abstract models of mathematics that make a mystery of its growth and fail to explain how finite, embodied, naturally evolved creatures can understand it.[2]

Since PMP arises from a dissatisfaction with traditional philosophy of mathematics, it will be useful to begin with some considerations about that tradition and the subsequent dissatisfactions that it gave rise to. Philosophy of mathematics has been traditionally (in traditional departments of philosophy and mathematics) understood to be a special branch of philosophy of science. For instance, the Stanford Encyclopedia of Philosophy (SEP) entry on philosophy of mathematics begins as follows:

> If mathematics is regarded as a science, then the philosophy of mathematics can be regarded as a branch of the philosophy of science, next to disciplines such as the philosophy of physics and the philosophy of biology. However, because of its subject matter, the philosophy of mathematics occupies a special place in the philosophy of science.[3]

[2] Larvor (2014)
[3] Horsten (2014)

Such phraseology suggests that what underlies this traditional approach to philosophy of mathematics is a commitment to the following two propositions:

(1) Mathematics is a science just as physics and biology are sciences, and therefore philosophy of mathematics is a branch of philosophy of science just as philosophy of physics and philosophy of biology are branches of philosophy of science.

(2) Mathematics occupies a special place in science that sets it apart from physics and biology, and therefore philosophy of mathematics occupies a special place in the philosophy of science that sets it apart from philosophy of physics and philosophy of biology.

Although these propositions are not mutually contradictory, emphasizing one proposition over the other does result in a correspondingly different picture of philosophy of mathematics. Consider the first proposition, according to which philosophy of mathematics is a branch of philosophy of science. If philosophy of science is "the branch of philosophy that is centered on a critical examination of the sciences: their methods and their results,"[4] one might conclude that philosophy of mathematics is the branch of philosophy that is centered on a critical examination of the mathematics, its results and its methods. However, philosophy of mathematics has traditionally laid far more emphasis on proposition (2), according to which mathematics occupies a special place in science. For instance, the above quoted passage from the SEP entry on philosophy of mathematics continues as follows:

> Whereas the natural sciences investigate entities that are located in space in time, it is not at all obvious that this also the case of the objects that are studied in mathematics. In addition to that, the methods of investigation of mathematics differ markedly from the methods of investigation in the natural sciences. Whereas the latter acquire general knowledge using inductive methods, mathematical knowledge appears to be acquired in a different way: by deduction from basic principles. The status of mathematical knowledge also

[4] Audi (1999), p.700

> appears to differ from the status of knowledge in the natural sciences. The theories of the natural sciences appear to be less certain and more open to revision than mathematical theories. For these reasons mathematics poses problems of a quite distinctive kind for philosophy. Therefore philosophers have accorded special attention to ontological and epistemological questions concerning mathematics.[5]

Thus typical expositions of traditional philosophy of mathematics begin with the rather uncritical assumptions that mathematics is a special science (and indeed a special human activity) with a special object of investigation, that it has a special methodology and that mathematical knowledge enjoys a special kind of objectivity and certainty. This is PMP's first point of dissatisfaction with traditional philosophy of mathematics: *traditional philosophy of mathematics does not adopt a sufficiently critical attitude towards mathematics.* PMP on the other hand seeks to maintain a critical attitude towards mathematics by treating mathematics as an ordinary (as opposed to special) scientific practice just like physics or biology, and as an ordinary human activity.

In order to adopt such a critical attitude towards mathematics, Van Kerkhove and Van Bendegem.[6] In doing so, they wish to combat the idea that "mathematics is a free creation of the human spirit."[7] According to them, this overused quotation "expresses the cherished beliefs that many share: mathematics stands on its own, free from any societal influence, individualist and immaterial, beyond space and time, in short, it occupies a universe of its own."[8]. Against this view, PMP calls for mathematics to be treated as an ordinary scientific practice and therefore an ordinary human activity that is embedded within a determinate social and historical context.

Besides the uncritical assumptions regarding the so-called "special" status of mathematics, what truly sets traditional philosophy of mathematics apart from the rest of philosophy of science is the extensive use of mathematical methods and formalisms; as the SEP entry points out: "it has turned out that to some

[5] Horsten (2014)
[6] Van Kerkhove and Van Bendegem (2007), p. vii
[7] Van Kerkhove and Van Bendegem (2007), p. vii
[8] Van Kerkhove and Van Bendegem (2007), p.vii

extent it is possible to bring mathematical methods to bear on philosophical questions concerning mathematics."[9] In particular, the development of mathematical logic in the late nineteenth and early twentieth centuries provided a mathematical apparatus that was used by philosophers and mathematicians to systematically investigate the foundations of mathematics (i.e. to investigate the nature of mathematical objects, the laws that govern them and how we acquire knowledge about them). The SEP entry continues:

> In the twentieth century, research in the philosophy of mathematics revolved mostly around the nature of mathematical objects, the fundamental laws that govern them, and how we acquire mathematical knowledge about them. These are foundational concerns that are intimately connected with traditional metaphysical and epistemological questions.[10]

However, an unfortunate consequence of this emphasis on foundations of mathematics and mathematical logic was that philosophy of mathematics became focused almost entirely on these foundational studies to the exclusion of all other areas and aspects of mathematics. Moreover, the focus has always been on the mathematical and formalisms themselves as opposed to the practice of mathematics that produced such formalisms:

> Twentieth-century research in philosophy of mathematics was mainly focused on foundational studies. ... All these approaches were mainly, if not exclusively, focused on the outcomes or "products" of mathematical practice ... however, it has become clear to an increasing number of scholars that a full understanding of mathematics also involves a grip on mathematical activity itself, as a process.[11]

Thus, these developments in philosophy of mathematics systematically alienated the sympathies of the average mathematician. This is PMP's second major

[9] Horsten (2014)
[10] Horsten (2014)
[11] Van Kerkhove and Van Bendegem (2002)

point of dissatisfaction: *traditional philosophy of mathematics is inadequate for actual mathematical practice, because it pays attention only to a few key areas of logic, and it further assumes that all that is interesting about logic can be captured in a system of formal representations (e.g. mathematical logic, set theory, etc.), but it ignores most of the actual practice of mathematics.* PMP on the other hand seeks to take into account various forms of mathematical practice and to provide an explanation for the practices and procedures of the mathematical community.

These philosophical and mathematical investigations into the the foundations of mathematics were bolstered by the implicit assumption that everything that is interesting about mathematics can be reduced to and represented in the foundational framework of mathematical logic (or category theory, or some other such formal symbolism). This idea blinded philosophers of mathematics (among others) to large swaths of mathematical practice.

> The idea that all of mathematical activity can in principle be represented by sequences of formal statements in some adequate system of logic obstructed the view towards what mathematicians are really doing. In fact, sociology of science mostly ignored mathematics presumably under the assumption that the human component of mathematical research is negligible.[12]

This suggests that the philosophical puzzles about mathematics can be solved purely by a mastery of the formalisms themselves, and that therefore there is no place for social scientists, historians and other academicians in addition to philosophers and mathematicians. This leads to PMP's third major dissatisfaction with traditional philosophy of mathematics: *traditional philosophy of mathematics is limited by the disciplinary boundaries of philosophy and mathematics.* (PMP on the other hand is committed to an interdisciplinary methodology that includes not only mathematical and philosophical perspectives, but also historical, social scientific and other related perspectives.) Löwe and Müller emphatically call attention to this goal of interdisciplinarity in their manifesto:

[12] Löwe and Müller (2010)

> But the solution to the puzzle of the objectivity of mathematical knowledge cannot be solved by philosophers alone. Involvement with mathematical practice means that other disciplines, such as the history of science, the fields of science education, sociology of science, cognitive science, and possibly psychology hold parts of the answer to our questions. Interdisciplinary exchange of ideas is a necessity in our attempts to understand the special nature of mathematics. The purpose of the research network PhiMSAMP ("Philosophy of Mathematics: Sociological Aspects and Mathematical Practice") was to catalyze this interdisciplinary exchange and to create a basis for communication between the involved research areas.[13]

Thus in the twentieth century, philosophers, mathematicians, social scientists and other academics who were dissatisfied with traditional philosophy of mathematics began to address these dissatisfactions by moving away from traditional philosophy of mathematics. The SEP entry on philosophy of mathematics concludes by signaling some important shifts in twentieth century philosophy of mathematics:

> In the second half of the twentieth century, research in the philosophy of science to a significant extent moved away from foundational concerns. Instead, philosophical questions relating to the growth of scientific knowledge and of scientific understanding became more central. As early as the 1970s, there were voices that argued that a similar shift of attention should take place in the philosophy of mathematics.
>
> For some decades, such sentiments remained restricted to a somewhat marginal school of thought in the philosophy of mathematics. However, in recent years the opposition between this new movement and mainstream philosophy of mathematics is softening. Philosophical questions relating to mathematical practice, the evolution of mathematical theories, and mathematical explanation and understanding have become more prominent, and have been related

[13] Löwe and Müller (2010), p.vii - viii

> to more traditional questions from the philosophy of mathematics (Mancosu 2008) ... This trend will doubtlessly continue in the years to come.[14]

The work cited in the passage as an exemplar of this shift is a volume edited by Mancosu titled "The philosophy of mathematical practice"[15] which is a collection of essays unified by "the shared belief that attention to mathematical practice is a necessary condition for a renewal of the philosophy of mathematics."[16] [17] Thus PMP began to emerge in the late twentieth century as new movement in the philosophy of mathematics — "opposing with great sensitivity the ahistorical received view in the philosophy of mathematics to a recently emerging trend of studies in contextualized mathematical practices."[18] Although it is still quite far from achieving mainstream status, the PMP movement has now acquired a formidable scholarly literature[19], regular conferences[20], and an international society — the Association for the Philosophy of Mathematical Practice (APMP).[21]

In brief summation, I think that PMP can be defined by three key characteristics:

(1) *Critical attitude towards mathematics.* Traditional philosophy of mathematics does not adopt a sufficiently critical attitude towards mathematics; instead it begins with the rather uncritical assumptions that mathematics is a special science (and indeed a special human activity) with a special object of investigation, and a special methodology, and that mathematical knowledge enjoys a special kind of objectivity and certainty. PMP on the other hand seeks to maintain a critical attitude towards mathematics by treating mathematics as an ordinary (as opposed to special) scientific practice just like physics or biology, and as an ordinary human activity.

[14] Horsten (2014)
[15] Mancosu (2008)
[16] Mancosu (2008), p. 2
[17] It is also worth noting that many in the PMP community do not consider Mancosu's volume to be representative of the field. For example, in Larvor's review of the volume, he says that "Some Philosophy of Mathematical Practice" would have been a better title for this volume, because it is on the conservative end of the literature. See Larvor (2010), p. 359
[18] Van Kerkhove (2008), p. v
[19] For a representative list of the scholarly literature, see Larvor (2010)
[20] For a representative list of the conferences, see Larvor (2014)
[21] Association for the Philosophy of Mathematical Practice (2014)

(2) *Focus on mathematical practice.* Traditional philosophy of mathematics is inadequate for actual mathematical practice, because it pays attention only to a few key areas of logic, and it further assumes that all that is interesting about logic can be captured in a system of formal representations (e.g. mathematical logic, set theory, etc.), but it ignores most of the actual practice of mathematics. PMP on the other hand seeks to take into account various forms of mathematical practice and to provide an explanation for the practices and procedures of the mathematical community.

(3) *Interdisciplinary perspective.* Traditional philosophy of mathematics is limited by the disciplinary boundaries of philosophy and mathematics and is thereby limited in its scope and methodology. PMP on the other hand is committed to an interdisciplinary methodology that includes not only philosophical and mathematical perspectives, but also historical, social scientific and other related perspectives.

2.2 Philosophy of logical practice

Inspired by the philosophy of mathematical practice (PMP) movement described in the above section (2.1), I believe that an analogous movement which should be called "philosophy of logical practice" (PLP) is necessary in the field of logic today.[22] The reasons for which I advocate such a movement are roughly the same as the reasons for which PMP was advocated: (1) traditional philosophy of logic does not adopt a sufficiently critical attitude towards logic, (2) traditional philosophy of logic is not adequate for actual logical practice, and (3) traditional philosophy of logic is confined by the disciplinary limitations of philosophy, logic and mathematics.

To begin with, consider this passage from the Cambridge Dictionary of Philosophy, which offers the following brief definition of philosophy of logic:

[22] As noted in the footnote at the beginning of this chapter, I am not the first to propose such a movement. Catarina Dutilh Noaves' attempt is discussed towards the end of this section.

> [T]he arena of philosophy devoted to examining the scope and nature of logic. Aristotle considered logic an organon, or foundation, of knowledge. Certainly, inference is the source of much human knowledge. Logic judges inferences good or bad and tries to justify those that are good. One need not agree with Aristotle, therefore, to see logic as essential to epistemology. Philosophers such as Wittgenstein, additionally, have held that the structure of language reflects the structure of the world. Because inferences have elements that are themselves linguistic or are at least expressible in language, logic reveals general features of the structure of language. This makes it essential to linguistics, and, on a Wittgensteinian view, to metaphysics. Moreover, many philosophical battles have been fought with logical weaponry. For all these reasons, philosophers have tried to understand what logic is, what justifies it, and what it tells us about reason, language, and the world.[23]

Such a conception of philosophy of logic might have been adequate to the practice of logic in its early days, but it is certainly not adequate to the practice of logic today, and it suffers from many of the same problems that plagued philosophy of mathematics. It is fundamentally inadequate to the practice of logic because it is tied to a very limited conception of logic. To see this contrast, let us compare this limited definition of logic with a modern mission statement for logical practice. Here is the limited definition of logic that the Cambridge Dictionary of Philosophy uses:

> Logic might be defined as the science of inference; inference, in turn, as the drawing of a conclusion from premises. A simple argument is a sequence, one element of which, the conclusion, the others are thought to support. A complex argument is a series of simple arguments. Logic, then, is primarily concerned with arguments.[24]

According to this narrow conception, logic is simply the science of inference, and therefore research in logic is limited to an analysis of arguments and inferences.

[23] Audi (1999), p.679
[24] Audi (1999), p.679

It is true that initially (and up until the 19th century) logic was a branch of philosophy that was concerned primarily with the validity of arguments and inferences in philosophical debates. However, already in the late 19th and early 20th centuries logic took on more technical and mathematical motivation, was developed significantly by mathematical methods, and it began to play a central role in the creation of the disciplines of computer science and artificial intelligence. More recently it has also found applications in linguistics, cognitive science. In contrast to the traditional conception quoted above, for a modern conception of logic and its practice, consider the scientific mission statement of the Institute for Logic, Language and Computation (ILLC), one of the world's leading institutions for research in logic:

> The scientific mission of the Institute for Logic, Language and Computation (ILLC) is to study formal properties of information, viz. the logical structure and algorithmic properties of processes of encoding, transmitting and comprehending information. Information here is to be viewed in its broadest sense, from the flow of information in natural and formal languages to the information contained in music and graphics. The research aim is to develop logical systems that can handle this rich variety of information, making use of insights across such disciplines as linguistics, computer science, cognitive science, artificial intelligence and philosophy.[25]

The actual practice of logic today is much wider in scope than the outdated dictionary definition of philosophers suggests. Far from being limited to an analysis of argumentation and inference, modern logic seeks to analyze all forms of information from natural language discourse to music and graphics. Johan van Benthem, one of the architects of the ILLC and its mission, expresses dissatisfaction with the philosophy of logic that is quite similar to PMP's dissatisfaction with philosophy of mathematics:

> One of the things that strikes me in the philosophy of logic is its great distance from actual research practice. To give one example, after almost a century of model theory and recursion theory, it is

[25] Institute for Logic, Language and Computation (2014b)

> still routinely claimed that logic is essentially about consequence and proof, rather than (also) about truth, meaning, and computation. And the gap gets even wider with modern logics of agency and interaction. ... by doing all this, the philosophy of logic 'deep-freezes' an old, traditional image of the field, making the (real?) logician a theorem-proving applied mathematician, — and nostalgia for the grand old age of foundational research in the 1930s the yardstick for the 21st century. [26]

Thus it seems that logic has expanded in scope beyond its traditional philosophical and mathematical roots, but philosophy of logic has not kept up with these developments. Moreover, there is also a sociological dimension to this issue. Logic was originally treated as a sub-discipline of philosophy or of mathematics, and therefore the logician was treated as a sub-species of mathematician or philosopher. However, it is increasingly developing its own scientific community — rooted in places such as the ILLC. In addition to the ILLC, there are several research institutes, regular conferences, summer schools and graduate programs devoted to research in logic. It is possible to get a master's degree, a PhD degree and a professorship exclusively in the field of logic and dedicate one's entire academic research career to the field (which would include the training of future logicians). The ILLC's critical reflection on its Master of Logic program contains an extensive list of similar departments at other universities that offer degrees in the field of logic.[27]

With these developments, it is also quite clear that over the past few decades logic has spawned a lively scientific community, which naturally comes with its own social norms, rules of behavior, and procedures for generating new results. Consequently, I believe that an adequate philosophy of logic needs to account for logical practice and to provide an explanation for the procedures and practices of the logical community. Therefore I advocate the development of a PLP field that can do so by combining historical, philosophical, and social studies of logic. Much like PMP, the characteristic features of the proposed PLP field would be as follows:

[26] van Benthem (2009)

[27] Institute for Logic, Language and Computation (2013), Appendix A

(1) *Critical attitude towards logic.* Traditional philosophy of logic does not adopt a sufficiently critical attitude towards mathematics; instead it begins with the rather uncritical assumptions that logic is a special science (and indeed a special human activity) with a special object of investigation, and a special methodology, and that logical knowledge enjoys a special kind of objectivity and certainty. PLP on the other hand seeks to maintain a critical attitude towards logic by treating logic as an ordinary (as opposed to special) scientific practice just like physics or biology, and as an ordinary human activity.

(2) *Focus on logical practice.* Traditional philosophy of logic is inadequate for actual logical practice, because it pays attention only to a few key areas of logic and the products or results of those areas, but it ignores most of the actual practice of logic. PLP on the other hand seeks to take into account various forms of logical practice and to provide an explanation for the practices and procedures of the logical community.

(3) *Interdisciplinary perspective.* Traditional philosophy of logic is limited by the disciplinary boundaries of philosophy, logic and mathematics and is thereby limited in its scope and methodology. PLP on the other hand is committed to an interdisciplinary methodology that includes not only philosophical, logical and mathematical perspectives, but also historical, social scientific and other related perspectives.

Finally, I would like to point out that PLP is an approach to philosophy of logic that is not completely uncharted, and neither is it completely without precedent. In particular there has been one prior attempt to chart this domain of inquiry, and there are two philosophical traditions (precedents) to which such an approach is indebted. The prior attempt to chart this domain of inquiry was an approach to philosophy of logic and philosophy of mathematics pioneered by Catarina Dutilh Novaes called "Practice-based philosophy of logic and mathematics." A few years ago she convened a workshop titled "Practice-based philosophy of logic and mathematics"[28] and followed this up with a paper in which she advocates an approach that incorporates the study of actual human and scientific practice (of logic) into the philosophical analysis.[29] Although we

[28] Dutilh Novaes (2009)
[29] Dutilh Novaes (2012)

use different labels (my choice of label seeks to emphasize the close relationship to the PMP movement), I believe that we are essentially proposing the same kind of movement.[30]

Then there are two philosophical traditions which serve as precedents or precursors to the PLP movement: (1) The first precursor is the "psychology of reasoning" approach to logic pioneered by Michiel van Lambalgen. In "Human Reasoning and Cognitive Science"[31], Stenning and van Lambalgen argue that the formal study of logic should be closely related to empirical studies of how human beings actually reason in the wild. PLP certainly wants to follow in this tradition of interdisciplinary inquiry into logic with a focus on empirical studies. (2) Then, in a broader and more general philosophical sense (i.e. not limited to logic) the experimental philosophy movement pioneered by Knobe and Nichols can be considered to be a precursor to PLP. In their "manifesto" for experimental philosophy, they proclaim:

> Unlike the philosophers of centuries past, we think that a critical method for figuring out how human beings think is to go out and actually run systematic empirical studies. Hence, experimental philosophers proceed by conducting experimental investigations of the psychological processes underlying people's intuitions about central philosophical issues.[32]

Following in the same tradition, I believe that this thesis can be loosely considered as an exercise in "experimental philosophy" although the focus of this thesis is more on social and historical processes rather than psychological processes.

[30] Unfortunately I was introduced to the work of Dutilh Novaes only at a very late stage of my thesis writing, and therefore I don't have anything very interesting to say about the relation between her project and mine. Otherwise, I would have liked to comment on this in detail, considering that our goals are nearly identical.

[31] Stenning and van Lambalgen (2008)

[32] Knobe and Nichols (2008)

Chapter 3

Introducing the case study in formal semantics

In this thesis I demonstrate one possible approach to philosophy of logical practice (PLP) by way of a case study in formal semantics, which is a particular form of logical practice. Formal semantics was chosen as the subject of this case study primarily due to the fact that there was a serious meta discussion about the successes and failures of this discipline within the formal semantics community. In particular, Martin Stokhof and Michiel van Lambalgen (two prominent (former) formal semanticists) initiated a significant debate in the journal *Theoretical Linguistics*, where they raised the question of whether formal semantics could be an example of a *failed discipline*. Starting with this debate as a point of departure, the case study seeks to discuss the question "is formal semantics a failed discipline?" by drawing primarily on two methodological frameworks: (1) qualitative research in the social sciences — in particular, this case study is structured as an interview study featuring interviews with critics, insiders and outsiders of formal semantics — and (2) Thomas Kuhn's framework for the understanding the history and philosophy of science. I take this to be an exercise in PLP as opposed to traditional philosophy of logic, because it adheres to the three characteristic features of PLP described in the previous chapter: (1) The case study adopts a critical attitude towards formal semantics, treating it as an ordinary scientific practice and an ordinary human activity. (2) The case study is focused on the practice of formal semantics as opposed to the formal models

that it produces. (3) The case study adopts an interdisciplinary approach by drawing on the Kuhnian perspective in the philosophy of science together with the interview methodology of qualitative research in the social sciences.

In this chapter, I will first (in section 3.1) describe the context and motivation for the case study, and then (in section 3.2) I will describe the methodologies adopted and the research design of the case study. Finally (in section 3.3) I will issue some disclaimers about the limitations of such a case study.

3.1 Setting the context

In this section I will describe the context of this case study in formal semantics. In order to do this, I will first (in subsection 3.1.1) give a brief overview of formal semantics as a discipline and as a scientific community, and then (in subsection 3.1.2) I will discuss the debate around its scientific status, which gives rise to the question of its failure (the question that is the subject of my case study).

3.1.1 Formal semantics as a discipline

In this subsection, I would like to give a very brief overview of formal semantics as a discipline and as a scientific community. The major part of this description — in terms of the nature, goals and methodology of the discipline — will emerge from the debate itself and from the analysis of the interviews (i.e. from the following chapters of the thesis), but it is useful to have some context and background in advance.

Formal semantics began as an interdisciplinary approach to the study of language, drawing upon sources from logic, philosophy and linguistics. The most important figure for this development was Richard Montague, who pioneered a formal logical approach to natural language semantics, wherein he specified a logical formalism and a method of translating natural language sentences into that formalism so as to apply a logical calculus to it — this came to be known as "Montague Grammar:"

> There is in my opinion no important theoretical difference between natural languages and the artificial languages of logicians; indeed I consider it possible to comprehend the syntax and semantics of both kinds of languages with a single natural and mathematically precise theory.[1]

Thus his goal was to construct a mathematically precise language that could be used to analyze the syntax and semantics of natural language. As we shall later see in the interviews (especially section 5.1), most formal semanticists today continue to acknowledge his influence even today, and view their work as a continuation of his project. Montague can be considered as the originator of this paradigm, the goal of which is to analyze natural langauge by means of mathematical - logical frameworks. Barbara Partee, a notable formal semanticist in the United States, writes about the history of formal semantics:

> Formal semantics has roots in several disciplines, most importantly logic, philosophy, and linguistics. The most important figure in its history was undoubtedly Richard Montague (1930-1971), whose seminal works in this area date from the late 1960's and the beginning of the 1970's ... The development of formal semantics over the past forty and more years has been a story of fruitful interdisciplinary collaboration among linguists, philosophers, logicians, psychologists, and others, and by now formal semantics can be pursued entirely within linguistics as well as in various interdisciplinary settings, including cognitive science, informatics, and computational linguistics. In the U.S. formal semantics is mostly within linguistics departments now, but in parts of Europe (e.g. Amsterdam) it's strongly embedded in the context of logic and philosophy.[2]

In this passage, Partee calls attention to the practice of formal semantics in Amsterdam where it is strongly embedded in the context of logic. This reference is owed to the fact that Amsterdam (together with the context of logic) was

[1] Montague (1970)
[2] Partee (2011), p. 4

of great importance to the development of formal semantics as a discipline. In section 2.2, I discussed the development of logic as a discipline and as a scientific community — beginning as a sub-species of the philosophical and mathematical communities, it slowly established itself as an independent scientific community with its own degrees, journals, conferences and colloquia. The development of formal semantics as a discipline also follows a similar path (although it has not achieved the degree of independence that logic has — one can obtain master's and PhD degrees in logic, but there are no degrees awarded in formal semantics specifically). Moreover, the development of the logic community and the development of the formal semantics community are intimately related to each other — with the ILLC in Amsterdam playing a central role for both developments.

Barbara Partee recently presented a paper as part of a festschrift for Jeroen Groenendijk, Martin Stokhof and Frank Veltman, in which she described the significance of these three figures, and more generally, the significance of the ILLC / Amsterdam community for the development of formal semantics. She writes:

> In 1980 the Amsterdam Colloquium became fully international — that was the first time I attended, and for a decade, that was where I always presented my main new work in formal semantics, because that was where there was the best audience: they could handle formal semantics and they wanted new results and new ideas.[3]

It was only in Amsterdam that the formal semantics community truly flourished with full fledged international colloquia and an eager audience for the latest developments in formal semantics. Partee goes on to point out that this environment was the only one that could have given rise to a serious textbook in logic and formal semantics:

> The Dutch original of the Gamut textbook [a textbook in logic and formal semantics] came out in 1982; that in itself was a very

[3] Partee (2012), p. 187

> special achievement. Only in this environment could one find a textbook that combined logic and formal semantics so thoroughly and beautifully.[4]

Therefore, it was only in Amsterdam that formal semantics was developed truly as a form of *logical practice* (as opposed to the United States where it is in some instances a form of linguistics practice). It is partly for these reasons that I have chosen to focus my case study on the formal semantics community at the ILLC specifically. Also due to the fact that the debate I am about to consider (in the next subsection — subsection 3.1.2) originates within this same community. In this sense, it could be considered an *internal* critique of the discipline.

3.1.2 A debate in theoretical linguistics

The 37th volume (2011) of the journal *Theoretical Linguistics* features a double issue devoted entirely to a debate initiated by Martin Stokhof and Michiel van Lambalgen regarding the scientific status of modern linguistics (especially formal semantics). The volume opens with an article titled "Abstractions and idealizations: the construction of modern linguistics,"[5] in which Stokhof and van Lambalgen raise the question of whether formal semantics could be an example of a *failed discipline*: "These observations give rise to a fundamental question with regard to linguistics as such: Could modern linguistics perhaps be an example of a 'failed discipline'?"[6] The rest of the volume features various responses to this article, and is finally concluded with Stokhof and van Lambalgen's responses to these responses.[7]

Having raised the question of failure, Stokhof and van Lambalgen go on to confirm the suspicion that, for a variety of reasons, formal semantics (and modern linguistics more generally[8]) could indeed be an example of a failed

[4] Partee (2012), p. 187
[5] Stokhof and van Lambalgen (2011a)
[6] Stokhof and van Lambalgen (2011a), p. 3
[7] Stokhof and van Lambalgen (2011b)
[8] It is not sufficiently clear from the article exactly which branches of linguistics are implicated this critique. The authors make it clear that Chomskyan linguistics and formal semantics are implicated, but they also make it clear that computational linguistics is not implicated. For the purposes of this thesis, it is sufficient to note that they think formal semantics in particular is implicated in the critique.

discipline. Crucially, they allege that the models used in formal semantics are far too abstract and idealized, and that therefore they cannot be compared to reality for empirical verification or falsification. The debate centers around the way in which modern linguistics "conceptualizes its central objects of study so as to fit a particular methodology."[9] A major factor that explains the success and prestige of modern linguistics is that it has succeeded to come up with scientific characterizations of its core concepts; this allows researchers to formalize these concepts and embed them into full-fledged models of language. However, Stokhof and van Lambalgen call into question the adequacy of these characterizations. In particular, they allege that the central objects of study of formal semantics (i.e. 'language', 'grammar', 'meaning', etc.) have been deliberately constructed through a process of 'idealization' which does not meet the standards of a rigorous scientific inquiry.

For example, formal semantics replaces the intuitive conception of language (what we encounter in our everyday usage of language) with a mathematical-logical concept of language as "a potentially infinite set of well-formed expressions generated by a finite, or finitely characterisable, set of rules (i.e., a grammar)."[10] Thus the concept of "language," which is one of the primary objects of study for formal semantics is a deliberate construction (and allegedly an idealization) that does not readily correspond to any phenomenon in reality. It is arrived at by starting with an observable real world phenomenon (what we encounter in our everyday usage of language) and idealizing away (i.e. ignoring) what are considered to be "irrelevant" aspects such as human finitude. Another example is the concept of linguistic "competence," which is another crucial object of study for modern linguistics. The concept of competence is an idealized construction that is arrived at by starting with the observable phenomenon of linguistic performance and then idealizing away (i.e. ignoring) the so-called "grammatically irrelevant" conditions such as speech errors, etc.[11] For a final example, consider what is perhaps the central and most important object of investigation for formal semantics — the concept of "propositional meaning" — is also the result of idealization according to Stokhof and van

[9] Stokhof and van Lambalgen (2011a), p. 1
[10] Stokhof and van Lambalgen (2011a), p. 3
[11] Chomsky (1965), p.3

Lambalgen: "the concept of propositional meaning as such is yet another example of a construction that is not so much an abstraction as an idealization."[12] It begins with the concept of natural language meaning as it occurs in ordinary discourse and idealizes away (i.e. ignores) the so-called semantically irrelevant (i.e. pragmatic) conditions such as context.

Stokhof and van Lambalgen allege that these theoretical constructions, although they take some real world phenomena as their point of departure, ignore several qualitative dimensions of the phenomena and end up with an idealized object that is ontologically speaking quite unrelated to the phenomena: "They [idealizations] change the object of study, and one of the consequences of this is that there no longer is an immediate relation between the idealised object and the original, natural phenomenon."[13] They discuss how these theoretical constructions idealize the object of study by once again citing Chomsky's notion of competence:

> What happens here is that competence, regarded as the proper object of study of linguistics, is constructed from what we can observe, i.e., everyday use of language, by stripping it from a number of features, such as memory limitations, mistakes, (communicative) goals, attention shifts, and so on. In other words, Chomsky constructs from observable language use a concept of linguistic competence by simply ignoring a number of its actual, real properties. In that we a new object of study is created, i.e., an object that has an ontological status that differs from that of the original one.[14]

Therefore, the claims made by a theory based on such idealized objects (such as the claims made by formal semantics) can never be empirically verified or falsified, because these claims cannot "actually be compared with observations and the outcomes of experiments."[15] This is because there is a fundamental mismatch between the theory and the empirical observations; the theory makes claims about idealized phenomena, while the observations tell us about real

[12] Stokhof and van Lambalgen (2011a), p. 15
[13] Stokhof and van Lambalgen (2011a), p. 13
[14] Stokhof and van Lambalgen (2011a), p. 11
[15] Stokhof and van Lambalgen (2011a), p. 9

observable phenomena, and there is no way to use the observation to verify or falsify the theory:

> [T]here is a serious lack of empirical validation of the theory about the idealized object. Apart from the fact in the case of linguistics the original phenomenon is hard to fit into an experimental design, there is the problem that, without an independently verified bridging theory, no theory about the idealized object will lead to predictions that can be tested on the original phenomenon (via observation or any any other means).[16]

Therefore, according to Stokhof and van Lambalgen, formal semantics could be seen as an example of a failed discipline, because it does not meet the basic criterion for a rigorous scientific inquiry, i.e. empirical verifiability / falsifiability.

3.2 Research design and methodology

The goal of this case study in formal semantics is to take the debate from *Theoretical Linguistics* as a point of departure and to discuss the question "is formal semantics a failed discipline?" In order to discuss this question I draw primarily on two methodological frameworks: The first is the framework of qualitative research in the social sciences; in particular this case study is structured as an interview study featuring interviews with three different groups of people: (1) Stokhof and van Lambalgen (the critics), (2) formal semanticists (the insiders) and (3) computational linguists (the outsiders). The second is Thomas Kuhn's framework for understanding the history and philosophy of science; in particular formal semantics is treated as a *paradigm* for doing research in linguistics, and the question of the failure of formal semantics becomes the question of a *paradigm shift*.

Strictly within the context of the debate, the goal should be to investigate the the way in which formal semanticists make use of models, the relation between the models and the real world-phenomena that they seek to model, and

[16] Stokhof and van Lambalgen (2011a), p. 18

to determine if the procedure is empirically verifiable or falsifiable. However, focusing solely on this question would lead to a very narrow view of the debate, and would be based on a Popperian conception of scientific inquiry, according to which the success or failure of a scientific endeavor is determined solely by the criterion of verifiability / falsifiability. Although I suspect that this is indeed the conception of science that is implicit in the way that the critics frame the debate, my commitment to philosophy of logical practice demands that I take a broader Kuhnian view and that also take into account the sociological determinants of this question, by means of qualitative social scientific research.

In this section (3.2), I will first (in subsection 3.2.1) explain the way in which I make use of the Kuhnian framework, then (in subsection 3.2.2) I will explain the way in which I make use of the qualitative research framework and the interview methodology. Finally (in subsection 3.2.3) I will describe how this was actually implemented and executed at the operational level of nuts and bolts.

3.2.1 A Kuhnian revolution

> Is mathematics finally going through the Kuhnian revolution that the sciences or, more precisely, the philosophers, historians, sociologists, economists, psychologists of science, ... have been able to deal with ever since the magical year of 1962?[17]

Just as the proponents of the PMP movement are agitating for a Kuhnian revolution in the philosophy of mathematics, I would like to suggest that philosophy of logic is in need of a similar Kuhnian revolution. In advocating PMP as a new approach to philosophy of mathematics, van Kerkhove and van Bendegem pose the following questiosn: "If it is your ambition, as it is ours, to set the Kuhnian revolution in mathematics on its tracks, what to do (to quote a famous political philosopher)?"[18] The revolution ushered by Kuhn in philosophy of science was that of viewing science not as an abstract entity (e.g. as the collection of scientific theories that are detached from social and historical influences), but rather as a body of human practices that are embedded in a determinate social

[17] Van Kerkhove and Van Bendegem (2007), p. vii
[18] Van Kerkhove and Van Bendegem (2007), p. viii

and historical context. According to them, PMP is the way to usher a Kuhnian revolution in philosophy of mathematics.[19] Similarly, I would like to suggest that PLP is the way to usher a Kuhnian revolution in philosophy of logic.

The question under discussion in this case study — "is formal semantics a failed discipline?" — is intimately connected with questions of what it means for something to be a discipline and what it means for a discipline to be a failure? Therefore, Kuhn's framework for understanding the history and philosophy of science in terms of normal science and revolutionary science (i.e. paradigms and paradigm shifts) is exceptionally well suited for this case study. Unfortunately there is very little scholarly precedent for applying these Kuhnian concepts to formal semantics in particular, or to linguistics in general, but I did manage to find one interesting precedent. Writing in 1976, soon after the Chomskyan revolution in linguistics that ushered in the paradigm of generative grammar, Percival questions the applicability of Kuhn's concept of a paradigm to linguistics and concludes negatively:

> This paper examines the applicability to the history of linguistics of Thomas Kuhn's conception of the history of science. It concludes that ... the concept [the concept of a paradigm] cannot be applied either to the history or the present state of linguistics.[20]

However, Percival's reasons for the negative conclusion are disputable. "In the 1970 version of Kuhn's theory, a paradigm is said to have four components: symbolic generalizations, models, values, and exemplars."[21] Percival grants that generative grammar does in fact have all of these four components, but he concludes that it fails to meet the sociological criterion for a paradigm:

> What causes trouble, however, is the sociological dimension of paradigms. Generative grammar does not command uniform assent

[19] Actually, their precise choice of words indicates that they are after a Kuhnian revolution in mathematics, but the context in which they make this point suggests that they are after a Kuhnian revolution in the philosophy of mathematics. I have chosen to adopt the latter interpretation, because I cannot make sense of the former.

[20] Percival (1976), p. 285

[21] Percival (1976), p. 286

> among linguists all over the world: it is not a conceptual framework shared by all the members of the profession.[22]

This leads him to make some radically counter-intuitive claims. As a consolation to linguists, he notes that "if linguistics falls outside the purview of Kuhn's theory, so do some fields [such as evolutionary biology] which are commonly regarded as legitimately scientific "[23] Therefore, I feel that it is not necessary to impose such a strict requirement of uniform worldwide assent for something to count as a paradigm. If we drop this strict requirement, generative grammar could indeed be considered a paradigm, and for similar reasons, so too can formal semantics.

Once we accept that formal semantics can be treated as a paradigm, the question of whether or not it is a failed discipline becomes the question of whether we are at the cusp of a paradigm shift or scientific revolution. According to Kuhn's theory, the development of sciences proceeds in two distinct stages: normal science (wherein researchers are working towards solving puzzles within a paradigm, without calling the paradigm itself into question), and revolutionary science (wherein a paradigm dies and is replaced by a rival paradigm). If formal semantics is a paradigm, then the activity of formal semantics researchers should be seen as an exercise in what Kuhn calls "normal science" and the critics should be seen as attempting to incite some kind of scientific revolution by calling this paradigm into question. In line with this Kuhnian framework, the case study has three main chapters: In the chapter "the critics" I interview researchers who are calling the paradigm of formal semantics into question. In the chapter "the insiders" I interview researchers engaged in the activity of "normal science" within the paradigm of formal semantics. Finally, in the chapter "the outsiders" I interview researchers engaged in the rival paradigm of computational linguistics, and consider the possibility that this rival paradigm may displace formal semantics and usher a paradigm shift.

[22] Percival (1976), p. 289
[23] Percival (1976), p. 291

3.2.2 Qualitative research in the social sciences

This case study draws on the framework of qualitative research in the social sciences; in particular the case study is structured as an interview study featuring interviews with three different groups of people: (1) Stokhof and van Lambalgen (the critics), (2) formal semanticists (the insiders) and (3) computational linguists (the outsiders). Qualitative research is a particular research methodology within the social sciences, and the interview study is one particular form of qualitative research: "In the social sciences today, qualitative interviews are increasingly employed as a research method in their own right, with an expanding methodological literature on how to carry out interview research."[24]

Originally qualitative research in the social sciences was purely a negative research program, in the sense that it was defined as "research that is *not* quantitative research." However, over time, qualitative research developed its own research profile and positive characteristics such as the use of text as empirical data, an interest in the qualitative perspectives of participants, etc. Quantitative research on the other hand is primarily concerned with numbers as empirical data and seeks to make statistical and probabilistic generalizations based on large data sets. Uwe Flick, the editor of the Qualitative Research Kit[25] (which is, together with Silverman's textbook[26], the primary reference manual for this study), describes it as follows:

> The term 'qualitative research' was for a long time used in a distinctive way to describe an alternative to 'quantitative' research ... However, qualitative research has a long history in many disciplines, where social research in general began with approaches that would now be summarized under qualitative research. The longer the development proceeded, the more a profile of what was meant by this term became clear. This profile is no longer defined *ex negativo* - qualitative research is *not* quantitative or *not* standardized or the like — but it is characterized by several features. Thus, qualitative

[24] Kvale (2009), p. 6
[25] Flick (2009b)
[26] Silverman and Marvasti (2009)

> research uses text as empirical material (instead of numbers), starts from the notion of the social construction of realities under study, is interested in the perspectives of participants, in everyday practices and everyday knowledge referring to the issue under study.[27]

Given these goals — especially the interest in the perspectives of participants with respect to everyday practices and everyday knowledge — the interview study seems particularly well suited to meet these goals. The interview study by default is a versatile format for research and it may be used for qualitative as well as quantitative research. The primary difference between qualitative and quantitative research interviews is the format of questions: "Questions may be either closed or open ended. With closed ended questions, a limited number of response alternatives are given; with open ended questions, respondents are free to answer in any way they like."[28] Closed ended questions are suited towards quantitative research (because the responses can be quantified over large data sets) while open ended questions are suited towards qualitative research, because they allow for deeper insight into the perspectives of the subjects. Moreover, there are many interesting things to be learned about a person's world-view that are not always amenable to quantification or to being pigeonholed into closed-ended questions.

> The qualitative interview is a key venue for exploring the ways in which subjects experience and understand their world. It provides a unique access to the lived world of the subjects, who in their own words describe their activities, experiences and opinions.[29]

In this case study, I followed the qualitative interview format with open-ended questions. Researchers were asked to give their perspective on a variety of open-ended questions pertaining to the debate about formal semantics. These interviews were then transcribed, analyzed and reported in accordance with the standards of qualitative research. The implementation and execution of the research design is elaborated in the following section (3.2.3).

[27] Flick (2009a), p. 1 - 2
[28] Cozby and Bates (2011), p. 134
[29] Kvale (2009), p. 9

3.2.3 Implementation and execution

In this section I will explain the implementation and execution of this qualitative interview study in five steps: (1) Planning the research project. (2) Considering the ethical issues involved. (3) Conducting the interviews. (4) Transcribing the interviews. (5) Analyzing and reporting the interviews.

(1) Planning. Having decided upon the context, goals and methodology, the next step was to select interview subjects and design an interview script that would further this agenda. The selection of interview subjects was rather straightforward. Martin Stokhof and Michiel van Lambalgen were selected as the critics because they are the originators of the theoretical linguistics debate. Jeroen Groenendijk, Floris Roelofsen, Frank Veltman and Katrin Schultz were selected as representative members of the formal semantics community at the ILLC. Khalil Sima'an, Ivan Titov, Philip Schulz and Remko Scha were chosen as representative members of the computational linguistics community at the ILLC. All of the interview subjects (with one exception) are senior staff members at the ILLC in their respective fields.

I then prepared an interview script, featuring a sequence of themes to be covered: Introduction, professional background and motivation, the nature and goal of formal semantics, the role of modeling in formal semantics, the success or failure of formal semantics, conclusion. The interview script also contained many pre-prepared questions for each theme. Some examples: "What is formal semantics? How would you describe it to a university student who has never heard of the field?", "What are the kinds of problems that researchers try to solve in formal semantics?", "What are the standards for success and failure in modeling?", "Do you think that formal semantics as a discipline has been successful thus far in achieving its goals?", "Are you familiar with the criticisms made by Martin Stokhof and Michiel van Lambalgen? If so, how would you react to that?"

(2) Ethical concerns. The next step was to address the fact that such a study raises some important ethical and epistemological concerns:

> Ethical guidelines for social science research emphasize the need to obtain the subjects' informed consent to participate in the study,

> to secure the confidentiality of the subjects, to consider the consequences for the subjects of participation in the research project and to be attentive to the researcher's role in the study.[30]

In order to address these ethical concerns, I sent a detailed email to each potential interview subject explaining the nature of my research and requesting their consent to be interviewed and recorded for the purposes of my thesis research. Then, after I conducted and transcribed the interviews, I emailed the transcriptions back to the interview subjects to request consent for me to publish them as an appendix to my thesis. With the consent of the interviewees, the interviews are appended to this thesis semi-anonymously — in the sense that it is not possible to identify any particular interview with any particular interview (furthermore, everywhere in the thesis I have used the pronoun "he" as a gender-neutral pronoun for referring to any particular interviewee).

(3) Conducting the interviews. Before conducting the interviews, I conducted a few test interviews with master's and PhD students in formal semantics at the ILLC, in order to test the interview script to see whether some questions resulted in confusion and needed to be reworded, and also to get a feeling for the length of the interview following the script. Then, during the months of March, April and May 2014, I met each interview subject (typically in their office, but sometimes at a public venue such as a cafe) and conducted the interviews in person with the use of an audio recording device.

> A qualitative interview is usually semi-structured; it has a sequence of themes to be covered, as well as some prepared questions. Yet at the same time there is openness to changes of sequence and question forms in order to follow up the answers given and the stories told by the interviewees.[31]

Subjects were briefed and debriefed before and after the interviews. During the interviews, I followed the interview script — I went through each of the themes, asking the list of prepared questions, while also following up on the particularities of individual responses to the questions.

[30] Kvale (2009), p. 31
[31] Kvale (2009), p. 65

(4) Transcribing the interviews. After conducting the interviews, the next step was to transcribe the interviews from the audio files. I did this manually by listening to the audio files and typing up the text, but there are many procedural questions that such an exercise gives rise to:

> Should the statements be transcribed verbatim and word by word, retaining frequent repetitions, noting 'mh'-s and the like, or should the interview be transformed into a more formal, written style? ... There are no correct, standard answers to such questions; the answers will depend on the intended use of the transcript, for example, whether for a detailed linguistic conversational analysis or for reporting the subject's accounts in a readable public story.[32]

Since detailed linguistic conversational analysis was not the intended use of these transcripts, I chose the latter option of transcribing the interviews in a manner that results in a readable text. I have not retained frequent repetitions, 'mh'-s, pauses, and sentence fragments and phrases that don't lead anywhere. The grammar and syntax has also been cleaned up with the addition of punctuation such as commas, dashes (—), parentheses, etc. to make the text more readable.

(5) Analyzing and reporting the interviews. The final step was to analyze and report the interviews: "Analysis as bricolage and as theoretical reflection goes beyond following specific techniques or approaches to interview analysis and draws in a variety of techniques and theoretical concepts."[33] Given the goals of my research, I structured my analysis as individual chapters discussing the question "is formal semantics a failed discipline" from the point of view of critics, insiders and outsiders, while drawing upon a broadly Kuhnian perspective.

In terms of reporting, the standard practice is to render the interview quotes in a contextualized and readable style to the degree possible.[34] I decided to publish the interviews in full — with consent from the subjects — as an appendix to the thesis (for maximum contextualization), and then to use some relevant quotes in the main chapters to make a readable argument. Such a

[32] Kvale (2009), p. 95
[33] Kvale (2009), p. 119
[34] Kvale (2009), pp. 129 - 135

style of reporting also has the added benefit that the appendices can serve as data for future studies of the same or related subjects.

3.3 Disclaimers

In this section I issue some disclaimers about the research design and methodology of this thesis. In particular, this section is structured as a list of some salient (hypothetical) objections that might be raised about the research design and methodology of my thesis, together with my responses to these objections

1. *Objection.* The case study is too narrow in scope. First of all it is limited to one particular form of logical practice (formal semantics) at one particular institute (Institute for Logic, Language and Computation). Therefore it is not possible on the basis of such a case study to make any interesting generalizations about logical practice. Moreover, the case study is quantitatively limited to ten interview subjects (all of whom are from the same institute), therefore it is not even possible to make generalizations that extend to the field of formal semantics in general.

   Reply. This is the nature of qualitative research. The methodology of qualitative research allows us to take a delve deeply into one particular community and its corresponding world-view (as opposed to quantitative research, which seeks to cover a lot more ground with less depth). Moreover, this particular community is of central significance to the development of logic and of formal semantics as scientific disciplines over the past few decades. Not only did it champion the cause of logic as a discipline with it's own research program, conferences and schools, but it was also responsible for some of the most significant advances in formal semantics since Montague. See section 2.2 and section 3.1.1 for more details.

2. *Objection.* The Kuhnian framework is not being used in the way that Thomas Kuhn intended it to be used. Kuhn himself did not consider the concept of paradigm as appropriate for the social sciences or any other discipline that is not a natural science. He explains in his preface to

The Structure of Scientific Revolutions that he concocted the concept of paradigm precisely in order to distinguish the social from the natural sciences [35]. In Kuhn's view, the existence of a single reigning paradigm is characteristic of the natural sciences, while philosophy and much of social science were characterized by a "tradition of claims, counterclaims, and debates over fundamentals."[36] Therefore it is inappropriate to apply Kuhn's framework to linguistics.

Reply. I believe it is not necessary to adhere to the overly strict use of the framework that Kuhn prescribes. The Kuhnian framework of analysis has taken on a life of its own (beyond what Kuhn intended), and indeed there are many instances where scholars have relaxed these requirements to make use of the Kuhnian framework to analyze disciplines other than the natural sciences.[37]

3. *Objection.* The analysis of verbal interviews does not yield philosophically interesting results. For instance, it is entirely possible that the considered and reflective views of the interview subjects would be different than the answers given in a verbal interview.

   Reply. Once again (like the first disclaimer) this is not a problem, because it is part of the goal of qualitative research. It aims to get insight into the perspectives of the subjects in terms of "everyday practices and everyday knowledge referring to the issue under study."[38] In terms of this goal, interviewing is a valid strategy, since it directly reveals the everyday knowledge and everyday practices as opposed to the kind of response one might get if the subjects were given time to prepare their response in advance.

[35] Kuhn (1970b), p. x
[36] Kuhn (1970a), p. 6
[37] For a fairly comprehensive review, refer to Heyl (1975)
[38] Flick (2009a), p. 1 - 2

Chapter 4

The critics: Stokhof and van Lambalgen

The collaboration between Stokhof and van Lambalgen (henceforth "the critics"), at first glance, seems rather unlikely. One is a classically trained formal semanticist with a background in philosophy, logic, linguistics and the philosophy of language. Motivated primarily by a deep fascination with language, he sought to use the formal methods of mathematical logic for a philosophical and scientific analysis of the phenomena of language and linguistic behavior. The other critic has, as he puts it, a rather "checkered"[1] academic background. He found his way into formal semantics by way of working in the fields of probability theory, mathematical logic and artificial intelligence. As an artificial intelligence researcher, he was studying mathematical theories of vision, and he was "trying to link up these theories with the semantics of perception reports."[2] His primary motivation was not philosophy or linguistics, but rather it was to better understand how the brain works; in particular, he wanted to understand the cognitive and neural processes associated with language processing. Yet, despite these differences, they shared a common interest in formal semantics for a long time, and now they have come to share some fundamental criticisms of the discipline. For various reasons, they both believe that formal semantics as it is practiced today is able to live up to neither its scientific nor

[1] Interview A2, Appendix A
[2] Interview A2, Appendix A

its philosophical aspirations. Their dissatisfaction with formal semantics has led each of them in different directions. One has largely given up his scientific aspirations in linguistics and is now a philosopher, while the other has gone in the direction of more empirical research in the cognitive sciences.

In this chapter, I will first (in section 4.1) discuss the nature and the goals of formal semantics as a discipline as it is understood by the critics, then (in section 4.2) I will explain their conception of role of modeling in formal semantics and the problems that it gives rise to, and then (in section 4.3) I will clarify their stance on the successes and failures of formal semantics from a scientific as well as a sociological point of view.

4.1 On the nature and goals of formal semantics

Formal semantics is a form of logical practice that seeks to apply logical techniques (or broadly speaking — formal methods) to the analysis of natural language. At this broad level of description both of the critics as well as the formal semanticists themselves are all in agreement, but as usual the devil is in the detail, and in particular there is disagreement among all parties regarding what is (or should be) meant by the term "natural language." According to the critics, natural language is simply what one encounters in ordinary, everyday uses of language in the observable form of conversations, texts, recordings, etc. (in the jargon of modern linguistics, this is what is known as "linguistic performance" or "linguistic behavior"). Therefore, according to one critic, observable linguistic behavior is precisely what formal semantics should be trying to analyze. According to the other critic, formal semantics should also be trying to analyze what goes on in one's brain when one is engaged in linguistic behavior (i.e. the observable cognitive and neural processes associated with language processing).

However, the formal semanticists have a rather different idea of what formal semantics should be about. In reality (as demonstrated below) formal semantics is concerned neither with actual linguistic behavior nor with its cognitive / neural counterpart, but it is instead concerned with natural language meaning, which is an abstract, structural feature of natural language. This notion of

meaning is closely related to a conception of natural language as "linguistic competence," which is an unobservable abstraction that is constructed from observable linguistic behavior. The critics allege that this abstraction is not really an abstraction, but in fact an idealization which has no basis in reality. When asked about the nature and goals of the enterprise of formal semantics, every single one of the formal semanticists that were interviewed cast their response in terms of meaning (as opposed to linguistic behavior or its cognitive / neural counterpart). Here I cite four different excerpts from the four different interviews in which the formal semanticists clarify the nature and goals of their enterprise (emphasis added):

> I would say it's a theory or an approach to meaning where you use the tools from logic (and build logical tools — it's not only using tools, but also building new logical tools if you need them) to get a detailed and formal explanation of the workings of *meaning*.[3]

> Formal semantics is concerned with *meaning* of expressions in natural languages in general ... the focus in this field is about how the meaning of complex expressions —– like sentences or parts of sentences, but expressions that contain multiple words in some complex configuration —– how the meaning of those complex expressions can be computed from the meaning of the simpler atoms.[4]

> [The goal of formal semantics is] to understand how *meaning* arises —– how expressions of language get their meaning. It's a kind of a miracle in a sense how meaning is built up, for instance, of larger parts by combining in some way the meanings of smaller parts and how that goes.[5]

> For me it's a lot about interpretation ... it's about *meaning* – — how we understand things and how we code our intentions in language.[6]

[3]Interview B1, Appendix B
[4]Interview B2, Appendix B
[5]Interview B3, Appendix B
[6]Interview B4, Appendix B

Although the critics agree that this is a good characterization of what formal semantics actually is, the critics do not agree that this is a good characterization of what it should be or what it could be. As noted, one of the critics asserts strongly that formal semantics should be linked to cognitive science:

> For me formal semantics is strongly linked to cognitive science, and ultimately the most interesting work done in that area is the work that gives some insight into language processing (both producing and comprehension).[7]

The other critic asserts that formal semantics should be about observable linguistic behavior (linguistic performance) rather than the notions of meaning or linguistic competence:

> [L]inguistics should be about linguistic behavior; actually I think it should, but it [modern linguistics] is actually not about that — it's doing something else.[8]

What is that "something else" that formal semantics is doing? The critic goes on to point out that formal semantics is constructing natural language as an idealization:

> Natural language [according to formal semantics] is then a somewhat idealized concept. The question of what natural language is is actually philosophically a very interesting question, because what we observe is linguistic behavior and the products of that behavior — e.g. what we are recording now and what is written down. That is what natural language is, but from the point of view of modern linguistics it's a more abstract type of object. That nevertheless has an important cognitive counterpart — namely the ability that is characteristic of competent language users [linguistic competence]. I think most formal semanticists are still knowingly or unknowingly subscribing to that conception of what it is that they are modeling.[9]

[7] Interview A2, Appendix A
[8] Interview A1, Appendix A
[9] Interview A1, Appendix A

Already we see that the dispute is of a fundamental nature, because it even extends to something as basic as the nature and goals of this discipline. Formal semantics as it is practiced is about abstract, structural features of natural language (such as meaning). The critics allege that these so-called abstractions are in fact idealizations that cannot tell us anything interesting about the real-world phenomenon of natural language, and they insist that formal semantics should be concerned with the observable phenomenon of linguistic behavior. Yet formal semanticists do maintain that the abstraction they wish to study is very useful, and that it does tell us something about the real-world phenomenon of natural language[10].

4.2 On the role of modeling in formal semantics

Logical practice in general presupposes and depends crucially on the processes of modeling and formalization; in particular, the logician will produce a formal model of some real world phenomenon in order to represent it and to reason about it. "This transformation of the real world data into formal representations — modeling — gives us a representation of the phenomena we want to understand in a formal language that we can manipulate and apply the deductive method to."[11] Formal semantics, being a particular form of logical practice, is no exception to this general rule of thumb. In this case the relevant real world phenomenon to be modeled is natural language. The formal semanticist will produce a formal model of natural language (or at least a model of some particular region of natural language) in order to represent it and to reason about it by using the deductive method and other logical calculi. Already in the above section (section 4.1), we saw that the goals of formal semantics were characterized in terms of "modeling natural language meaning"[12]. In fact, this quest for modeling is built into the very heart of the formal semantics enterprise. As one critic points out:

> There are many ways of describing what semantics does (especially the role of formal languages in the enterprise) where the

[10] The last point is discussed in more detail in chapter 5, especially section 5.1
[11] Löwe (2012)
[12] Interview A1, Appendix A

> formal languages themselves are described as models, and I think that that's certainly inherent in the Montague paradigm. ... Basically you look at the formal language with its interpretation as modeling the properties of the natural language expressions in the relevant respects. So I think in that sense much of formal semantics is actually modeling...[13]

This goal is agreed upon by not only the critics, but also the formal semanticists themselves (this is discussed in more detail in the following chapter — especially section 5.2). Given that formal semantics is concerned with modeling, and given that modeling is a process that replaces one thing with another — in particular, it replaces the real world phenomenon of natural language with its formal representation in a formal language — we must take care to make sure that the model stays close to the modeled object.[14]

But how do we make sure that our model is actually a good model for the thing we are trying to model? There must be some external benchmark against which we can test our formal representations. So what are the standards for success and failure in this enterprise of modeling? Traditionally in the humanities, the answer to this question has been that intuition (specifically *expert intuition*) is the benchmark against which models and theories are to be measured, whereas in scientific disciplines empirical data serves as the benchmark.[15]

Does formal semantics follow in the footsteps of the humanities and set intuition as its benchmark? Or does it follow in the footsteps of the natural sciences and set empirical data as its benchmark? As we shall see, there is in fact no definitive answer to this question, and it seems like at some level formal semanticists want to have their cake and eat it too — in the sense that they want it to be a simultaneously philosophical as well as a scientific inquiry. I would like to suggest that this is the root cause of some fundamental tensions in the discipline. This point will be discussed in further detail in the next chapter. For now, suffice it to say that it is not uncommon for empirical verification to be cited as the benchmark:

[13] Interview A1, Appendix A
[14] Löwe (2012)
[15] Löwe (2012)

> It should get the facts right that it wants to account for. ... If it gets all those facts right (whatever those facts are taken to be) then it is considered successful.[16]

> The models should agree with what is known about the process it is modeling, and with other aspects of cognition, and it should give predictions about as wide a range of data types as is appropriate. By this I mean not just the so-called appeals to intuition, but also corpus research or neural imaging data.[17]

If the benchmark for success or failure in modeling is that the models should agree with the facts and the data, this immediately leads to the question of what are the data? According to one critic, there are a wide range of data types that should be taken into consideration:

> let's run through the variety of data that you have when you study language comprehension. (1) There's eye tracking which traces the movement of the eyes while reading, and that can give information in case backtracking occurs. (2) Then there are ERP data which is a kind of average form of EEGs taken when a subject is processing linguistic material. (3) Then there is what is properly behavioral data —– namely people's answers to verbal questions. I haven't exhausted all the data types...[18]

The data types mentioned here pertain to linguistic behavior and to the brain processes associated with linguistic behavior. But the critics allege that it is quite impossible for such data to serve as a benchmark for the theories of formal semantics, because there is a categorical mismatch between the kind of data we have and what the theory or model wants to talk about. Since the data are about observable features of linguistic behavior — while the models and theories of formal semantics are about abstract, structural features of language — it is unclear what the relation is between the two.

[16] Interview A1, Appendix A
[17] Interview A2, Appendix A
[18] Interview A2, Appendix A

What adds a final intriguing twist to this narrative is that formal semantics is — according to the critics, but also according to the formal semanticists (as we shall see in the next chapter in section 5.2) — currently going through some kind of identity crisis with respect to the questions of modeling and data:

> It used to be simply intuitions and then people started to incorporate data that they got from corpora and data that they got from doing experiments (asking people about judgments and stuff like that), so now it's very much mixed.[19]

> Since they are increasingly using that [data from linguistic behavior] as their data — the move away from intuitions to more behavioral data creates a tension between these two conceptions, and that is something that is very much happening now. The outcomes of that are still sort of undecided.[20]

I would like to suggest that this is one of the fundamental tensions that lies at the heart of formal semantics. Is formal semantics about observable linguistic behavior (linguistic performance) or is formal semantics about some hypothetical ideal form of linguistic behavior (linguistic competence)? The critics' point is that if it is about linguistic performance, then it should discard it's idealizations and instead adopt a rigorously scientific methodology. Whereas if it is about linguistic competence, then it should acknowledge that this is a philosophical question and it should give up its rigorously scientific aspirations. Should formal semantics try to emulate philosophy and use intuition as the benchmark for success and failure, or should formal semantics try to emulate the natural sciences and use empirical data as the benchmark? As we shall see in the following chapter, these tensions remain very much unresolved and they have a polarizing effect on the discipline.

[19]Interview A1, Appendix A
[20]Interview A1, Appendix A

4.3 On the success or failure of formal semantics

If formal semantics is alleged to be a failed discipline, the allegation must rest on some particular notion of "failure" (as in "failed discipline"). This requires a goal or a definition of success together with the observation that this goal was not met. In this case, what is the notion of success or failure that is in play? The allegation of failure would be frivolous if it depended on some arbitrary goal. For example, I might argue that formal semantics is a failure because no formal semanticist has ever won a Nobel Prize. Such an argument would not be interesting, because it depends on an arbitrary goal or standard of success (winning a Nobel Prize) that I have stipulated. What makes the critics' argument interesting is that the allege formal semantics to be a failure *by it's own standards.*

In this sense, there are two standards of success or failure that are worth considering. One is a scientific standard — given that formal semantics would like to satisfy the standard criterion for being a scientific discipline, the critics attribute to it the implicit goal of 'being scientific' which involves the possibility of empirical validation / falsification. The other is a sociological standard — given that formal semantics would like to satisfy the standard criterion for being a discipline (scientific or not), the critics attribute to it the implicit goal of 'being a discipline' which involves having a thriving community of students and researchers. Although their paper only focuses on the first goal (the scientific goal), it is clear from the interviews that they allege formal semantics to be a failure with respect to both goals (scientific and sociological).

Firstly, the critics allege that since the central objects of study of formal semantics are idealizations, it does not meet the scientific criterion of empirical verifiability / falsifiability:

> [Formal semantics] is not able to provide a good account of how what it does deal with is connected to what it should be dealing with ... In an idealization you simply forget about that [various empirical features], and that means that you need something else to translate your results back into results that are relevant for what

> you idealized away from, but linguistics doesn't contain its own bridging theory. In that sense, yeah, it's a failed discipline.[21]

Here the critic argues that formal semantics has been a failure by it's own standards of being scientific (because the critic attributes to formal semantics the goal of being scientific). We will see what the formal semanticists themselves have to say about this point in the following chapter (in section 5.2). Besides the scientific standard, since this is an exercise in philosophy of logical practice, it is equally important for us to pay attention to the sociological dimensions of the success and failure of this logical practice. Interestingly, one of the critics brings up this discussion in a rather off-hand manner:

> I think there is no uniform set of goals for formal semantics, so it's very hard to speak of success or failure. You may have attended the discussion meeting at the Amsterdam colloquium, where Angelika Kratzer said that the success of formal semantics can be measured by the amount of young researchers that work in the area. That is not exactly an intellectual criterion (to put it mildly). I think that the goals of formal semantics are unclear, and as far as I can see, too narrow in scope. There is some testing of the model against data, but it is usually all of the intuitive kind.[22]

The same critic then goes on to say that he thinks most formal semanticsts would disagree with the assessment that formal sematnics is failed discipline by these standards: "I think they would mostly disagree and take Angelika Kratzer's side. ... that as long as formal semantics keeps attracting intelligent and enthusiastic young people the subject is thriving.[23]" Here the critic is attributes to formal semantics the goal of having a thriving community of researchers and students. Although the formal semanticists think that they have been successful with respect to this sociological goal (according to this critic), the critic goes on to point out that this may not actually be the case:

[21] Interview A1, Appendix A
[22] Interview A2, Appendix A
[23] Interview A2, Appendix A

> There are I think sociological processes at work. The subject started out in the 1960s in a rather passionate manner where the clear goal was defined by Montague grammar. As time wore on more goals appeared and disagreements arose, but subjects very seldom simply die out. Just to give you an example of how we filled vacancies here — you know that Martin [Stokhof], Jeroen [Groenendijk] and Frank [Veltman] retire. Frank has already retired and Martin will follow a year from now. It could have been decided to reallocate the open positions to some different fields, but it was felt that the tradition of formal semantics in Amsterdam was so important that the positions need to be allocated to formal semanticists. And it's those type of decisions that keep a subject going, but it doesn't necessarily lead to intellectual coherence.[24]

Part of his point is essentially that formal semantics is able to sustain a thriving community of students and researchers, but that it is only able to do so in an artificial (as opposed to organic) manner. He argues that this will eventually threaten the intellectual coherence of the discipline, at which point it will no longer be able to meet its sociological standard for success. Thus the critics argue that formal semantics has been a failure by its own standards. Formal semantics *wants* to be a scientific discipline — it wants to be scientific and it wants to have a thriving community of students and researchers — but it is failing (or will fail) at both of these goals.

4.4 Conclusions

Logical practice in general presupposes and depends crucially on the processes of modeling and formalization; in particular, the logician will produce a formal model of some real world phenomenon in order to represent it and to reason about it. Formal semantics, being a particular form of logical practice, is no exception to this general rule of thumb. In this case the relevant real world phenomenon to be modeled is natural language. The formal semanticist will

[24] Interview A2, Appendix A

produce a formal model of natural language (or at least a model of some particular region of natural language) in order to represent it and to reason about it. However, there is a dispute concerning what is meant by "natural language." Formal semanticists seem to be interested primarily in abstract, structural features of language such as the concept of propositional meaning. The critics on the other hand question the usefulness of such abstractions, urging instead that they focus on the observable phenomena associated with linguistic behavior.

Given that formal semantics is concerned with modeling, and given that modeling is a process that replaces one thing with another — in particular, it replaces the real world phenomenon of natural language with its formal representation in a formal language — care must be taken to make sure that the model stays close to the modeled object. This requires some external benchmark against which the models can be measured. Typically in the humanities expert intuition is used as the benchmark, whereas in the natural sciences empirical data is used as the benchmark. What about formal semantics? Here formal semantics seems torn between its philosophical and scientific aspirations: Traditionally formal semantics was not very interested in linguistic performance or real-world linguistic behavior, therefore it was happy to rely on intuition rather than any empirical data; however, this "armchair" methodology is of late seen to be in conflict with its scientific aspirations, and we are now beginning to see a shift in formal semantics away from intuitions towards empirical data. This is creating an identity crisis for the discipline and results in some tension between the philosophical aspirations of the discipline (which may be judged by intuition) and the rigorously scientific aspirations of the discipline (which demands empirical data for judgment). What about the question of failure? Given that formal semantics would like to satisfy the standard criterion for being a scientific discipline, the critics attribute to it the implicit goal of 'being scientific' which involves the possibility of empirical validation / falsification.

Additionally, we see that there are sociological processes at work, and that they are equally important to the future of this debate. Given that formal semantics would like to satisfy the standard criterion for being a discipline (scientific or not), the critics attribute to it the implicit goal of 'being a discipline' which involves having a thriving community of students and researchers. On the one hand we see that some people (e.g. Angelika Kratzer) think that formal

semantics continues to be successful in this regard, but we also see (as one critic pointed out) that this might be an artificial (as opposed to organic) phenomenon which might threaten the intellectual coherence of the field. While this all leaves the future of formal semantics very much uncertain, I think that we can look to the careers of the two critics as an example of what the future might hold. One critic notes:

> [We] are engaged in a foundational critique [of formal semantics] which also affects the way I used to do formal semantics, but that critique has not sufficiently crystallized into a new set of practices ... we have not yet hit upon the proper way to study this scientifically ... until then I keep to the set of goals that I had.[25]

This critic has retreated from formal semantics back to his previous research agenda, which is a rigorously scientific and empirical investigation of the brain processes involved with linguistic behavior (one might locate such a research project in the field of cognitive linguistics). The other critic however, has gone in an entirely different direction. He has retreated from formal semantics not to any other rigorously scientific inquiry, but to philosophy proper — he now spends his time researching the philosophical musings of Wittgenstein and Heidegger (among others) on the phenomenon of language.

I would like to suggest that this is the critics' solution to the identity crisis of formal semantics — between its philosophical and its scientific aspirations — the way to solve the problem is for the formal semanticist to abandon his philosophical aspirations and enter into a rigorously scientific discipline (as one critic has done) or to abandon his scientific aspirations and enter into philosophy proper (as another critic has done). In the following chapters, we also see some evidence a similar trend among the other (former and current) formal semanticists.

[25] Interview A2, Appendix A

Chapter 5

The insiders: formal semanticists

> *In so far as he is engaged in normal science, the research worker is a solver of puzzles, not a tester of paradigms. Though he may, during the search for a particular puzzle's solution, try out a number of alternative approaches, rejecting those that fail to yield the desired result, he is not testing the* paradigm *when he does so.* [1]

The formal semanticists interviewed in this study all acknowledge the criticisms made by the critics, yet their responses to the criticism are surprisingly dismissive of it. As this chapter reveals, the concerns raised by the critics — especially regarding the idealism of the models and the tension between the philosophical and scientific aspirations of the discipline remain unresolved. The formal semanticists maintain that their discipline is as rigorously scientific as it gets. Although they recognize that there are serious problems with respect to the empirical validation of their models, they explain this problem away by delegating this empirical work to some point in future or to some other group of researchers. They remain optimistic about meeting the goal of being a scientifically valid discipline; however, some of them are quite pessimistic about meeting the goal of sustaining a thriving scientific community for the discipline (especially within the larger community of linguists).

[1] Kuhn (1970b), p. 144

In this chapter I first (in section 5.1) present the direct responses given by formal semanticists to the criticism of the critics, and I make an attempt at a roughly Kuhnian explanation for this puzzling response. Then (in section 5.2) I explore their view of the problems relating to modeling and data in formal semantics, revealing some fundamental tensions between the philosophical and scientific aspirations of the discipline. Then (in section 5.3) I discuss the success or failure of formal semantics in the scientific dimension as well as the sociological dimension.

5.1 On the criticism of formal semantics: reaction to the critics

Having explored the criticism of formal semantics in the previous chapter, I would like to begin this chapter by allowing the formal semanticists themselves to directly respond / react to that criticism. In my interviews with the formal semanticists, towards the end of each interview I brought up the subject of the criticism of formal semantics made by Stokhof and van Lambalgen. I asked each subject if they were familiar with the criticism, and if so how they would react to it. In each case the answer was almost the same: they were familiar with the criticism, but it didn't seem to bother them. Each subject agreed that there was some truth to the criticism, but then went on to say that it didn't have any impact on their practice of formal semantics. The following excerpt from interview B1 is in fact paradigmatic of the general response.

> I: Are you familiar with this criticism?
>
> S: Yes, yes, I'm familiar with that.
>
> I: And how would you react to that?
>
> S: To be honest, I don't care very much.
>
> I: Why not?
>
> S: Well ... I mean, of course, it's true. You are idealizing, and maybe you are idealizing in a way that's not good. That can always happen, but ... I think it's just a matter of obvious fact

> that we have learned interesting things about language by these logical tools, and that's all I want — to learn more.[2]

The other formal semanticists interviewed also gave very similar responses. Here is one that is almost identical:

> I: Are you familiar with this criticism?
> S: I'm certainly familiar.
> I: And how would you react to it?
> S: I disagree.
> I: Could you elaborate?
> S: It's such a scholastic kind of criticism. So yes, I think that formal semantics uses idealizations the same way that mechanics uses idealizations and that doesn't matter. ... So there's nothing wrong with it from a methodological point of view ... I disagree [with the critics] ... I certainly think that the theories as we have them do have some cognitive value...[3]

And here is a third response, following in a very similar vein:

> I: I'm wondering if you are familiar with this criticism, and if so how you would react to it?
> S: Yes, but I don't have something very sensible to say about that I think. Somehow it just doesn't seem to affect the way that I see myself working.
> I: Ok. Why do you think that is?
> S: It just doesn't ... I think they make good and helpful distinctions at a meta-level looking at what's happening in the field, but I don't necessarily see it as a criticism, but just as a reflection.[4]

In each case we see that the formal semanticist acknowledges the criticism, and acknowledges the critics' point about idealization, but then they refuse to

[2]Interview B1, Appendix B
[3]Interview B3, Appendix B
[4]Interview B2, Appendix B

accept that this should have any impact on their research. It's rather puzzling at first that the formal semanticists don't seem to care so much about a criticism that claims to invalidate the enterprise of formal semantics. However, this fact becomes less puzzling when we recall Kuhn's point (quoted at the beginning of this chapter) that these researchers are engaged in the activity of normal science within the paradigm of formal semantics — in this capacity, they are what Kuhn calls "solvers of puzzles" and not "testers of paradigms." What the critics are doing is calling the entire paradigm into question. What we learn from Kuhn is that this only happens during times of crisis, when we are at the cusp of a scientific revolution — it does not happen during the course of normal science. "Popper famously complained that psychoanalysis could not be scientific because it resists falsification. Kuhn's account argues that resisting falsification is precisely what every disciplinary matrix [disciplinary matrix = paradigm] in science does."[5] Testing the paradigm, by comparison to reality or by any other means, is usually not part of the activity of normal science, and is actively resisted by normal science: "As we shall shortly see, when I turn from the experimental to the theoretical problems of normal science, there are seldom many areas in which a scientific theory, particularly if it is cast in a predominantly mathematical form, can be directly compared with nature."[6]

Kuhn tells us that normal science means "research firmly based upon one or more past scientific achievements, achievements that some particular scientific community acknowledges for a time as supplying the foundation for its further practice"[7] It is precisely these achievements that Kuhn calls paradigms:

> Their achievement was sufficiently unprecedented to attract an enduring group of adherents away from competing modes of scientific activity. Simultaneously, it was sufficiently open ended to leave all sorts of problems for the redefined group of practitioners to solve. Achievements that share tees two characteristics I shall henceforth refer to as paradigms.[8]

[5] Bird (2013)
[6] Kuhn (1970b), p. 27
[7] Kuhn (1970b), p. 10
[8] Kuhn (1970b), p. 10

In this sense formal semantics is certainly a paradigm that originates with the work of Montague. Just as Chomsky can be considered the originator of the generative grammar paradigm (discussed in section 3.2.2), Montague is the originator of the formal semantics paradigm (discussed in section 3.1.1) This is evident in the interviews with the formal semanticists, who continue to acknowledge his influence and continue to cite Montague's work as exemplary for the field. Here are some excerpts:

> The big issues are "what is the right framework?" and then the normal science is to develop theories within one of these framework. For a long time Montague grammar was the framework.[9]

> Montague's work in the early 1970s — using typed lambda calculi to provide a framework in which the meaning of sentences can really be derived compositionally from the meaning of the words all the way up to the meaning of complex sentences. That's probably seen as *the* major breakthrough in the field.[10]

> The subject started out in the 1960s in a rather passionate manner where the clear goal was defined by Montague grammar.[11]

Based on this analysis I would like to suggest that formal semantics researchers can be seen as continuing to work on normal science within a particular paradigm in a scientific community that acknowledges Montague grammar and its goals as supplying the foundations for the discipline. They do not question the paradigm, because they do not believe that the discipline is going through a crisis. What this ultimately reveals is that the critics have not managed to convince them that the discipline is going through a crisis, and therefore the formal semanticists are reluctant to call the paradigm into question. Of course, we also learn from Kuhn that in the history of science it is very rare for critics to convince mainstream practitioners of a discipline to challenge their paradigm — For Kuhn, such a crisis or revolution is almost always driven by a competing paradigm (rather than mere criticism).

[9] Interview B3, Appendix B
[10] Interview B2, Appendix B
[11] Interview A2, Appendix A

5.2 On the role of modeling in formal semantics

In the previous chapter we noted that the criticism of the critics is due in large measure to the the idealism of the models in formal semantics. Moreover, there were concerns regarding the standards (or lack thereof) for success and failure of the models — in particular, it was pointed out that there seemed to be an unresolved tension between the philosophical and the scientific aspirations of the discipline. Originally the discipline was more closely aligned with its philosophical aspirations, and intuition was used as the benchmark. But now the discipline wants to become more closely aligned with its scientific aspirations, and therefore seeks to use empirical data as the benchmark, although there is no obvious relation between the models and the data. In the above section, I quoted three out of four responses to the critics. The fourth response (quoted here) touches upon the issue that we are now discussing:

> I: Are you familiar with this criticism?
>
> S: A little bit. I haven't read these papers.
>
> I: Then I was going to ask if you have any reaction to it?
>
> S: I think there is something right . . . [in] their idea of the idealism of the data. That's true, but this is also linked to the general development of the field —– to getting closer to the data and really doing serious data studies before you start modeling.[12]

Here the subject agrees (provisionally) with the criticism, but suggests that the criticism will become more and more irrelevant with the "development of the field" as it gets "closer to the data." Thus we see this idea of development as getting closer to the data. This suggests that the contrast is between an earlier point in time and the present. However, it turns out that even in the present day the contradiction persists. First of all, we see that formal semanticists distinguish themselves from linguists by the criteria that the linguists are primarily interested in empirical data about linguistic behavior, but formal semanticsts are less interested in the data and more interested in the conceptual work of building formal frameworks. Here are two excerpts from semanticists:

[12]Interview B4, Appendix B

> [I would describe myself] as a semanticist, and I think that's something in between being a logician and being a linguist. If I would be a full blown logician I would also do things like making interesting proofs concerning systems, and if I were a linguist I would be very much interested in empirical data, but neither the proofs nor the empirical data concern me very much. I see myself as building conceptual frameworks — logical frameworks — which are of course stimulated by certain issues that are there in the empirical field, but the data you need for that are usually just some small, very basic, characteristic things that you want to have a model for ... You will typically have just two sentences: "every farmer who owns a donkey beats it," and "a man is walking in the park, and he is wearing a hat." You don't have an enormous database of examples there. That's enough to motivate what you are after — what your formal framework is going to do...[13]

> I think the difference is that linguists are primarily interested in the data (in the data themselves), and I am not. I am primarily interested in the theory — in formalizing the data.[14]

So we see that formal semanticists are typically not motivated by the real world data of linguistic behavior in the same way that the linguistis are motivated by that data. Instead formal semanticists are motivated by certain puzzling sentences drawn typically from their own intuition. Then they proceed to build conceptual frameworks with in order to solve those puzzles. Not only is it the case that they are not motivated by the empirical data directly, but what is more troublesome is that the models and frameworks that they go on to build might not have anything to do with the data at all. Some semanticists see the models as having no relation to the observable phenomena of linguistic behavior or its counterpart in the brain. In response to the question of what the models are about, one semanticist responds as follows:

> I: What are these models actually about? What is the natural phenomenon that we are trying to model?

[13] Interview B1, Appendix B
[14] Interview B4, Appendix B

> S: Well, let me say what it is not. It's not a model of the mind or something like that. It's also not a model of the real world. It's not that either. What I think formal models always do — they just have a good framework that has a conceptual motivation, and they can explain certain phenomena. What are the phenomena? I'm afraid that for me the phenomena are basically my intuitions.[15]

As the critics point out, this lack of interest in observable phenomena of linguistic behavior stands in a stark contrast to the scientific aspirations of the discipline, according to which the field is getting closer to the data as it develops. "Things are changing there, because people now do lots of more extended tests — on the internet, by corpuses, and so forth."[16] The research manifesto of the ILLC's Logic and Language program, which is the research program housing the formal semanticists interviewed explicitly notes that empirical ratification of analytical work is their main ambition (emphasis added by me):

> The research program of the group encompasses a broad range of topics at the intersection of philosophical logic, philosophy of language, linguistics, and cognitive science. Major themes are human reasoning and interpretation of natural language, and the methods we use for investigation are mostly based on logical and philosophical analysis. *Empirical ratification of analytical work is our main ambition and touchstone for success.* Our research strategy is non-monolithic, allowing for different approaches, but demanding philosophical reflection and internal and external debate.[17]

Thus we see once again that there is a tension or a contradiction between the analytical or philosophical aspirations or the discipline and its empirical and scientific aspirations. This tension or contradiction cannot be explained away as a contradiction between the past and the future of the field. Neither can it be explained away as a contradiction between two different factions within the field. There are instances where the contradiction lives within the same individual. For example, in one of my interviews[18], the formal semanticist

[15] Interview B1, Appendix B
[16] Interview B3, Appendix B
[17] Institute for Logic, Language and Computation (2014a)
[18] Interview B4, Appendix B

began by describing himself as a logician rather than a linguist, because he is not so much concerned with the data. Then, when the discussion turned to the topic of modeling, and the standards for success and failure in modeling, he began to stress the importance of data and empirical studies:

> I: I'm curious about the fact that at the beginning of this interview when I was asking you if you would describe yourself as a logician or as a linguist, you said that you would describe yourself as a logician rather than as a linguist, because you are not so concerned with the data, but now it seems that this is also of fundamental importance to your work. How do you reconcile that?
>
> S: Yes, it is very important. I'm aware that it's very important, but I don't want . . . it's not . . . I'm not interested enough in the data to invest a lot of work in getting the data, and I think that's the big difference between me and the typical linguist...[19]

The same formal semanticist goes went on to say that for this reason he is moving away from the field of formal semantics into philosophy proper. So the tension that was alluded to in the previous chapter between the scientific and philosophical aspirations of formal semantics remains unresolved. As a result, the proper role of philosophical intuitions on the one hand and empirical data on the other hand, and how these relate to the formal models remains unclear.

5.3 On the success or failure of formal semantics

In the previous chapter we discussed the success or failure of formal semantics in the scientific as well as the sociological sense. The critics alleged that it was a failure from a scientific point of view, because it did not meet the criterion of empirical verifiability / falsifiability. They further alleged that it might also be a failure from a sociological point of view due to a lack of intellectual coherence. In this section I will discuss success or failure of formal semantics first from a scientific point of view, and then from a sociological point of view.

[19] Interview B4, Appendix B

"Is formal semantics a rigorously scientific discipline?" This was one of the questions that I asked in my interviews with the formal semanticists. None of the interviewees answered in the negative. All of them affirmed, to some degree, that it was a rigorously scientific discipline. Although they did also acknowledge that as it currently stands, there is a problem with respect to empirical verification / falsification of the models, this problem was explained away either by claiming that it will be solved in the future (as the field gets closer to the data), or by claiming that it will be solved by a different set of people (it is not the job of formal semanticists, but of empirical linguists). The first response argues that the problem of empirical verification / falsification will be solved as the field develops and gets closer to the data:

> I: Do you think that formal semantics is a rigorously scientific discipline?
>
> S: I think it's on the way. It's still very young. For instance, the data ... how to check our predictions and how to check our models using data . . . this is something that really just developed during the last twenty years. At the beginning it was possible to write a whole PhD thesis about the intuitions of the writer and his sister about presuppositions, and that was fine, but now it is no longer acceptable. So you really have to underpin all of your models using serious studies of the data.[20]

Another variant of the same response was to claim that the theory is ahead of the experimental corroboration. In other words, that the experimental corroboration would occur in due time.

> I: Would you say that formal semantics is a rigorously scientific discipline?
>
> S: Yes. I would say yes. I think that I would say it's as scientific as it gets. I do feel that sometimes people get the impression, from outside or from people peeking over the shed from the neighbors', that they might get the impression that it's not as scientific as other disciplines may seem. That maybe because, as I said, the theory is

[20] Interview B4, Appendix B

> often ahead of the experimental corroboration of all the predictions that those theories make.[21]

Yet another idea, that is closely related to the above suggestions is that there should be a separation of theoretical and empirical work in semantics. One semanticist made an analogy with physics — which is divided into the sub-disciplines of theoretical and experimental physics — arguing in effect that formal semantics should be seen as the theoretical branch of linguistics, while the experimental branch should be left to others:

> In natural science you have the experimental people —– they are doing experiments in big laboratories, etc. And then you have the theoretical physicists —– they are just sitting behind the desk ... they have a computer, and what they do is type and scribble on paper, and they build concepts. ... They are not doing the empirical research, they are doing the formal stuff, and then they hand it over to the empirical people. ... So in semantics you need the same stuff. There is the work that is being done by logicians ... purely analytical work that is of course motivated a bit by data, but it is not data description, it is really conceptual work, making a new system and then hand it off to the linguist who can play with it, and the linguist can (and should) then try and get all sorts of empirical testing and corpus or whatever, but that's not the job of the [formal semanticist].[22]

However, it is not at all clear that delegating the empirical work to some future point in time or to some other group of researchers is actually a solution to the problem rather than an avoidance of the problem. Such a proposal does not address the fact that there is no established methodology for measuring the success and failure of theories and models in formal semantics. Currently the field is lacking a proper understanding of how empirical work should relate to the theoretical models. As one semanticist points out, it is currently not possible for empirical work to verify or falsify theoretical models:

[21] Interview B2, Appendix B
[22] Interview B1, Appendix B

> So it's not that somebody did an experiment and said "These are the data, so this theory is correct." No, it's one experiment, and then other people have done another experiment with different results and a slightly different methodology. So the field is still finding out how to better combine empirical and theoretical methods. That's fair to say, but I would say it's as scientific as it gets.[23]

Thus it is not at all clear that the optimism of the semanticists is warranted when they claim that "it's as scientific as it gets"[24] or that "it's on the way to becoming rigorously scientific."[25] The critics would argue that it is neither scientific nor is it on the way to becoming scientific, so long as it is unable to provide a bridging theory of how the idealizations are related to the real world:

> In an idealization you simply forget about that [various empirical features of language], and that means that you need something else to translate your results back into results that are relevant for what you idealized away from, but linguistics doesn't contain its own bridging theory. In that sense, yeah, it's a failed discipline.[26]

So, although the formal semanticists have an optimistic outlook regarding the scientific status of their discipline, there remain some fundamental tensions between their theoretical models and possible empirical work that remain unresolved. Until these tensions have been resolved, it seems premature to claim that the discipline is rigorously scientific or that it is on the way to becoming rigorously scientific.

Next we come to the sociological dimension of the issue. In this dimension, the formal semanticists describe their situation far more pessimistically. The crucial problem seems to be that formal semantics, according to its traditional conception, is having a very hard time fitting into the linguistics community. One semanticist describes this difficulty in quite some detail:

[23] Interview B2, Appendix B
[24] Interview B2, Appendix B
[25] Interview B4, Appendix B
[26] Interview A1, Appendix A

> I am maybe not even very optimistic about how things go. First of all, what I have noticed over the last couple of years ... because when I and people like me (the other formal semanticists here and around me . . . at the ILLC, and maybe in Holland . . . we have this sort of community of logical semanticists), when we go, we mainly go to conferences that have linguists ... I notice that it gets harder and harder for us — when we write abstracts that you want to get accepted — to get accepted. What it means to me is that the real logical conceptual work is apparently no longer that much appreciated, and you can also see this by the sort of commentaries that you get when a thing is not accepted. What very often plays a role is empirical motivation. If you don't start your abstract with a concrete empirical problem that you want to solve, and then give the machinery, and then show how you've solved it, then you are in trouble. ... That's the linguist's typical procedure, and he is not going to accept it if you do it in the other direction. And for me, if I do it in the other direction then I am cheating, because I did not start with these empirical issues ... empirical methods in semantics are sort of rising, and if your story is not built upon hardcore empirical stuff, then you are not doing something that's relevant. And I think that's very very wrong.[27]

This pessimistic expression touches upon many of the themes that we have been discussing so far such as the criticism of the critics, the tension between its philosophical and scientific aspirations and the possibility of its failure. According to the traditional conception of formal semantics, it is not so much concerned with empirical data, and it's main focus is instead conceptual or philosophical analysis. Linguists on the other hand are keen to move away from the "armchair linguistics" stereotype towards a more empirical and rigorously scientific methodology. The upshot of all this is that formal semanticists are beginning to feel displaced from what was once their home community (linguistics). As a result of this some people are beginning to abandon the field of formal semantics. Already we saw that the critics have done so, but so are some formal semanticists:

[27] Interview B1, Appendix B

> I'm on the way [to] leaving the field. So why is that? Basically because I really . . . I'm not interested enough . . . on the one hand I'm not interested enough in the data I think as what is necessary right now in the field. So if you now want to do serious formal semantics you really should get into systematically studying data, and that's just not what drives me. That's one of the things, and the other thing is that I'm still not satisfied with the . . . how to describe it? . . . with the level of reflection on the theory building and on the methodology that is done in the field. So I want to get deeper than just focus on the data and just building a system that covers these data . . . I want to get deeper. So that's why I'm moving to philosophy more and more I think.[28]

Formal semantics as a field is attempting to set more rigorously scientific and empirical standards for itself as part of an attempt to adapt to the standards of the larger linguistics community. In some cases, those who are unable or unwilling to adapt in this manner are abandoning the field altogether. This also raises some interesting questions about logical practice in general. If logic is an independent discipline, why do some of its practitioners (such as formal semanticists) need to seek validation from other communities such as the linguistics community at large? What is it about logic that makes its practitioners seek validation by showing that their work is relevant to other academic disciplines? It would be interesting to see how this compares to the state of other disciplines such as mathematics. Do mathematicians need to seek validation by showing that their work is relevant to other academic disciplines? Although interesting, such a discussion is beyond the scope of this thesis.

5.4 Conclusions

The formal semanticists interviewed in this study all acknowledge the criticisms made by the critics. They recognize that there is some degree of idealization involved in their process of modeling natural language by means of formal representations, and they further acknowledge that this also poses problems for the

[28] Interview B4, Appendix B

scientific validity of their models. However, they are reluctant to acknowledge that these criticisms pose an existential threat to their academic discipline. They continue to assert that despite acknowledging these criticisms, formal semantics remains to be a legitimate academic and scientific endeavor.

This might seem puzzling at first, but it begins to become less puzzling when we invoke the Kuhnian framework, according to which the development of sciences proceeds in two distinct stages: normal science (wherein researchers are working towards solving puzzles within a paradigm, without calling the paradigm itself into question), and revolutionary science (wherein a paradigm dies and is replaced by a rival paradigm). "In so far as he is engaged in normal science, the research worker is a solver of puzzles, not a tester of paradigms." Thus it is not surprising that the research worker in formal semantics, being a solver of puzzles in this paradigm, is not inclined to test or call into question the paradigm itself.

Nevertheless, the concerns raised by the critics remain valid, and there seems to be no satisfactory response from the formal semanticists. The tension that was identified between the philosophical and the scientific aspirations of formal semantics — giving rise to a so-called "identity crisis" — remains unresolved. The critics urged that formal semantics should concretely determine their own identity by choosing between one of their two aspirations: either become rigorously scientific by focusing on observable linguistic behavior, or give up the aspiration to be a rigorously scientific discipline. The formal semanticists are unable to decide on an identity, and seem to suffer from the syndrome of "wanting to have their cake and eat it too." Some formal semanticists explicitly deny that their models should be about anything observable, and continue to insist that it is sufficient to be guided by intuition. This is in stark contrast to the scientific aspirations of the discipline, according to which "empirical ratification of analytical work is our main ambition and touchstone for success."[29]

What does this say about the question of failure? In the scientific dimension, formal semanticists maintain that their discipline is rigorously scientific. Although they recognize that they are lacking an experimental methodology which would allow for empirical verification or falsification of their models, they

[29] Institute for Logic, Language and Computation (2014a)

explain this problem away by delegating this task to some point in future or to some other group of researchers. By the critics' lights, this would count as an avoidance rather than solution to the problem, because formal semantics is not able to give a clear account of how such an experimental methodology could be developed, or even if it is possible to do so in principle.

Moreover, in the sociological dimension, some of the formal semanticists are themselves quite pessimistic about the future of the discipline. It is pointed out that what was once their home community — linguistics — is now reluctant to accept the work of formal semanticists, because the linguists want to move away from the stereotype of "armchair linguistics" towards a more rigorously empirical, data-driven approach. Therefore it seems that formal semanticists are finding it increasingly difficult to fit into this community, and as a result some people are abandoning the field altogether. As a consequence, the goal of sustaining a thriving community of students and researchers is becoming increasingly difficult to live up to.

Chapter 6

The outsiders: computational linguists

> *Competition between segments of the scientific community is the only historical process that ever actually results in the rejection of one previously accepted theory or in the adoption of another.*[1]

As we see in this chapter, part of the criticism (made by the critics) of formal semantics concerned its failure to deliver on any of the practical applications that it implicitly promised — such as a machine or computer program that can understand natural language. They point out that the discipline of computational linguistics on the other hand has enjoyed tremendous success with respect to such applications. The computational linguists also claim that in addition to practical applications, their models also have far greater explanatory and predictive power in a scientific sense. Within the framework of a Kuhnian analysis, this naturally raises the question of whether computational linguistics might be a competing paradigm?

In this chapter I first (in section 6.1) argue that computational linguistics and formal semantics can be considered — in the Kuhnian sense — competing segments of the scientific community. In the following section (in section 6.2) I present the competition from the point of view of formal semanticists who argue

[1] Kuhn (1970b), p. 8

that computational linguistics is a mere engineering tool, and then (in section 6.3) from the point of view of computational linguists who argue that it is in fact a rigorously scientific discipline. Finally (in section 6.4) I present a dissenting voice from within the computational linguistics community who argues that computational linguistics cannot fulfill the more theoretical aspirations behind formal semantics.

6.1 Computational linguistics vs. formal semantics

Kuhn tells us that competition between segments of the scientific community is the only process by which scientific revolutions (i.e. the death of one paradigm and its replacement by another paradigm) occur. In this chapter, I wish to explore the possibility that computational linguistics and formal semantics are two such competing segments of the scientific community, and what potential this competition has (if any) for displacing the paradigm of formal semantics. First we acknowledge that computational linguistics and formal semantics are in fact competing paradigms. As one computational linguist (and former formal semanticist) points out:

> Certainly they [computational linguistics and formal semantics] do fit the definition of what one calls a paradigm. They're just different ways of operating and different ways of looking at the problem to begin with.[2]

Although this point may not be obvious, and although there is very little scholarly discussion of such developments in the literature, the interviewees of this study seemed to be in agreement with this suggestion. For instance:

> I: Do you view these two disciplines as sort of competing ways of analyzing the same phenomenon?
>
> S: Yes. Mostly people just view them [formal semantics and computational linguistics] as completely distinct and not having anything to do with each other, but actually my advisor would

[2] Interview C4, Appendix C

> always say that we are in competition for explanatory power. Just because we are using different formalism — on the one side logic and on the other side statistics — doesn't mean we are doing different things. We are doing the same thing and we are in competition for explanatory power. It's basically up to your standard of evaluation to say that one is better than the other, but they are in competition definitely.[3]

Here the subject begins by pointing out that mostly people view the two disciplines as completely distinct and not having anything to do with each other. This is also corroborated by the fact that there is very little scholarly literature on this topic. There is however one volume on the subject, which was the result of a workshop called "Computational Linguistics and Formal Semantics."[4] The volume contains an epilogue on the relation between computational linguistics and formal semantics which surprisingly (surprising given the title of the workshop and the volume) concludes that there is no intimate connection between the two fields:

> [C]alling a workshop "Computational Linguistics and Formal Semantics" assumes that these two areas of academic endeavor have something to say to each other; may, even, be inextricably related. Superficially, of course, this seems likely to be true: an investigation of language and its use could be seen as the core interest of both disciplines. But, on looking a little more closely, the intimate connection tends to evaporate.[5]

The main reason cited in support of this surprising claim was that computational linguistics is concerned primarily with computer applications, while formal semantics is concerned primarily with the formalisms themselves. However, whether or not they have something to say to each other is a separate question than the question of competition. My interviewees certainly gave me the impression that they are in some sense considered to be in competition

[3] Interview C3, Appendix C
[4] Rosner and Johnson (1992)
[5] King (1992), p. 283

with each other, with the computational linguists winning (at least in terms of practical applications). "In terms of practical applications, certainly [they are competing], with the statistical ones winning."[6] The "statistical ones" here refers to the statistical methods of computational linguistics. One computational linguist jokingly dismisses the question of competition, by saying: "at the moment there is no competition — we basically won!"[7] Although this was intended as a joke, such comments suggest quite strongly that the notion of competition is not completely out of place in this discussion.

6.2 From the point of view of formal semanticists

In our discussion thus far, we have focused only on the philosophical and scientific aspirations of formal semantics. However, it is worth noting that formal semantics does also have some practical / engineering oriented aspirations. Part of the criticism of the critics was that formal semantics has been a failure with respect to these aspirations while disciplines like computational linguistics have had great success with their stochastic methods:

> As the theoretical models of the generative tradition [including formal semantics], based as they are on the notion of a grammar as an system of explicit rules, failed to deliver in applications such as machine translation, question-answer systems, and the like, people started to use other constructions of central concept such as 'language', 'meaning', and so on. Often these new constructions were based on stochastic properties and patterns derived from large corpora of actual text (and, later, speech). These constructions were based on other, often less far-reaching idealizations, i.e., they stayed closer to the original phenomenon and hence were more amenable to empirical testing.[8]

One might object (against the critics) that formal semantics perhaps does not have any such engineering oriented aspirations, and it is therefore unfair to

[6] Interview C4, Appendix C
[7] Interview C1, Appendix C
[8] Stokhof and van Lambalgen (2011a), p. 18

criticize them for failure on these grounds. But the fact that formal semantics does have these engineering oriented aspirations is quite clear even from their own descriptions of their work. One semanticist notes that the engineering aspiration is always (to some degree) in the back of his mind when he is doing research in formal semantics:

> When I make something it has to be of the nature that in principle it can be used, say, in building computer systems that can understand language and that can do some technical things, but I am not *working* on that (that's not what I'm *doing*), but I do want to do things in a fairly explicit way such that in principle that is possible. I'm not really steered by those issues ... I have my own agenda ... whether they can really use it or not — that's not my business, but I do want to do it in a way such that in principle that is possible.[9]

Another semanticist notes that the engineering aspiration of building a machine that can understand natural language is part of the goal of formal semantics in the sense that it would be the ultimate test of the success or failure of formal semantics:

> Well, if you really understand it, of course, you could build a machine that understands natural language for instance. Because you could then implement in the machine what it is to know the meaning of the expressions of the language. So that would be the ultimate test I think.[10]

But then he goes on to say that he would not have any interest in actually building such a machine:

> The goal is just giving a description of how meaning arises, but you can test whether you've reached the goal if you can build a machine. I would never build the machine myself, and I would probably think it's a waste of time ...[11]

[9] Interview B1, Appendix B
[10] Interview B3, Appendix B
[11] Interview B3, Appendix B

Then he follows this up with a claim that formal semantics will play an important role in the construction of such a machine (despite the fact that he has no interest in building such a machine, and despite the fact that we are quite far off from realizing such a machine):

> We are actually very far from what I said [earlier] that we can build a machine that understands natural language, but I am sure that what formal semantics has to say will play an important role there.[12]

Thus it is apparent that formal semanticists are at some implicit or explicit level motivated by engineering oriented practical applications for which they believe formal semantics is relevant (despite the fact that formal semanticists try to distance themselves from such aspirations at a personal level).[13] At any rate, even the formal semanticists don't dispute the fact that computational linguistics has had far more success in this arena than formal semantics. In my interviews with the formal semanticists, I asked the interview subjects for their reaction to this fact. The usual reaction was for the formal semanticists to be mostly dismissive of the successes of computational linguistics. The gist is that although these are successful in engineering oriented applications, they are not successful in the scientific or philosophical sense of achieving "understanding" (of language presumably):

> They [computational linguists] have far more success, but the machine doesn't understand anything, so I don't believe in that approach at all ... when it comes to the idea of "let's build a machine that understands natural language."[14]

> Yes, that's interesting, because that approach [computational linguistics] doesn't focus so much on understanding, and it focuses completely on coverage of data. In this respect it's much more

[12] Interview B3, Appendix B

[13] This simultaneous interest and lack of interest in building a machine or program that can understand natural language, would make for an interesting sociological study in it's own right, but is beyond the scope of this thesis.

[14] Interview B3, Appendix B

> successful than we are. While we say the benefit of doing formal modeling is to gain a deeper understanding of the problem, but our coverage is extremely less strong than what stochastics offers. What can I say? I mean ... I see the benefits of the stochastic approach, but there are limits, and I think that sometimes now they even reach these limits. There are limits that they just can't pass because they just don't understand what they are modeling. I hope that maybe this is going to be even more pressing in the future, and then it will be more important again to have models that really understand the phenomena that you try to model.[15]

The allegation is that computational linguistics is exclusively application oriented, and that it does not contribute to a deeper philosophical or scientific understanding of language. The nature and tone of the reaction also suggest to me that formal semanticists are a bit defensive in their reaction to the success of computational linguistics. Such a reaction supports the hypothesis that they are in fact competing disciplines.

6.3 From the point of view of computational linguists

In my interviews with the computational linguists, I asked them to comment on the relation between computational linguistics and formal semantics in general, but I also asked them to comment in particular on some of the claims made by the formal semanticists — for example, the claim that stochastic models used in computational linguistics are mere engineering devices which don't contribute to a deeper understanding of natural language. Predictably, computational linguistics did not agree with this assessment:

> The claim that stochastic models haven't contributed to our understanding of natural language is just plain false, because basically all modern research that is being done in psycholinguistics or

[15] Interview B4, Appendix B

> any kind of behavioral linguistics where you examine people, where you conduct experiments, where you gather data . . . is done in some sort of statistical evaluation, and statistical evaluation includes building stochastic models.[16]

> I would just say the same thing that a physicist would answer: Look, you don't know what it means to understand what's going on. When you understand a phenomenon, you actually can predict what's going to happen. Understanding demands from you to predict human behavior, because that's the gold standard. ... if you're talking about [formal] semantics, I dare claim that none of the algebraic [logical] representations that we've been occupied with over the last 40 or 50 years in semantics are actually any kind of evidence for understanding the phenomenon of human language processing...[17]

The computational linguists dispute the claim that they are only focused on engineering. They claim that their enterprise is legitimately scientific in the sense that physics is legitimately scientific, insisting that their models should be judged by the gold standard of human behavior, as measured with experiments and data. This evidences one of the major differences between the two fields: computational linguistics evaluates itself by means of a clearly testable empirical methodology. As the critics noted: "they [computational linguists] stayed closer to the original phenomenon and hence were more amenable to empirical testing."[18] The stochastic models give predictions (a prediction for how to translate a sentence for example), and the model is considered to be successful if the predictions are in agreement with human judgment, and there are clear guidelines on how to measure this agreement.

> It [computational linguistics] has a clear methodology and it has certain research standards that you can follow, so it's much easier to evaluate your results in computational linguistics (just as

[16] Interview C3, Appendix C
[17] Interview C1, Appendix C
[18] Stokhof and van Lambalgen (2011a), p. 18

> any kind of empirically driven science) against what other people are doing, against what has been found out before ... maybe also to falsify findings that other people have claimed. There is this expression "armchair linguistics" — it's where people just sit on their armchairs and think about stuff, and that makes it really hard to falsify them.[19]

Thus computational linguists define themselves directly in opposition to so-called "armchair linguistics" (in which formal semantics would certainly be included). "Generally you are not supposed to choose examples with a phenomenon you like, and then show that you are doing something reasonable on that. Generally you start with the data set."[20] It seems that the field has also developed a rigorous methodology to use these data sets for testing its models: "The field has developed over the last 20 years a methodology of testing where you're standing in terms of modeling ... So there are benchmarks for experiments and for testing where your model is."[21] In this sense, there are some clear advantages over formal semantics where such standards do not exist.

> There are clear evaluation standards in computational linguistics. Those standards just don't exist in formal semantics. Usually there are no standard data sets that everybody agrees upon. So, basically every author is picking his own example sentences. Sometimes they are the same, but mostly it's just random sentences that the authors themselves made up. ... In computational linguistics, people are really investing their whole research time into developing metrics for evaluation. Such an effort is not even being made in formal semantics.[22]

Thus computational linguistics claims to model the phenomenon of natural language in a way that has not only produced many engineering successes, but also in a way that addresses many of the concerns raised by the critics of formal semantics.

[19] Interview C3, Appendix C
[20] Interview C2, Appendix C
[21] Interview C1, Appendix C
[22] Interview C3, Appendix C

6.4 A dissenting voice

As a counter-point to the views represented in the previous section, which argued that computational linguistics is in fact a rigorously scientific discipline, I would like to present a dissenting voice in this section which argues that computational linguistics is severely limited by its engineering goals and that it cannot be properly construed as a scientific discipline. Although this interviewee was one of the pioneers of the field of computational linguistics and an early advocate for the use of statistical methods in linguistics, he is now more sympathetic to formal semantics than to modern day computational linguistics. He is dissatisfied with the fact that computational linguistics has become an independent paradigm, rather than a tool for enhancing existing paradigms:

> I've been actively involved in launching the statistics stuff, and I've only been disappointed in that it became this independent paradigm which didn't have any connection with the [formal] modeling anymore. That was never my intention, but I believe in the importance of the statistics.[23]

Due to the fact that computational linguistics has severed its connection with the aspirations of formal modeling, he argues that it can no longer be considered to be a scientific discipline, and should instead be seen as a mere tool for practical / engineering purposes:

> I: Do you think that computational linguistics is a rigorously scientific discipline?
>
> S: No . . . because . . . certain versions of it are more practical where the goal isn't even scientific. You just try to make something that works, and if it works it works.[24]

He is particularly worried that the means by which computational linguistics has achieved its significant engineering successes — especially the rigorous testing and evaluation criteria — might have turned in to a kind of intellectual straitjacket:

[23] Interview C4, Appendix C
[24] Interview C4, Appendix C

> The thing is people tend to focus very much on just the performance on the test set. It's very difficult to publish a paper which contains new ideas if it doesn't give better performance on the standard test sets. I don't think that that's a very good situation ... because it sort of makes the test set as a goal in and of itself, as if there are no other intellectual goals anymore. The questions of whether something makes sense from a cognitive point of view, or a linguistic point of view, or even from a practical point of view — those questions aren't on the table anymore if you focus completely narrowly on the test set behavior.[25]

Computational linguistics has severed connections with some of the intellectual goals that motivated formal semantics. Thus certain questions and intellectual aspirations of formal semantics are no longer on the table. Being dissatisfied with the direction in which computational linguistics is developing and with the fact that it has severed its connection to the questions that motivated formal semantics, the subject expresses a desire that these two paradigms should come together collaboratively rather than competitively:

> Then perhaps the most interesting question is whether they are competing approaches if we talk about possible future theories of cognition. There it seems obvious that we need some kind of synthesis ... that we do need both approaches, but in a way where they are not competitive, but where they are integrated.[26]

Although some of the other interviewees did also pay lip service to this collaborative goal, there is very little concrete work being done towards addressing this goal.

> I think in general you should talk more to formal semanticists about the fact that they don't often come to our conferences, etc.

[25] Interview C4, Appendix C
[26] Interview C4, Appendix C

> Maybe that's not right, but maybe you can blame us for not learning enough formal semantics to be able to meaningfully speak to them.[27]

As things stand, there is very little meaningful dialogue between the two disciplines. One could argue that this is in fact typical for competing paradigms, and that this might also related to Kuhn's controversial incommensurability thesis, according to which "science guided by one paradigm would be 'incommensurable' with science developed under a different paradigm, by which is meant that there is no common measure for assessing the different scientific theories"[28] — at least there is no common measure that both parties can agree upon.

6.5 Conclusions

We have so far been discussing formal semantics only in terms of its scientific and philosophical aspirations (and the conflict between them), but it turns out that the field also has some engineering oriented aspirations for building practical applications. Although these goals are not really emphasized, they are implicit in the project of formal semantics, and some of the interviewees have even called attention to them explicitly. In particular, the goal of building a machine that can process natural language is espoused by formal semanticists as a practical use case of formal semantics. In terms of such applications however, the stochastic models of computational linguistics have been extremely successful (not only in academia, but also in industry), while formal semantics has had almost no success (neither in academia nor in industry).

This raises the question of whether computational linguistics can be seen as a competing paradigm and a potential threat to formal semantics. Many of the computational linguists and formal semanticists interviewed do certainly give the impression that these two can be seen as competing paradigms (with the computational side winning — at least in terms of engineering and practical applications). The critics also point to this as an example of formal semantics'

[27] Interview C2, Appendix C
[28] Bird (2013)

failure, and they note that an important reason for the success of computational linguistics is the fact that their models stay closer to the original phenomenon (of natural language) and hence are more amenable to empirical testing. In response to this, the formal semanticists allege that these stochastic models of computational linguists are *merely* engineering devices and that they do not contribute to a deeper scientific or philosophical understanding of natural language.

The computational linguists on the other hand do not agree that their discipline is a mere engineering discipline. Although it emerged as an engineering discipline, they now view it as a rigorously scientific discipline. The especially emphasize the fact that computational linguistics has a clear methodology and clear metrics for evaluation in the form of standardly agreed upon benchmarks and experimental methods for empirical verification or falsification their models. They oppose this (much like the critics) to formal semantics, which lacks clear standards for evaluation and model testing. They also criticize formal semantics for the fact that it is not motivated by empirical data (it is instead motivated primarily by the intuitions of the researchers), and is thus associated with the negative connotations of "armchair linguistics." Thus computational linguistics claims to model the phenomenon of natural language in a way that has not only produced many engineering successes, but also in a way that addresses many of the concerns raised by the critics of formal semantics.

There are, however, some dissenting voices from within the computational linguistics community. In particular, one of the interviewees who was a former formal semanticist and early pioneer of statistical and computational linguistics finds himself more sympathetic to formal semanticists. He points out that the features of computational linguistics that made it successful with respect to the engineering goals — in particular the rigorous evaluation and testing procedures — have become a kind of intellectual straitjacket for the discipline. Therefore it is unable to take on all of the intellectual aspirations of formal semantics (especially the more theoretical aspirations). Therefore, although computational linguistics could be a replacement for some aspects of formal semantics, some of the crucial intellectual ambitions of formal semantics are not shared by computational linguistics — especially in terms of its philosophical aspirations, these have no place in computational linguistics.

Chapter 7

Conclusion: towards a philosophy of logical practice

In this chapter I first (in section 7.1) conclude the case study in formal semantics, and then (in section 7.2) I indicate some possible directions for future work towards a philosophy of logical practice more generally.

7.1 Concluding the case study

I began this case study with the context of the development of formal semantics as a scientific and philosophical discipline, and its subsequent criticism by a couple of its prominent practitioners. The goal of the case study was to shed new light on this debate, by allowing all parties to voice their concerns, criticisms and responses in their own words.

In the first chapter of this case study (chapter 4: the critics), I presented the criticisms of the critics in terms of a conflict between the philosophical and scientific aspirations of formal semantics — in particular, I analyzed the conflict between the idealizations involved in the models of formal semantics (which demands intuitive and conceptual analysis of linguistic competence) and the scientific aspirations of the discipline (which demands an empirical analysis of data from actual linguistic behavior).

In the second chapter of this case study (chapter 5: the insiders), I probed the responses of the formal semanticists to this criticism — I discovered that although they are confident that the discipline is a scientifically and philosophically valid endeavor, they expressed some pessimism with respect to the goal of sustaining a thriving community of students and researchers — especially in the context of the external climate of the larger linguistics community, which is presently reluctant to accept work that is not data-driven.

In the third chapter of this case study (chapter 6: the outsiders), I explored another dimension of the criticism of formal semantics — namely, its failure to deliver any practical applications, and I considered the possibility that computational linguistics is a competing paradigm. Here I discovered that although it meets the criteria of bringing a data-driven and empirically rigorous approach to the study of natural language, computational linguistics has no place for some of the more overtly theoretical and philosophical motivations behind formal semantics.

Finally, I concluded that the outcome of these developments is still very much unclear at this point, but I speculated that we might look to the critics' own careers as an example of what the future might hold — one has given up formal semantics in order to pursue rigorously scientific research, while the other has given up his scientific aspirations to pursue philosophy proper. It is possible that others will follow a similar trend (some already have done so), but it is also unlikely that any major crisis will occur until some alternative paradigm gains momentum (at the moment, there is also no obvious candidate).

Besides making a contribution to the debate around formal semantics, I also hope to have shed some light on questions of general interest to philosophers of logical practice — concerning the relationship between formal models and the real world phenomena they claim to be modeling and concerning the basis for verifying the accuracy of such models. Hopefully this encourages formal semanticists and logicians more generally to engage in some philosophical reflection on the way in which they make use of formal models.

7.2 Towards a philosophy of logical practice

Having concluded the case study, I would like to conclude this thesis with some broader considerations about the philosophy of logical practice (PLP). With this case study I hope to have demonstrated one particular approach to PLP, but with this thesis I also hope to have sparked some general interest in the PLP field, by convincing the reader that it is a worthwhile and fruitful academic endeavor. In this final section I would like to indicate some potential directions for future research in the PLP field:

1. *Additional exploration of this case study.* As noted earlier, the interview material contained in the appendices is far richer than my analysis in the main chapters of this thesis. A full analysis of the interview material would have been beyond the scope of a master's thesis (not to mention, beyond my level of expertise and capacity for analysis), but I believe that the interview material can serve as a useful data point for further analysis of this case study (perhaps even with an expanded scope that goes beyond the question of failure that was discussed here).

2. *Additional case studies.* Formal semantics is only one particular form of logical practice. Similar case study could be conducted with a focus on some other forms of logical practice (e.g. formal epistemology).

3. *Additional methodologies.* This case study implemented the methodology of qualitative research in the social sciences. However, it is also possible to make use of quantitative social science research, historical research, psychological research, and various other research methodologies. Each of these methodologies (among others) has some place in PLP and would be interesting in its own right.

4. *Additional engagement.* At the moment, Catarina Dutilh Novaes is the only other researcher I am aware of who is attempting to engage with this kind of research. It would be a great pleasure for me to see more logicians, philosophers, mathematicians, linguists, historians, social scientists and other researchers engage in such investigations.

Appendix A

Interviews with the critics: Stokhof and van Lambalgen

This appendix (Appendix A) contains transcriptions of the interviews conducted with the critics: Martin Stokhof and Michiel van Lambalgen. The interviews have been transcribed using the following schema:

> I: Lorem ipsum dolor sit amet, consectetur adipiscing elit. Aliquam ultricies lacinia euismod.
>
> S: Nam tempus risus in dolor rhoncus in interdum enim tincidunt. Donec vel nunc neque.

"I" denotes the interviewer, and "S" denotes the interview subject, i.e., the person who is being interviewed. The interviewer ("I") is in all cases the author of this thesis, while the subject ("S") denotes an anonymous member of the set {Martin Stokhof, Michiel van Lambalgen}.

NB: These interviews are semi-anonymous in the sense that the individual subject of each interview remains undisclosed. Therefore the reader is requested not to identify any particular interview with any particular individual in any future citations of this work.

A.1 Interview with a critic

I: I would like to begin with a bit of an introduction from you. Can you briefly describe your academic and professional background, and how you got interested in the field of formal semantics?

S: Yes, I'll try to keep that short. I studied philosophy and linguistics in the early 1970s, and at the time I got interested (I was always interested) in language. Then I got introduced to logic at the very start of the philosophy curriculum and that also caught my fascination ... the idea that you could use logic to represent things in and about language was immediately something that Jeroen Groenendijk and I were fascinated by — we studied together. Then an independent interest in the developments of linguistics started to form, which was mainly at that time generative grammar in the Chomskyan style. But we also did a lot of typological research. At some point someone showed us a mimeograph of a paper by Richard Montague who had been visiting the department of philosophy a couple of years before that. That was a mimeograph of "Proper Treatment of Quantification in Ordinary English." Johan van Benthem was then a newly appointed assistant professor in logic. So we asked him to read that material with us, and to explain things that we did not understand. So that's what got us first interested — a combination of an interest in logic as a representational tool, and linguistics, and all the philosophy of language that is in between.

I: Which of these would you say was your primary motivation? Or were they both equally important?

S: I think while the fascination of language goes back even further, so that must have been the deeper motivation.

I: Would you describe yourself primarily as a linguist or as a logician or something else?

S: Well, the convenient aspect of having this kind of double background is that you can say to the logicians that you're not a logician but a linguist or a philosopher and you can say to the philosophers that you're a logician or

a linguist, so that's really very convenient. But in reality I would say that certainly now, I'm basically a philosopher.

I: I now have a few questions for you about the nature of formal semantics. To start with: how would you answer the question "what is formal semantics?" How would you describe it to a university student who has never heard of the field?

S: I would say it's the application of logical techniques (or in a broad sense, the application of formal methods) to model natural language meaning. That would be my short characterization.

I: And if you were to give a slightly more descriptive characterization...

S: It's the use of logic, and maybe some techniques that are developed elsewhere that use logic in a broad sense, to model intuitions (and other kinds of data that we have) about core semantic properties like entailment, synonymy, analyticity and so on — the things that characterize natural language meaning from a truth conditional point of view. That would be my characterization of what formal semantics actually is, but not necessarily what it should be or what it can or cannot be. That's sort of the nature of the enterprise ... take an interesting range of phenomena from natural language (like tense and aspect / mass vs. count / quantifiers / anaphora), make an inventory of how basic semantic properties (like entailment, synonymy and so on) play out in that particular domain, and then come up with a logical theory that accounts for those phenomena.

I: Now that you've introduced this distinction between what formal semantics is and what you think it should be, maybe you also want to briefly say something about the latter?

S: Well, I think that that is actually very fine, and that is what formal semantics does best, but I don't think (unlike maybe I used to think) that thereby it is a straightforward empirical discipline — something like a science that not only describes but also explains. I think it's more complicated ... in the sense of what it contributes to our overall understanding of what language is and how language functions. It's not that it simply describes and explains a particular

aspect of that ... the kinds of descriptions that it makes are more higher level kind of accounts of certain aspects of what is involved in language.

I: So I will return to this point later in this interview. Could you describe what you think the goals of formal semantics as a discipline are?

S: Again as a discipline I think that the ultimate goal is to come up with an explanatory theory of core features of natural language meaning. I just was at a workshop this weekend which was about bridging the gap between formal semantics and conceptual semantics. Conceptual semantics is (at least there was understood as being) primarily concerned with lexical meaning and the internal structure of lexical meanings. That's something that formal semanticists don't have much to say about — and maybe even don't want to say anything about — because the toolkit that they apply simply doesn't seem to be geared towards describing and explaining what goes on inside the meanings of lexical items. So its core features are usually understood in terms of the more structural semantic properties of a natural language (like as I mentioned quantification / tense and aspect / stuff like that).

I: So what do you think are the main kinds of problems that researchers in formal semantics are trying to solve?

S: About 80% of the work (that's a rough estimate) or the bulk of the work goes into describing interesting and puzzling phenomena. "What is the distinction between mass nouns and count nouns from a semantic point of view?", or "How do we come up with a good theory of temporal expressions?", "How to deal with indexicality?" and stuff like that. That's descriptive. Then there's some stuff on the interface between syntax and semantics (that's particularly strong in the U.S. based work) where people say that semantics is part of the grammar and there are other parts of the grammar (like syntax and morphology) so how does it [semantics] interact with that [the other parts of the grammar]? Are these completely independent modules that you could click together? Presumably not, so there are all sorts of interface effects. On the other side there's the interface with pragmatics. How much of natural language meaning in a broad intuitive sense can be accounted for (or should be accounted for) in terms of the truth conditional semantics that formal semantics deals with, and what should be relegated to the pragmatic aspects (things that have to do with the use

of expressions in context for certain purposes, etc.)? Again there too there's lots of research that explores what belongs where and people will argue about where to draw the line exactly.

I: So when are these problems considered to be solved? For example, you were talking about giving a description of some interesting phenomena. When would the researchers think that they have solved this problem?

S: That's a good question. If you look at the facts then I think that hardly ever does it occur that the problem is solved, because usually solutions just create more questions. So in that sense I don't think there is [any] part of what formal semanticists have been studying that is considered [solved] (as in: "OK, we've done that; you can look up the results there; you don't need to bother about it"). That's because, of course, if you solve a particular problem in a particular way that will have repercussions (if only because of the assumptions of how you solved it) for what you can or cannot say about other things. If we give a description of, say, temporal expressions in a particular way and then we say "OK, I've done that, now let's turn to modalities", and then you observe that of course modalities and tense interact, and what looked like a good description of the temporal aspect actually turns out to have negative implications for how you can deal with modality and then you have to go back again.

I: So it sounds like somehow the overall goal would be to have a unified description or model that can account for all of these features [of language]? Or is that not the goal?

S: Well that used to be very much and very explicitly the goal in the early days, especially in the Montague tradition where people wrote very explicit fragments — a syntax that generated a set of expressions and a semantics that assigned meaning to those expressions, and then if you wanted to do something else you were supposed to do that on top of what was already there. The idea was that you could actually extend the fragment. That turned out to be not very practical. If you think about it, it was a bit of an impractical goal, but I think that certainly those formal semanticists who subscribed to the claim that they are not only giving nice formal models that account for a bunch of data but that they are describing something that has a cognitive reality — they must be committed actually to this "one big coherent picture" thing. Unless

you go for a highly modularized view on cognition — but even then, there must be chunks in the way that humans deal with language that are more comprehensive than a particular aspect like tense or anaphora. These things interact, and if we humans are able to do that in a more or less coherent and consistent way then there better be a theory that models that as a coherent and consistent set of principles and procedures. So I think that in that sense formal semantics is committed to that, but actually if you look at the practice then that is something that takes a back seat.

I: I'm moving on to some specific questions about what role modeling plays in formal semantics. Logicians are often concerned with modeling — in the sense that they try to create formal representations or models of some real world phenomena in order to represent them and to reason about them. Can you describe what role this notion of modeling plays in formal semantics?

S: I find that a very difficult question. First of all there isn't much discussion actually about this, and what discussion there is (what explicit statements you can find about this) is of course not in the actual work that's being done, but you find that in textbooks and in notes of lecture courses where people start to introduce students into the view — then of course they have to say what it is that they're actually doing. So I looked at that on a number of occasions to try to find out what semanticists themselves think that semantics is. There are many ways of describing what semantics does (especially the role of formal languages in the enterprise) where the formal languages themselves are described as models, and I think that that's certainly inherent in the Montague paradigm. Doing semantics is translating sentences generated by a particular grammar into formalisms / formulae of some formal language, and then interpreting them in that formal language. Given that the translation is a homomorphic function, we can look at the meanings assigned to the formulae as meanings of the expressions that were translated into those formulae. Basically you look at the formal language with its interpretation as modeling the properties of the natural language expressions in the relevant respects. So I think in that sense much of formal semantics is actually modeling, but in reality (and certainly if you look at the plurality of logical tools that is increasingly being used to describe and analyze natural language phenomena) the practice looks more like that of using formal methods as a toolbox to get something done

without thereby considering the tool as itself a model. So it's still modeling, but in a difference sense. The formal language is itself no longer a model but it is something that is used to describe a particular range of phenomena. To what extent that models something becomes unclear, unless of course you take the thing one step further and say that basically what I'm modeling is cognitive structures and procedures.

I: That actually leads me to my next question. What exactly is the natural phenomenon that formal semantics seeks to model?

S: Yeah, that's the question "what are the data?" It used to be like in linguistics that people mainly used intuitions as their data points. You look at a bunch of sentences and you ask yourself "does one imply the other?" or "do they exclude each other?" etc. Then you talk with other people about it and they say OK. So the data are your own intuitions sort of verified by checking with intuitions of other.

I: Sorry, I think my question was slightly different. It's not about how they're checking it, but what exactly is the goal? What exactly is the formal semanticist trying to model? Is it cognitive processes? Or language? (Even if we say language, how do we disambiguate that)?

S: I think in fact people are modeling these data — either intuitions or corpus data or experimental data and stuff like that. The idea is that these data reflect properties of natural language. Natural language is then a somewhat idealized concept. The question of what natural language is is actually philosophically a very interesting question, because what we observe is linguistic behavior and the products of that behavior — e.g. what we are recording now and what is written down. That is what natural language is, but from the point of view of modern linguistics it's a more abstract type of object. That nevertheless has an important cognitive counterpart — namely the ability that is characteristic of competent language users. I think most formal semanticists are still knowingly or unknowingly subscribing to that conception of what it is that they are modeling.

I: The second conception?

S: Yeah. Natural language or linguistic competence in this more abstract sense.

I: Not in the sense of what we are recording now or actual linguistic behavior?

S: No. But since they are increasingly using that [recordings / actual linguistic behavior] as their data — the move away from intuitions to more behavioral data creates a tension between these two conceptions, and that is something that is very much happening now. The outcomes of that are still sort of undecided.

I: I guess we've already been talking about this, but typically in formal semantics how do we check whether the model is actually a good model for the thing that we're trying to model? And how do we check the accuracy of the models?

S: It used to be simply intuitions and then people started to incorporate data that they got from corpora and data that they got from doing experiments (asking people about judgments and stuff like that), so now it's very much mixed. One of the interesting questions also with regard to what it is that we are modeling is, of course — can we simply combine these data that you get from running a very controlled experiment and what you get from looking at language in the wild (what you get from these large corpora studies) and your intuitions. To what extent are these data comparable? An intuition based description always aims for consensus, whereas if you go out and ask people you never get 100% scores. There are always people who think it means something else, etc.

I: What exactly are the standards for success and failure in this enterprise of modeling? What distinguishes a successful model from a failed model?

S: It should get the facts right that it wants to account for. Say somebody describes the interaction between tense and modals in a certain way, but cannot account for the fact that in a certain setting there is a scope ambiguity apparently regarding the intuitions — if they don't get that then that shows that the analysis is not completely successful. If it gets all those facts right (whatever those facts are taken to be) then it is considered successful. Then, of course, there are additional considerations — can you combine it with other successful models? This "larger whole" thing that it should fit into does play a role, although it is never really systematically checked, and is just left to the readers so to speak to come up with those kinds of objections. And there maybe

questions of elegance — some people are very much concerned with, say, the algorithmic complexity of the tools that you use and a lot of people don't care. Some people might object to, say, the use of possible worlds because they say "philosophically I don't know what a possible world is" and then the average formal semanticist would say "Whatever; if it gets the job done, we'll use it. After all it's just a tool." There you see the modeling and the tool perspective being a convenient space to wiggle.

I: What role, if any, does empirical evidence play in this process?

S: Well, the empirical evidence is the intuitions. The intuitions are taken to be empirical evidence. If a competent speaker of English says that A implies B, then your description better make sure of that.

I: By intuition you mean the intuition of the researcher himself?

S: Well, officially it's the intuition of competent speakers. As a researcher if you are not a competent speaker (a native speaker) of that language then you should check with native speakers, and sometimes you do that, because the distinctions can often be very subtle. You can use informants and their intuitions are the relevant empirical data, or you can look at what you find in a corpus and what you get from experiments.

I: Then I have some questions about the state of formal semantics. Do you think that formal semantics as a discipline has been successful thus far in achieving its goals? Why or why not?

S: I would say yes and no. Yes if you take the goal to be to come up with descriptively adequate ... [[*interruption*]] ... Successful in the sense that it gives you very interesting and rich descriptions that reveal interesting aspects of ranges of phenomena, but I don't think that it has been successful (this depends on what you take linguistics to be about) if you take formal semantics (or linguistics in general) to be an explanatory theory about linguistic behavior — then I think it's much less successful. Because it really doesn't have very strong links with what we know goes on in the cognitive dimension and it really lacks any connection with what else is involved in linguistic behavior — namely the social and physical environment. Language is an extremely complex phenomenon, and you can argue that you can't even expect there to

be an explanatorily adequate theory of language because it's simply so multi-dimensional, and there are so many different aspects that need to be described and explained in different ways that you can't have an overall theory. Formal semantics is sort of contributing to that, but it is not an explanation of what natural language meaning is in that broad behavioral sense.

I: However, in the paper "Abstractions and idealizations" you and Michiel have invited the reader to consider the possibility that formal semantics (and modern linguistics more generally) could be an example of a failed discipline. What is the notion of failure that you had in mind there?

S: Yeah, that was maybe a little too provocative, but... Well, the notion of failure that we had in mind there was this: if you take linguistics on it's own word — linguistics is descriptively adequate and explanatorily successful account of natural language — then it hasn't been successful, and it hasn't been successful because, as I just mentioned, it doesn't provide straightforward hypotheses about, for example, the underlying cognitive or even neural functions; there you see that the connection is actually quite weak. It also has not been very successful in computational applications. The best theories out there that deal with say translation, search, and other forms of natural language interfaces are by and large based on stochastic methods and not on the kind of rule based methods [of formal semantics]. Of course you can try to incorporate them into hybrid approaches, but linguistics according to it's own self image hasn't been very successful there. Of course there are lots of sciences that are still progressing — that doesn't mean that they are failures. The failure is that a lot of these limitations are sort of self-imposed and very difficult to get away from. That's why this notion of an idealization is relevant in that argument.

I: My understanding of your argument in this paper was that you were criticizing formal semantics as potentially a failed discipline because it does not meet the scientific standards for empirical verifiability or falsifiability because according to you the claims that formal semantics makes are about these idealizations and therefore they cannot be compared to any real world data, because the data we have are about the actual phenomenon (and not about the idealization).

S: Yeah, that's true to a certain extent, but if you formulate it like that then you say that linguistics should be about linguistic behavior; actually I think it should, but it [modern linguistics] is actually not about that — it's doing something else. The something else that it is doing is not something completely arbitrary with no connections to empirical reality — it has, but it is not able to provide a good account of how what it does deal with is connected to what it should be dealing with according to us. In that sense it's a failure. The most important distinction between an abstraction and an idealization according to us is that if you do an abstraction you still retain the thing that you abstract away — you say "Let's forget about it for the moment, but not for good." So you always in principle know what you should be doing in order to no longer make that abstraction. In an idealization you simply forget about that, and that means that you need something else to translate your results back into results that are relevant for what you idealized away from, but linguistics doesn't contain its own bridging theory. In that sense, yeah, it's a failed discipline. And then there's lots of ideology also involved. Many of these idealizations are driven by a very specific conception of what it is to be scientific. This is basically modeled on the natural sciences. You can argue (and many people have argued) that what linguistics deals with — human being — cannot be treated in that way. Maybe aspects of it can be treated in that way, but not as such. In that sense it's also a failure, because it simply starts from the wrong model of what it is to be a science.

I: Can you maybe expand on this a little bit? What do you think it takes to be a science? And why does formal semantics not meet that standard?

S: This is something that is ingrained already in people doing linguistics right from the start — it's the idea that linguistics should follow the lead of the natural sciences and provide law like generalizations, and what are often called hypothetical-structural explanations (explanations in terms of postulated structural features), and is concerned with language and linguistic ability basically as they are exemplified and manifested in individuals. So it is committed to individualism in a very strict sense, and it is committed to a particular form of explanation which works fine (actually not completely) in the natural sciences, but it is very doubtful whether that actually is a fruitful way of looking at what linguistics should do with its object. Physics for example deals with physical

reality (which is complex enough), but minimally it has the advantage of being a more or less homogenous domain where reductions of complex to more basic stuff actually makes good sense, but the domain of linguistics is much more varied and is very heterogenous. It is of course concerned with properties that reside in the individual (we have certain basic cognitive structures that are conducive to the development of language — that's true), but it also has a social dimension. We have language obviously for the purposes of (broadly) communication and interaction, and to establish and maintain social cohesion. The way we use language is subject to broadly societal forces and so on and so forth. What things mean is not something that can be ascribed to an individual (division of linguistic labor, etc.). So, it's really a very complex and heterogenous phenomenon (or field of phenomena). To think that you could get that all into the mould of one explanatory model is really a misconception.

I: If a scientific discipline like physics is a wrong role model for formal semanticists, could you suggest perhaps what would be a better role model?

S: David Marr famously made this distinction between three levels of explanation, but he also made (and that's maybe less famous, but nevertheless very interesting) a distinction between what he calls type 1 and type 2 theories. That really has to do with this issue of the heterogeneity or homogeneity of the phenomena that a theory is trying to capture. He says that although some theories deal with very complex phenomena, as long as all the various aspects or dimensions that are involved are relevant for all the manifestations that the theory wants to capture, it is in principle possible to come up with one comprehensive theory. It will be complex, and we may never get there, but it is in principle possible. But if you are dealing with a range of phenomena which has a similar kind of complexity but not every element is relevant for every dimension then you are basically dealing with something that defies the possibility of coming up with one comprehensive theory. You have a field of phenomena where you say "this part of it is explained in this way and that part of it is explained in that way, and there is no way that these two kinds of explanations can somehow be unified." That seems to be the reality of areas like economics, linguistic behavior, and some people have argued that even in biology you are confronted with that kind of predicament. So if that's true, to go back to your question, then formal semantics shouldn't try to be that one big theory that

arguably doesn't exist. It should be satisfied with doing what it does best — namely, coming up with these nice, consistent, interesting descriptions — and say "This is what we have to contribute broadly to the understanding of the phenomenon; so we drop the pretense of coming up with something that is the most fundamental or completely explanatory. If we abstract away from lots and lots of things, then this is what we're left with, and that we can explain in this particular way." Then the measures of success and failure can be viewed from the proper perspective.

I: I was going to ask you what, if anything, could cause you to change your mind about your criticism of formal semantics, but it seems like if they were to change their expression of the goals of this discipline, that would address a large part of your criticisms?

S: When I talk about these things, my perspective here is basically Wittgensteinian. In the Philosophical Investigations Wittgenstein says "Whatever philosophy is, it's not something that interferes with actual practice." He makes that very explicit in his remarks on the philosophy of mathematics, and says that we're not going to try to change the practices of mathematicians, because they're perfectly alright as they are. Logicians and philosophers shouldn't be trying to meddle with that [with the practice of mathematicians]. Then he generalizes that to philosophy — basically it doesn't produce anything new or doesn't want to change things, but it just wants to come up with a better picture of what it is that things are and how things are. That's also characteristic of my way of trying to answer these questions. I don't want to change formal semantics — formal semantics is perfectly fine. Of course there can be internal criticisms, and I think that for example this issue of comparability of data that comes from different sources is something that formal semanticists really should think about — that's an internal matter. But apart from that, I just want to know and to understand better what it is that they are doing. If they find that interesting, then so much the better, but if I'm the only one who finds it interesting then...

I: My final question to you is: do you think that formal semanticists themselves would agree or disagree with the views you have expressed in this interview?

S: Well, I know for a fact that many of them don't agree.

I: What do you think are the main points of agreement or disagreement?

S: I think that the main points of the disagreement ... this is an ongoing process ... as I said, I was at this workshop where I gave a presentation on Saturday, and I said that this hypothetical-structural model of explanation may not be the best one (it's actually also not the one that plays a role at the very fundamental level of quantum theory). Maybe we should step back and look at language and the ability to use language more from a dispositional point of view. Then people started to react to that saying "There must be underlying structural properties." I said "Of course there are underlying structural properties that play a role in any disposition, but can the disposition be reduced to that?" Then people said "Of course, it must be like that!" I said "Well, why? Why must it be like that?" Then you run into these very basic, almost unaware, views that people hold that define for themselves what it is that they're doing. Then you get these references "But that's how it is in science." My first question: "And why should it be the same here?" Second observation: "That's actually not how it is in science; but forget about that — the first question is the most important one." So I think many of the objections come from a certain unwillingness ... well, it's hard to reflect on your own presuppositions, and it's not always a very comfortable thing to do, and it's not obvious why you should do it, and when somebody comes along and starts to nag you with these questions you get irritated: "Go away! Let me do my thing!" And I think that's a perfectly fine reaction.

I: And what are the key presuppositions in question here?

S: That explanation should be in terms of structural properties, for example. I think that's due also to the way in which people have been trained — you get your philosophy of science 101 where we say "Of course we have dispositional properties of physical objects like fragility, but these are explained ultimately in terms of structural properties like surface tension, etc." "If they get rid of dispositional properties in physics, then we should get rid of dispositional properties in linguistics." This holds for many humanities disciplines and even for some of the social sciences. That's so very much ingrained, that people feel uncomfortable with it being questioned.

I: Do you think there are any other underlying reasons for which this disagreement exists?

S: Like what?

I: Perhaps a different philosophical outlook on the world, or a different view of language, or something like that?

S: Yes, that's certainly true. Most formal semanticists nowadays are trained as linguists, and they are trained to look at the object of linguistics in this fairly idealized way. So again, if someone comes and questions that, then people think "Isn't that obvious?" Of course there is a lot of critical discussion of linguistics from a philosophical point of view, but people say that that's just philosophy and we don't need to bother with that. I happen to be also a formal semanticist — at least I used to be one. Maybe that's also what irritates people, but I don't know about that — I can't answer that question really.

I: Those are all the questions that I wanted to ask you. Is there anything else that you would like to add before we end the interview?

S: [[*long pause*]] No, I think we covered the various dimensions and aspects of what you are concerned with. I think (but this is just a thought for maybe a possible future) doing an ethnographic study of formal semanticists at work could be very interesting. How do people actually do their research and discuss their research amongst themselves?

I: Yes, I hope to provoke some interest at least in this direction.

S: I think this is a very interesting project, and I'm very happy that you're doing it.

I: Thank you very much.

A.2 Interview with a critic

I: To begin with, I would like you to briefly introduce yourself. Can you please describe your academic and professional background and how you got interested in the field of formal semantics?

S: I have a rather checkered history, and I've worked on many topics — originally on a topic related to probability theory and mathematical logic; I worked on this until 1993. Then I worked in artificial intelligence for some seven years, and that is actually what got me interested in formal semantics, because I was studying mathematical theories of vision and I was trying to link up these theories with the semantics of perception reports. That first aroused my interest in applying sophisticated mathematical models to do semantics, so the result ultimately was a book together with Fritz Hamm — "The proper treatment of events" — which applied logic programming to tense and aspect. Later, the psychologist Keith Stenning and I wrote a book applying non-monotonic reasoning as a model for human reasoning behavior.

I: What was your primary motivation for studying formals semantics? Was it part of your attempt to understand human reasoning behavior? Or something else?

S: No, it was very different. It was a case of "I have the tools; now can I find a problem to solve?"

I: Would you describe yourself primarily as a linguist, or as a logician, or as something else?

S: I'm definitely not a linguist; I'm more the formal person, but I've been lucky in collaborating with people who have all the linguistic details at their fingertips. That holds both for Fritz Hamm, and later — my student and ultimately collaborator — Giosue Baggio. with whom I did experiments with EEGs based on predictions by the formal model.

I: OK. So that's about your background and motivation for studying formal semantics. Next I have a few questions about what the nature of formal semantics is according to you. To begin with I want to ask: What, according to you,

is formal semantics? For example, how would you describe it to a university student who has never heard of this field?

S: You mean, you want to hear my own idiosyncratic definition?

I: Sure. Like, if there is a university student who is interested in exploring a new field, and he finds about formal semantics, and he asks "What is this field?" — what would you say to him?

S: For me formal semantics is strongly linked to cognitive science, and ultimately the most interesting work done in that area is the work that gives some insight into language processing (both producing and comprehension).

I: What do you think are the goals of this enterprise of formal semantics?

S: I already announced that goal — in the sense that I said I would like to have some insight into language processing. So, the deeper level indicated by your question is "What is this insight?"

I: Right. Or "What kind of insight?"

S: I would say that one has achieved insight when one has a formal model that predicts a variety of behavioral data.

I: You originally said that the goal was to understand language processing, which suggests that the goal is to understand the cognitive process, but in terms of the predictions of the model you are saying that they should be predictive of behavior?

S: Sorry, I'm using the lingo of the field. I should say just *data* — so, let's run through the variety of data that you have when you study language comprehension. (1) There's eye tracking which traces the movement of the eyes while reading, and that can give information in case backtracking occurs. (2) Then there are ERP data which is a kind of average form of EEGs taken when a subject is processing linguistic material. (3) Then there is what is properly behavioral data — namely people's answers to verbal questions. I haven't exhausted all the data types, because there is also fMRI data (which personally I find less interesting because it gives you less information about the time course or events). To come back to your question of what is the goal — the goal is to

collect a whole raft of data of these very diverse types, and to devise a model that has something to say about all these data types, which is such that it predicts what is observed.

I: Can you be a little more specific? Could you give some examples of the kind of problems you try to solve using these models, or the kind of research questions that you would try to answer?

S: Yes, let me give two examples. One example is a border line case because it has to do with psychiatry. We were interested in the reasoning patterns of people with autism. So, answers to reasoning problems were counted as behavioral data, and we had a hypothesis based on a certain formal analogy between a reasoning task called the suppression task and a non-verbal task having to do with flexibility of behavior. The formal analogy was such that we could predict that people with autism will behave very differently on the suppression task than neurotypical people. We did this experiment, and the hypothesis was confirmed, and the formal model that we used had some ramifications because it could also be viewed as a neural network and as such could be linked up with research on neural abnormalities in the brains of people with autism. So the goal here was to come up with a model that has many ramifications — not just a model that explains one isolated phenomenon.

I: OK. I would like to ask you some more questions about how exactly this process of modeling works in formal semantics. Let me tell you what I mean by modeling (or how I understand it) — logicians are often concerned with modeling in the sense that they try to create formal representations (i.e. models) of some real world phenomena in order to represent them and to reason about them. Do you think that this notion of modeling plays some role in formal semantics?

S: I'm not sure; somehow I seem to miss the idea of prediction in your description. Have I misheard?

I: No, my notion [of modeling] was to represent and to reason about some real world phenomena.

S: Yes, I would think that that's definitely not enough.

I: OK. So maybe you can say what your notion of modeling is?

S: The models should agree with what is known about the process it is modeling, and with other aspects of cognition, and it should give predictions about as wide a range of data types as is appropriate. By this I mean not just the so-called appeals to intuition, but also corpus research or neural imaging data.

I: What exactly are the phenomena that these models are trying to model?

S: Here I have to take a stance on what happens when language is processed. I believe in an account of language comprehension which is the psychological form of discourse comprehension (it's called situation models), which says that we always have a model of the discourse encountered so far, but the model is much richer than the verbal material. Meaning that causal relation will be encoded (even though they are not overtly present in the material), expectations for future events will be encoded, and in general all the information that the semantic memory has available will be stored in the situation model. So, what a formal model has to do is to give a description of the computations that are going on in that situation model.

I: The computations that are going on in the brain?

S: Yes.

I: So, according to you then, what the model is actually trying to model is the brain processes that are involved in language processing?

S: Yes. You can situate this question at different levels: (1) You can talk about the function that has to be computed when you consider the process leading from a discourse to a situation model. (2) You can ask about the algorithm that executes this function. (3) You can ask about a neural implementation of this algorithm. Ideally a formal model should address all these layers, but in practice you are often happy when it addresses only the first of these levels.

I: What is the first level?

S: That you specify a function which tells you how discourse is mapped onto a situation model.

I: And the deeper levels would be...

S: The algorithmic level — when you write down the algorithm where you write down all the steps that have to be taken to actually do the computation to implement the function. Then the last step is to associate the algorithmic steps with neural computations.

I: OK. So, given this notion of modeling, how do we define success and failure in modeling? What distinguishes a successful model from a failed model?

S: That's a thorny issue. A straightforward answer would be that a failed model is one which leads to a failed prediction. Now it doesn't always go like this, and falsification in the sense of Popper is dead. You don't often falsify a theory completely; it's often possible to make ad-hoc changes.

I: Are you using the word theory as interchangeable with model?

S: Yes, good point, so far I do. Yes.

I: OK.

S: I'm trying to think of an example from my own research where something definitely failed ... In the autism research ... I told you a while back about the experiment with the suppression task (which is a kind of pragmatic inference task) — people realized that no information may make a previously endorsed inference invalid. We concluded from this that there might be a general pragmatic impairment in people with autism, which should show in their incomprehension of Gricean maxims. So for example, you have a sentence like "a tree has branches or leaves" and that should be considered nonsense, and we thought that people with autism would behave differently from neurotypicals in that they wouldn't call this sentence nonsense. There the prediction failed dramatically — the people with autism behaved exactly as the neurotypicals.

I: So, the reason your model was considered a failure was because your data (of the behavior of the subjects) did not match the predictions of the model?

S: Yes.

I: At some basic level then, success or failure of the model depends on whether the predictions made by the model are confirmed by the data or not?

S: Yes, that's one reason to abandon a model or a theory. Another reason may be that the theory proposed is too distant from neural reality, so there's no way you can envisage an implementation.

I: What role does (1) empirical evidence and (2) intuition play in this process of modeling?

S: Good question. Let me think for half a minute ... I would say it is neither. What drives me personally is the conviction that mental processing can be described by and large by mathematical models. When it comes to finding the right models we look in slightly greater detail at how cognition operates, and for instance the models that we used to predict neural imaging data or the behavior in the suppression task — these models derive from the hypothesis that language is ultimately an outgrowth of the human planning faculty, and there are some biological reasons supporting this; that is the empirical component. But there is another component — the a priori component — which is neither empirical nor intuition, but is just the conviction that there must be mathematics behind it.

I: But how can this conviction be tested against reality?

S: If a model is any good it makes predictions, and nature is generally unkind, so predictions get falsified. That's one way, and next to the neural implementability issue there's always the issue of scope. It doesn't make sense to devise a mathematical model which does what it should do on the suppression task (say) and is useless anywhere else.

I: So, in some sense, the model that we are seeking should be a unified model? In the sense that it explains not only discrete phenomena, but it explains various phenomena in a unified manner?

S: Yes, and that will happen automatically when the model has its roots in some known biological properties of the brain, which in this case was the observation that important language areas of the brain appear to have migrated in the course of human evolution from an original position in the planning center.

I: OK. Then I would like to move on to some questions about formal semantics as a discipline. We talked a bit about the goals and the methods of formal

semantics. Now I would like to ask you if you think formal semantics as a discipline has been successful thus far in achieving its goals? Why or why not?

S: The presupposition of this question is buried in the phrase "its goals". I think there is no uniform set of goals for formal semantics, so it's very hard to speak of success or failure. You may have attended the discussion meeting at the Amsterdam colloquium, where Angelika Kratzer said that the success of formal semantics can be measure by the amount of young researchers that work in the area. That is not exactly an intellectual criterion (to put it mildly). I think that the goals of formal semantics are unclear, and as far as I can see, too narrow in scope. There is some testing of the model against data, but it is usually all of the intuitive kind.

I: What do you mean by "the intuitive kind?"

S: The data come from peoples' intuition.

I: In your paper "Abstactions and idealizations... ", you and Martin have invited the reader to consider the possibility that formal semantics (and modern linguistics more generally) might be an example of a failed discipline. What is the notion of failure that you had in mind there?

S: Various notions. One of the most striking failures I find is Chomskyan linguistics, which every couple of years seems to take on a different form. That is not a good sign. I also think that it lacks a set of central research questions as you have for instance in physics — everybody wants to unify the basic forces of nature, and everybody wants to know what happened in the first few seconds after the big bang. I cannot see that the identity of formal semantics is defined by a common set of research questions, and as such I find the discipline a bit incoherent.

I: Do you think that it is a rigorously scientific discipline?

S: No, because for me rigorously scientific means that there are agreed criteria for when to abandon a theory, and I don't see this anywhere.

I: What do you mean by "theory" here? A particular model?

S: Say, a proposal from the Chomskyan school for a particular grammatical analysis.

I: But when I asked this question earlier, we seemed to agree that one criteria for discarding a model or a theory is when we have some data that refute the predictions of this theory or model.

S: Yes, but I thought you were asking me a general question. What I doubt is whether that criterion is at all operative in formal semantics.

I: OK. Generally when there is a claim that "x has failed", this requires a postulation that x had some goal and an observation that x did not meet this goal. But this is not the way you would look at formal semantics?

S: No. Maybe it's even the proliferation of goals that make it a failed discipline.

I: The proliferation of goals?

S: Yes. When I go to a semantics conference, I'm always struck by how many different topics are treated with so many different techniques, and how each speaker seems to work in isolation and not as part of a program trying to achieve some universally recognized goals. That is a very odd situation. If you go to other fields like, say, developmental psychology, then it's very clear what people want to achieve, and there is much mutual comprehension and mutual criticism because of the shared goals. At these semantics conferences there is some polite interaction after a lecture, but there's no evidence of a passion for something you want to achieve.

I: So you think it's very fragmented, and there is no unified manifesto or statement of goals for formal semantics?

S: Yes, that's a fair summary.

I: I would like to then ask you if you think that practitioners of formal semantics would agree or disagree with the views that you have expressed in this interview? And what do you think would be the main points of agreement or disagreement?

S: I think they would mostly disagree and take Angelika Kratzer's side.

I: Which is?

S: That as long as formal semantics keeps attracting intelligent and enthusiastic young people the subject is thriving.

I: OK, but what about your notion of what it is that formal semantics should be modeling? You have this idea that it should be about modeling cognitive processes in the brain — that's not a view that is very common I think among practitioners of formal semantics.

S: No, it's extremely rare, and it's a very personal view. Personally I can't see what insights other approaches to semantics give. But I must conclude from what I see and hear that most practitioners of formal semantics must have an idea of insight which is beyond my comprehension.

I: OK. Do you have any idea what could be the underlying reasons for this disagreement about the discipline?

S: Yes. There are I think sociological processes at work. The subject started out in the 1960s in a rather passionate manner where the clear goal was defined by Montague grammar. As time wore on more goals appeared and disagreements arose, but subjects very seldom simply die out. Just to give you an example of how we filled vacancies here — you know that Martin, Jeroen and Frank retire. Frank has already retired and Martin will follow a year from now. It could have been decided to reallocate the open positions to some different fields, but it was felt that the tradition of formal semantics in Amsterdam was so important that the positions need to be allocated to formal semanticists. And it's those type of decisions that keep a subject going, but it doesn't necessarily lead to intellectual coherence.

I: OK. The last question I would like to ask you is: what, if anything, could cause you to give up your criticism of formal semantics? For example, what kind of change could happen within formal semantics that would lead you to retract your criticism?

S: That's a very difficult question. My immediate inclination is to say if it is shown after all that the structures that they use bear some relation to how the brain operates. Most of the formalisms in use were set up with no idea of a connection to cognitive processing, but it's not impossible that these formalisms turn out to have some relevance (as has been claimed for generalized quantifier theory for example). So, if it turns out that with these formalisms, you can after all do the cognitive investigations that I would most like to see, then I change my view and think it will be successful after all.

I: Those are all the questions that I wanted to ask you. Is there anything else that you would like to say before we end this interview?

S: I may come back to you. During this interview there were some half-baked thoughts that interfered with my other answers, and they may gestate later.

I: Thank you very much.

[[*After switching off the recording device, we chat for a couple of minutes about Michiel's collaboration with Martin. When I remarked that I found their collaboration to be rather unexpected, he asked me to switch on the recording device.*]]

S: Well, it's more complicated. Martin and I do have different views as regards practice. What I've been telling you about is the set of ideas that motivated me to do semantic work. At the same time, Martin and I are engaged in a foundational critique which also affects the way I used to do formal semantics, but that critique has not sufficiently crystallized into a new set of practices. Until then I keep to the set of goals that I had. It might well be that after talking for another year, we will be able to transform our ideas of what language is into a different set of research questions — different from what we think now and different from what other people are doing. So it may sound a bit schizophrenic, but I have a set of research goals which makes me operate fairly efficiently in research, but these research goals are motivated more by wanting to learn something about the brain than learning about language. The foundational work is Martin's — specifically directed towards learning something about language. There the research goal may differ.

I: You're saying that although your motivation is more about [understanding] the brain and cognition as opposed to language, this has no impact on the foundational critique of formal semantics? Because that [critique] was conceding the methodology and foundations of the discipline.

S: Yes. Where Martin and I agree is that the set of research goals that I have tell you very little about language as a social phenomenon, and we agree that that is one of the most important determinants of language. Only we have not yet hit upon the proper way to study this scientifically.

I: But you think that this is probably not the proper way to study it scientifically?

S: What do you mean by "this?"

I: Formal semantics as it is being practiced today.

S: That's definitely not the right way.

I: OK. I don't have any more questions.

S: This was helpful — this last bit.

Appendix B

Interviews with the insiders: formal semanticists

This appendix (Appendix B) contains transcriptions of the interviews conducted with the formal semanticists: Jeroen Groenendijk, Floris Roelofsen, Frank Veltman and Katrin Schulz. The interviews have been transcribed using the following schema:

> I: Lorem ipsum dolor sit amet, consectetur adipiscing elit. Aliquam ultricies lacinia euismod.
>
> S: Nam tempus risus in dolor rhoncus in interdum enim tincidunt. Donec vel nunc neque.

"I" denotes the interviewer, and "S" denotes the interview subject, i.e., the person who is being interviewed. The interviewer ("I") is in all cases the author of this thesis, while the subject ("S") is an anonymous member of the set {Jeroen Groenendijk, Floris Roelofsen, Frank Veltman, Katrin Schulz}.

NB: These interviews are semi-anonymous in the sense that the individual subject of each interview remains undisclosed. Therefore the reader is requested not to identify any particular interview with any particular individual in any future citations of this work.

B.1 Interview with a formal semanticist

I: To begin with let's start with a brief introduction of yourself. Can you please describe your academic and professional background, and how you got into formal semantics?

S: How I got into formal semantics ... that's a long time ago, so I don't remember anymore! I just was interested in the more formal aspects of philosophy from the moment I started studying philosophy (which is a long time ago), and philosophy of language and linguistics was my main interest and my main topic. We (when I say we it's usually me and Martin Stokhof, because we cooperated for a very long time, wrote a dissertation together, etc.) ... when we were studying, Montague grammar started to arise, and it was very new ... applying logical methods in the semantics of natural language. We studied that from the very start, when it began, and also later on we had courses about it, and we started to do work ourselves in that framework. In our thesis work — that was about formal pragmatics actually from the start — the idea was that Montague had formalized semantics and that we should do something similar for formal pragmatics (i.e. Gricean pragmatics). So we looked at the Gricean maxims, and we tried to formalize those; the things that have to do with informativity, etc. are not so difficult (that's basically entailment), but the hard part was what to do with relevance? Then the idea was to formulate the relevance of an utterance in terms of answering a question — it was relevant if the utterance was an answer to the question that was sort of the question in the discourse at that moment. So then we had to model questions, which is also important for philosophical reasons, because the Montague grammar is basically about assertions, but natural language also obviously has questions as at least a second major category. The whole logical semantics [of Montague] is in terms of truth conditions and entailment, and obviously questions are not true or false. So you can get the idea that formal semantics will not be applicable to natural language, because it [formal semantics] is only dealing with truth conditions, and there are very many other uses of language, and it [formal semantics] has nothing to say about that. So then the idea was to show that with the standard methods of logic you can just as well talk about questions and what they denote, and question-answer relationships, and that there are

also entailment relations between questions, and that you can do formal things (in the same formal way) as with declaratives ... So the whole thesis was about having a new model for questions in a logical grammar.

I: What was your primary motivation for studying formal semantics?

S: That's difficult ... The main motivation ... My primary interest was in language and how it works (even before I started studying philosophy). What you see is that these logical tools — they give insight into what's happening there, and that you can actually formalize and calculate with meaning and with language. That just sort of interested me.

I: Would you describe yourself primarily as a logician, or as a linguist, or as something else?

S: [I would describe myself] as a semanticist, and I think that's something in between being a logician and being a linguist. If I would be a full blown logician I would also do things like making interesting proofs concerning systems, and if I were a linguist I would be very much interested in empirical data, but neither the proofs nor the empirical data concern me very much. I see myself as building conceptual frameworks — logical frameworks — which are of course stimulated by certain issues that are there in the empirical field, but the data you need for that are usually just some small, very basic, characteristic things that you want to have a model for. That's a real step away from the existing stuff — you have to design something new. Of course it has to be logically OK (so you do have to study the logical features of it), but the real logical work is done by logicians, and I'm not doing that myself. You have to set up the semantics in a way that the logician will be interested in looking at it and making a proof system for it, and, of course, you try to put it in a way that the linguists can understand what you are doing, and that they can further apply that logical work to more detailed stuff than what you've done yourself in your framework building. Questions is one big example in my own history. The other big thing (that I also did together with Martin) is dynamic predicate logic. To sketch how this works: Montague grammar is very much a compositional theory of meaning, but then in the early 1980s there came alternative theories — concretely that was Hans Kamp's "Discourse Representation Theory" (DRT). There were some empirical things there too — stuff like the donkey sentences

and how to get anaphoric relations right there, or just anaphoric relations across sentence boundaries. So the DRT gave a framework to deal with those things, but it very clearly said of itself that it was not a compositional semantics — it was essentially representational. So what motivated Martin and me was that we didn't like the fact that it wasn't compositional, and we had a feeling that there must be a compositional treatment of these phenomena, and so we should stay in this methodology of Montague grammar (of which compositionally is the core methodological principle). So, try to get a new framework — a logical framework that can deal with the same stuff that DRT can deal with, but in a compositional way. So that's sort of how it works. Partly it's philosophy — it's a methodological thing, and then you have to get the right idea of getting a formal system that is compositional and that can deal with the phenomena, and then you're done (well, it took a couple of years to get it really done!). This sketches how I look at my task in the whole business. You will typically have just two sentences: "every farmer who owns a donkey beats it," and "a man is walking in the park, and he is wearing a hat." You don't have an enormous database of examples there. That's enough to motivate what you are after — what your formal framework is going to do, and then after you have made it you can look at other things and then see what predictions it makes.

I: That was about your background and motivation, then I have some questions about formal semantics as a discipline. To begin with, how would you answer the question: "What is formal semantics?" Suppose you had to explain it to a university student who has never heard of this discipline.

S: Then I would say it's a theory or an approach to meaning where you use the tools from logic (and build logical tools — it's not only using tools, but also building new logical tools if you need them) to get a detailed and formal explanation of the workings of meaning.

I: I guess you already said this, but more explicitly, what are the goals of formal semantics?

S: Developing a formal analysis of aspects of meaning. I don't have a big whole [picture] like "there should be a complete grammar of those and those languages and dealing with those and those constructions that can be used in applications for blah blah blah." Of course, you have these things roughly in

your mind; when I make something it has to be of the nature that in principle it can be used, say, in building computer systems that can understand language and that can do some technical things, but I am not *working* on that (that's not what I'm *doing*), but I do want to do things in a fairly explicit way such that in principle that is possible. I'm not really steered by those issues ... I have my own agenda ... whether they can really use it or not — that's not my business, but I do want to do it in a way such that in principle that is possible. It's just my experience that I myself — when I do it — I have the idea that I have achieved something, that I have got new insight into how language works. Of course, you could do it in thousands of other ways, and that's also fine, but you shed light on questions that have to do with meaning, and that's all — I just like it.

I: What are some examples of the kinds of problems that researchers try to solve in formal semantics? What are some of the research questions that they try to answer?

S: I mentioned them already a little bit. For example, I'm still very busy looking at questions (the inquisitive semantics project). What is steering me there? At a very basic level we are using logical tools and logical frameworks. Logical frameworks are designed traditionally for something very specific — for validity, argumentation reasoning, proof, etc. Then the observation is that argumentation may be *a* function of natural language, but it's not the primary function of natural language — the primary function of natural language is communication, is information exchange. Then the rough idea is that your notion of meaning should not be steered immediately from argumentation, but it should be steered from information exchange. Try to develop a notion of meaning — it has to be formal, logical — that is directly motivated from *this* function of natural language i.e. the communicative function rather than the argumentative function. Then see what new questions for logic you get, etc. It's a very big question, and of course you will need details ... you will typically start with the simplest thing that you can think of, so you are not immediately looking at natural language — you are looking at propositional logic ... If you have researched that, then you do first order logic, then if you research that you can try to generalize it to a full blown type-language that has some resemblance to natural language ... that's how it goes. So, you first get your

basic ideas for the semantics, and then you get an interpretation for this simple logical language, and then you start logical work on it — you study the logical properties of the system, and then you try to derive a proof system based on this, etc.

I: Are there any kind of "big problems" that everybody in the field is interested in solving, or some kind of major research questions that are shared by everyone in the field? Or is it just these separate projects?

S: There are many many things that are happening over the past decades in semantics for which I find this [inquisitive semantics] to be crucial and central. There are many things in pragmatics and in semantics that turn around this discussion of questions. All sorts of phenomena are being studied where this notion plays a role. Another is called "alternative semantics" — that's partly about disjunction, also about existential quantification, where people argue that it has to be analyzed in a slightly different way, etc. There are all sorts of small problems that are hanging around, so there's a lot of research that's done around these things — like focus, intonation, etc. What I then see is that what is lacking here is a formal system that is basically tuned towards these issues. People start making sort of small amendments to the standard systems — changing things a little bit — but there is not something that in principle starts from this idea of questions ... that that has to be modeled, and that that has to be a really hardcore element of your semantics ... that has to be at the core — not something that's built on top of it [your existing system], but it has to be in the center of things. You try to provide — at a very central, basic, logical level — new tools which all these theories can use to get a better formal representation of what they are doing, and to make everything more uniform.

I: When is a problem considered to be solved? Or when is a project completed?

S: This type of theory is, like any theory, never really complete. If I look at the past, then it's easier to see "When did I stop? When did I think that I have done my job?" If you look at dynamic predicate logic — well, actually, basically when the semantics for the first-order system was there, then I didn't stop, because that's just a first-order language. You also want to show that you can generalize this thing to a full blown type system. So then we wrote

a paper on dynamic Montague grammar showing how to generalize to a full blown type language. Then I didn't stop completely with it, but that was more or less what I did ... after that I lost my interest so to speak ... other people should do the details ... we just want to show the way. That was not really the full end of it (I am still back again at dynamic semantics), and it will never really stop (until I die). There was this other big issue: Frank Veltman was working on update semantics, and there was dynamic semantics, and these two things are conceptually very much alike. They both go from the idea that meaning is not really covered by truth conditions, but meaning is information change potential. In the most basic version of update semantics, the thing that shows its specific properties is the epistemic "might" operator. The conceptual idea behind dynamic predicate logic, and this might-system were completely the same, but both these things used completely different features of dynamic properties. They each had a particular property — that if you wanted to deal with, you need dynamics — but they were different properties. The might system was sort of classical where dynamic predicate logic was dynamic and [vice-versa]. The three of us worked together for a couple of years to develop a system where you can deal with both at the same time. It was pretty hard to do that, but in the end we succeeded such that it all worked together. You can have things like "There was a man walking in the park. It might be anyone." So you get the "might" and the quantifiers hanging together. So, for me it stops when this general thing is worked out. What you always end up with is that you can do certain things, and certain things that seem related you cannot do. You don't have a story for that. A big problem is always to decide if this is something really different or if it is something tied up. I don't have to deal with that in my basic story — the basic story can be restricted to those things, because for the additional stuff you need something extra at one level higher.

I: Like what?

S: Well for example, what you have in dynamic predicate logic is that the existential quantifier has this dynamic feature, and not the universal quantifier. So if you have universal quantification, then it doesn't leave you with something you can use for further reference. Negation also very much blocks it — for example you have a sentence like "No man is walking in the park." but you cannot call him "him." But there are certain examples where you do seem

to get that type of possibility, like "Everyone chooses a pawn. He puts it on square one." What you typically find in describing the rules of a game. Here you have the "every" and the "he". The basic system doesn't do that, because it doesn't deal with universal quantification. So this is one of the phenomena that you have on the table as you are looking at pronominal stuff, and the system that we have doesn't account for it. Does that mean that it is wrong? Do we need something bigger than what we have? Or is this a different phenomenon? That's a decision that you have to make. The real decision to say that this is the basic system is when you actually know what you need on top of that simple system to get that out also. If you know in principle how to do it (but then it might be really different — it's another layer). The reason of proceeding in this way is that you want to get at a simple thing (as simple as possible) if only because you want to study the logical properties of it. If you make it tremendously complicated from the very start, then you are making life very hard from the logical side. So you want to try and extract the minimal thing from what you are studying and an important motivation for that is because you want to have the logic right. If you put in all the complications immediately, that's not going to work. The linguist will typically work in another way, because he will have this [say] 175 examples ... a corpus of related stuff ... and he wants to build a machinery that can deal with all of that. Whether that machinery is still logically very simple and basic and whether you can prove things about it is not going to interest the linguist. He wants to have a descriptive coverage over a big range, and he is not so much interested in making this conceptual thing that is doing the basic trick. I think that's the big difference, and I think you need people from all these directions to do the job together. Of course, you need the people who collect this interesting data, and you need the people who do the conceptual work, and you need the people who do the logical work, and they should all cooperate.

I: Then I would like to ask you some questions specifically about how modeling plays a role in formal semantics. Logicians are often concerned with modeling in the sense that they try to create some formal representations (i.e. models) of some real world phenomenon in order to represent them and to reason about them. Can you describe what role this notion of modeling plays in formal

semantics? Are you trying to build [these kinds of] models?

S: Yes. There I'm doing what the logician is doing. Yes, precisely that.

I: Are your models also supposed to make any predictions?

S: I find that difficult to answer. Yes, I mean in a way, of course you hope to see new things that you didn't see before.

I: What would the predictions be about?

S: I'm trying to think of ... You are not building the model to make predictions. You just have a conceptual idea, and you motivate the conceptual idea about modeling something. Let's take questions [for instance], because I sketched the questions ... or more generally, you want to have a logical model that can deal with information and with issues in a uniform way — that's what you want to make. The way you get there ... of course, you also have examples that steer you, but basically you have a specific idea about what it would look like. So in this particular field ... in the work that Martin and I did in the late 1970s and 1980s, we also came up with a particular concept of what a question is. A question is a partition of logical space, and the parts in the partition are the complete exhaustive answers to the question. So later in this inquisitive semantics work we came up with a different story of what questions are. That's not something that we saw in the beginning, but we see that more and more sharply now. It's the relation between the question as a partition and the notion we have now. And it just turns out that this notion is more general than the other one, and that you can point out very precisely at the formal level what determines the restrictions of the partition semantics in relation to this new doctrine. Then in terms of predictions (I smile a little bit, because *is it a prediction?*), you get ... to give the crucial example, what we never looked at in the partition semantics was conditional questions — things like "If John comes to the party, will Mary come as well?" It's a very simple structure at the propositional level (you could do it in propositional logic with something like $p \rightarrow ?q$ or something like that), and you don't even meet that example in the earlier work that we did, or even not in linguistic work on questions in general. You do not find conditional questions as a topic. Apparently everyone thought it goes roughly the same ... "we will do that later, but it should not

be different whatsoever." But in the new system it's just a very basic thing — these conditional questions ... also the way implication is interpreted, it doesn't care whether it has a question in the consequent or whether it has an indicative in the consequent ... it's just one rule for implication, and it will just give you the right thing for conditional questions ... no difficulty there. But in terms of partitions, they are really a problem, because conditional questions (if p then ?q) gives two answers: yes (if p then q) or no (if p then not-q). Those two propositions are not mutually exclusive, but they overlap. So you do not get partitions, but in the partition semantics the whole conception is that the question is a partition, and these conditional questions just do not fit. They don't fit in it. So in terms of predictions, the prediction is now that under this more general picture (I haven't given this more general picture, but it's not so important and is not much more complicated than the partition view — just a slightly different conceptualization) you get conditional questions for free. So you are not bound to partitions. Partitions are also questions under this new view, but they are just a special case. So the prediction is that you get a wider range of questions that you can deal with — you have conditional questions, you have disjunctive questions (in the partition view you just cannot deal with them).

I: So more broadly, would you say that it's predictive of the way in which people *use* questions in language?

S: Yes, it just has a better coverage of all the types of questions that are actually being asked. Yes, and I would also say of the function.

I: More generally, beyond this project of questions, in general if we are building models in formal semantics, what are these models actually about? What is the natural phenomenon that we are trying to model?

S: I would say it's the ... Well, let me say what it is not. It's not a model of the mind or something like that. It's also not a model of the real world. It's not that either. What I think formal models always do — they just have a good framework that has a conceptual motivation, and they can explain certain phenomena. What are the phenomena? I'm afraid that for me the phenomena are basically my intuitions. About relations for instance — so "this implies that" or "these two things do not match" and stuff like that. Of course, that's

not the end of the story, because my intuitions are not everything that count, but in the sort of work that I am doing, I need these intuitions that I start from and that I try to account for. Then of course, the rest of the story is to see more broadly how this squares with how people actually behave linguistically. That's of course another story, and then you have to test it there as well, but that's not my work.

I: What do you think are the standards for success and failure in modeling? What distinguishes a successful model from a model that's a failed model?

S: Something fails if you can come up with a phenomenon or example that you typically would expect the thing to do given its set-up, and where it clearly goes wrong.

I: Sorry, I didn't follow ...

S: I've actually just in the past couple of weeks written a big review on a paper for a journal, where this sort of thing happens. That was again about questions (embedded questions), so there are certain rules, and of course a certain basic idea of what questions are (it's about all the words that embed questions ...). What you see at a certain point ... if you just apply the stuff as they described it ... and there are certain cases that clearly fall under the scope of the theory ... I mean you can always find things that are not under the scope of the theory and that it doesn't do; that's not nice. You have to find things that are clearly in the scope of the theory and that come out wrong under the analysis that is provided ... and then you try to repair it.

I: So, basically *counterexamples*?

S: Yes, counterexamples. The basic counterexamples are not things that are not dealt with by the story, but that you would expect to be dealt with by the story but don't work. In this particular case for example, what came out was that if you just take something simple like "John knows who is walking" and then if John had no idea about who was walking, and in fact nobody was walking, then the theory predicted that then he knew who was walking. That I think is clearly wrong. I think you want to say in that situation "he doesn't know who is walking" and not that "he knows who is walking." So there was a mistake somewhere in the system.

I: But these counterexamples — are they drawn from intuition typically?

S: Yes. Yes, I don't have to ask a language user what he would say in a situation like "If no one is walking and John has no idea about who is walking, is it then true that he knows who is walking?" Everyone would say [no]; I'm not going to ask that to a language user. There are other cases where it is not so clear, so maybe I should talk a bit about that so you can see what sort of things for me are the empirical stuff. Where I can decide that in this [case] I'm not going to say what is right and what is wrong, because I don't know, so maybe we should ask the language users and see what is happening. So a clear case is *surprise* ... no, not surprise ... the particular example I was thinking of is *predict* — "Mary predicted who would be at the party" or something like that. Say that Mary has only said "John will be at the party and Bill will be at the party" and for the rest she doesn't dare to make any predictions. She's not saying that they will be the only people at the party. Now, if it later turns out that indeed John and Bill were at the party, and they were the only ones at the party, no one else turned up. Are we going to say in this situation that Mary predicted who were at the party, and she was right? Or are we going to say that it is not true that Mary predicted who was at the party? It's weak vs. strong exhaustiveness — that's sort of the name of the phenomenon. For years and years, there's an argument in the literature about whether such sentences have this [property]. So, they really require that Mary has said something to the effect that "those and those, and no one else will be at the party" or whether it is sufficient for her to have given a list of positive instances and not saying about the list "that's all". If it turns out that the list was completely correct, everyone she mentioned was there, it was not the case that someone she didn't mention was there, but she didn't say "that's all". What I would like to have there is that the analysis that I have allows in principle for both of these possibilities — so that I can in principle account for the ambiguity, and then we can see what the language users say. So, we can try to set up experiments that sort of decide whether indeed both readings are possible in different situations ... maybe find out what triggers it. But for the system as such I would say leave it open. Of course you should really be able to see when you arrive here and when you arrive there, so there should be something in it that gives you these two possibilities, but what is most likely actually is that

it depends on the verb ... maybe for *predict* it's this, but maybe for another verb it's something else ... so that it would be small features in the lexical meanings of these verbs that determine it.

I: So these are some cases in which you would not rely just on intuition, but you would ask ...

S: Yes, there I really don't know. So I personally always think to go for the strong version, but then you just notice (amongst linguists for sure) that many many have the idea that there are these weak readings as well, so then I will just believe them and try to model it in such a way that it can go either way. So here again on the background is the partition theory, because the partition theory is only going to account for the strong case and not the weak one, because to have a partition you have to have something like "those and those people were walking and no one else". That's sort of in the answer; even if you don't explicitly say it, that's in the answer.

I: Lastly I have a couple of questions about how successful you think formal semantics has been so far. Broadly do you think that formal semantics has been successful so far in achieving its goals? Why or why not?

S: Yes, I think that a lot of things have been achieved. Yes. Then of course you are going to ask me to list the big successes ... [[*laughs*]] ... I also have my criticisms. I am maybe not even very optimistic about how things go. First of all, what I have noticed over the last couple of years ... because when I and people like me (the other formal semanticists here and around me ... at the ILLC, and maybe in Holland ... we have this sort of community of logical semanticists), when we go, we mainly go to conferences that have linguists, and ...

I: Sorry, you're saying it's only linguists?

S: No, it's a mix, but maybe the linguists are in the majority. I notice that it gets harder and harder for us — when we write abstracts that you want to get accepted — to get accepted. What it means to me is that the real logical conceptual work is apparently no longer that much appreciated, and you can also see this by the sort of commentaries that you get when a thing is not accepted. What very often plays a role is empirical motivation. If you don't

start your abstract with a concrete empirical problem that you want to solve, and then give the machinery, and then show how you've solved it, then you are in trouble. If you just present the thing as what it is (a conceptual idea or a logical system) ... when I write a paper like that I will start with the conceptual ideas, and motivate them, and then I will have at the end something that's called an illustration. And the illustration will give examples and apply the stuff to deal with the examples. So I do not do it in the other way, I do not start with the examples, and then say these are the examples I want to explain, and then come up with a story.

I: So the latter is the linguist's typical procedure you think?

S: Yes. That's the linguist's typical procedure, and he is not going to accept it if you do it in the other direction. And for me, if I do it in the other direction then I am cheating, because I did not start with these empirical issues — sometimes you do, of course, but I'm talking about papers where that's not the case — it's really a conceptual idea, and then you show that it can be fruitfully applied, but you are not going to suggest (because it's not the case) that you started with the phenomenon and that then you made up this formal machinery to get that right. That's not what you do. So there is more and more of this ... empirical methods in semantics are sort of rising, and if your story is not built upon hardcore empirical stuff, then you are not doing something that's relevant. And I think that's very very wrong.

I: Why?

S: Because there are two things that you are doing ... and that's not typical for linguistics, you have exactly the same for natural science ... in natural science you have the experimental people — they are doing experiments in big laboratories, etc. And then you have the theoretical physicists — they are just sitting behind the desk ... they have a computer, and what they do is type and scribble on paper, and they build concepts. Of course, they know about what the relevant empirical issues are, and they are trying to think of new models that are related to that, but you are not going to judge their work by [the criteria] that they should describe or have motivation for what they are doing that is given by empirical research. They are not doing the empirical research, they are doing the formal stuff, and then they hand it over to the

empirical people. We recently had this (I don't remember precisely what it was) ... something that has been thought up in physics in the Einsteinian period about ideas, and we don't know whether it is right, because we didn't have the experimental tools to check it, but now thirty years later (or whatever) there is empirical evidence that the theory that was cooked up then, that it is correct. So in semantics you need the same stuff. There is the work that is being done by logicians ... that's philosophical logic ... purely analytical work that is of course motivated a bit by data, but it is not data description, it is really conceptual work, making a new system and then hand it off to the linguist who can play with it, and the linguist can (should) then try and get all sorts of empirical testing and corpus or whatever, but that's not the job of the [semanticist]. But the linguist should see that that's also an element of linguistics (of formal semantics at least) that is important for their work. And also, if they design systems themselves, then they should try to make it up in such a way that also the logician is interested in it. So it should be cooked up in a principled enough way for a logician to be interested in the formal properties of the system that is made. So the linguist ... also to get money for stuff, you have to do experiments, and data ... and I find it all very important ... and maybe even the brain stuff, etc. ... but I don't think it gets only interesting when you can do this (also very costly) purely empirical work. I think the theoretical work is also important, and that they should feed each other. So there I am worried about the current situation. You could say that it goes too much in the empirical direction, and it doesn't see anymore what the importance of the real logical, analytical stuff is. The analytical component is equally as important as the empirical component is.

I: I was going to ask you if you think that formal semantics is a rigorously scientific discipline, but I guess you would say that it's scientific, but in a theoretical sense? You were comparing it to theoretical physics just now.

S: Yes.

I: How would you react to the fact that in terms of building models, that the stochastic models favored by computational linguists have generally had far more success than formal models?

S: This is a long time ago, and I forgot the names of the people involved, but this was about (what do you call it?) people who were building computer systems that you can talk to, and then it will understand ... not understand but give a transcription ...

I: Speech recognition?

S: Yes, speech recognition, that's what I mean. So there was a big guy in those days in speech recognition. Really in the stochastic business of course, and he made a small conference (quite a while ago), but where he got people from the more theoretical side (like myself) to go there. And his reasons were as follows: "We are doing better and better all the time. So we move from 70% success to 80%, to 90%, but even 98% is very bad, because you can do very little with it. That's still a lot of mistakes. Of course you can still use it if the user knows what makes it easier for the system to understand him. So like doctors for example, if they want to talk while they are operating and that it gets it in the right way, they can learn for example to use big words rather than small words. In specialized applications it can work quite alright, but not completely generally. And his idea was that to get it from what you can optimally achieve by stochastic means to something that can really be used, you need input from the theoretical approaches as well. So stochastics works quite alright, but it never reaches more than success in 95% maybe 98% of the cases, but the last 2%, to get it right, you cannot do it purely by stochastic means.

I: Then you need these formal methods?

S: Yes.

I: Ok. One more question. There have also been some recent criticisms made by Martin Stokhof and Michiel van Lambalgen about formal semantics — that it deals with idealizations, etc. Are you familiar with this criticism?

S: Yes, yes, I'm familiar with that.

I: And how would you react to that?

S: To be honest, I don't care very much.

I: Why not?

S: Well ... I mean, of course, it's true. You are idealizing, and maybe you are idealizing in a way that's not good. That can always happen, but the idealizations ... I know that they think about it in a way that the idealizations as they are made in linguistics or formal semantics are not of the same type as what in the other sciences happens ... I'm not so sure about that. I think if that's not a proper approach, then the whole idea that you can do anything in understanding things about language using logical tools sort of evaporates. And thats ... I mean ... even sort of something like the partitions view of questions or generalized quantifier theory ... at least things that on the one hand they have to obviously with linguistic stuff ... they have insights on it that are of a mathematical nature, and I think the mathematical stuff tells you something about the linguistic stuff ... so I don't see why we should not follow those lines anymore. I think it's just to be a matter of obvious fact that we have learned interesting things about language by these logical tools, and that's all I want — to learn more.

I: My final question to you is: what, if anything, could cause you to give up your commitment to formal semantics?

S: Death! [[laughs]]

I: Alright. Those are all the questions that I wanted to ask you. Thank you very much for your time.

S: It was a pleasure.

I: Before we end this interview, is there anything else you would like to add?

S: No, I think I've said enough. I wish you success.

I: Thank you very much.

B.2 Interview with a formal semanticist

I: To begin with I'm just going to ask you to introduce yourself. Can you describe your academic and professional background, and how you got interested in the field of formal semantics?

S: My bachelor's degree is in mathematics, and I love mathematics, but I wanted to try and find a specific area where there would be some grounding — so, not just abstract theorizing but also some empirical grounding. Then first I did a masters in artificial intelligence, and at least the way I was trained was a bit too practical (too application oriented), then I came to Amsterdam to do the master of logic here, and then I started a PhD. At that point it seemed like I would be doing more logic proper — not so much language related. But then actually during my PhD I became more and more interested specifically in applying logical methods in linguistic research like formal semantics. What was important in particular was finding this nice balance (between abstract theorizing on the one hand and empirical grounding on the other hand) in something that is so tangible and so ubiquitous as language.

I: What would you say was your primary motivation for studying formal semantics?

S: I think really the mix of these two elements and my fascination for both. Language itself and the structure and patterns that you find there; even at first sight if you don't see them, when you look closely, all of a sudden there are all these patterns and regularities that you then start wondering about how they might be captured (in one language, but also across languages). Then on the other hand the structure of the logical theories, and just the mathematical beauty of those theories and frameworks in themselves that are used in this particular field. So both sides fascinated me. Within the field typically people focus on either one of these sides, but for me it's really about both, and often even about bringing them together in new ways.

I: Would you describe yourself primarily as a logician, or as a linguist, or as something else?

S: Neither. I wouldn't see any of the two as primary. Of course, my background is more in logic, and my linguistics background is much more recent, and also I'm much more self trained in linguistics; in logic I've really had a proper training. But still I think it's more or less even.

I: Ok, so that's about your background and motivation. Now I'm going to move on to some questions about formal semantics. Just to begin with, how would you answer the question "What is formal semantics?" Suppose you had to explain it to a university student who has never heard of the discipline.

S: Formal semantics is concerned with meaning of expressions in natural languages in general. Often one focuses on one particular language, but it's really about natural languages in general. The focus in the field is not so much on the details of the meaning of each word in every language (that's a different field - I would say it's a related but a different field), but the focus in this field is more about how the meaning of complex expressions — like sentences or parts of sentences, but expressions that contain multiple words in some complex configuration — how the meaning of those complex expressions can be computed from the meaning of the simpler atoms. In that sense it's a bit like physics where you want to be able to predict the behavior of a certain object from the properties of the basic elements that it consists of — like the atoms, and the way that these atoms are put together in molecules and larger structures. That is also the main focus in formal semantics. So that's the topic. Then the method (or the conceptual and mathematical toolbox) that we use is a variety of logical systems that have mostly been developed independently in logic. Not necessarily (or not only) with this particular application in mind of studying language. We use — this would depend on the background of the student that I'm talking to — but we use (maybe the student has heard of) propositional logic and first order predicate logic. So those are the very basic logical systems that we use, but we also use extensions of these systems. For instance, ones that make it possible to analyze how meanings are composed at the sub-sentential level. As you can imagine, with predicate logic, there is only so much that you can do in that regard.

I: What would you say are the goals of formal semantics as a discipline?

S: I think that's actually a question to which I have an answer, but I should also say that people in the field might not all have the same view on. To me the goal is to understand the process of interpretation of natural language — not only interpretation, but just the concept of meaning, the nature of human language, and the many ways in which it can be used to express certain meanings, and also the constraints (how this variation across languages is ultimately constrained). We want to understand that, but at a rather abstract level. I think that many people (that holds for me personally) are not so much interested in the details of, say, the cognitive processing of language - to some extent [yes], but it's quite superficial. So I do look at cognitive experiments, and sometimes I'm even collaborating with people doing such experiments, because they can shed light on certain questions that we're concerned with, but I would say that it's still rather superficial. If it gets too detailed, then it's just too much detail to handle at the abstract conceptual level at which we tend to work.

I: Can I ask you to try and say a bit more concretely what you mean by "we try to understand . . . "?

S: The sense in which we try to understand, or the way in which we seek deeper understanding, is to construct a mathematical model of the interpretation process, and the idea is that this model is very specified in enough detail for it to make very concrete predictions. These predictions are tested against data that sometimes come from controlled experiments and sometimes just come from introspective thought experiments about language. "Does one sentence entail another?" or "Is this particular way of putting words together grammatical or felicitous?" The theories make predictions about these kinds of things, and they are tested. And then are made more and more adequate in two ways: they try to approximate the actual data as much as possible, and at the same time we try to make them more and more simple in the sense that we try to minimize the number of independent assumptions that need to be made in setting up these models — when looking at a particular phenomenon, but especially when considering several phenomena. Typically when we look at a certain phenomenon, we try to construct a theory (a model) that makes predictions about this phenomenon, but in constructing this model, in motivating the pieces that we use we typically seek motivation from analyses that have been given of other

phenomena. A very simple example: One thing that I'm studying is the interpretation of words like yes and no. At first sight the most simple and basic kind of words that there are and used all the time. Also, any language that I know has some form of expressing what yes and no express in English, but of course, it turns out to be more complicated than you would think at first sight (the interpretation of these words). Then in developing a theory of how the interpretation of these words works you seek to connect them to other types of expressions in natural language. The broader the class of words that you can connect them to (or that you can assimilate them to) the better. One broad class of expressions or constructions that exist in any natural language that I know of are so-called anaphoric expressions — expressions whose meaning is determined by some antecedent. So you can think of pronouns like "he" and "she". "Peter is going to the store. He will bring some food." Then this word "he" is interpreted as Peter, but in another sentence it might have another interpretation, so it is anaphoric in that sense. You can think of "yes" and "no" as being anaphoric in the same sense, because if the question is "Is Peter going to the store?" and the answer is "yes" then it means that Peter is going to the store. If the question is "Is Peter getting some food?" and the answer is "yes" then it means something else. But you can think of it as operating in the same way as picking up this antecedent. Pronouns like he and she have been studied a lot, and people have found out all sorts of things about how they work and don't work. If some of these ideas could be used in the domain of words like "yes" and "no" (and it seems to some extent that that is indeed the case) then a formal semanticists says that we've gained some insight. If we've not just developed an ad-hoc theory for these constructions, but if we've actually connected two classes of constructions, and if we've seen that the same mechanism governs (or at least affects) the interpretation of both (or multiple). That's when (at least for me) we have the feeling that some important insight has been gained.

I: Going back to what you were saying about predictions. Is it fair to say that formal semantics is trying to make predictions not about cognitive processes, but about how people use language?

S: There are (I would say) some predictions about cognitive processes, but not it much detail, because the theories are not stated in those terms and they

just don't make reference to very detailed aspects of the cognitive process. But there are predictions in the sense that some semantic theories for instance assume that in the interpretation procedure certain options are compared or certain preferences are weighed, and the interpreted sentence may be compared to something else that may have been said instead and the interpretation of what was actually said depends also on the interpretation of what could have been said instead. So this would be a reasoning process involving some additional ... it wouldn't be a direct interpretation procedure, but one that would involve some additional effort. So you just predict that there would be some additional effort involved in interpreting that particular type of sentence, and that can be tested, and that is tested. But from a cognitive point of view you get very simple predictions and very simple experiments (just maybe reading time experiments). This is also where people have different views ... [some] people actually try to make semantic theories or derive more complex and more detailed predictions about the cognitive process from semantic theories and also test them — that happens (more and more), but then you're already at the interface of formal semantics proper with psycholoinguistics. In formal semantics proper, the predictions that are made at the cognitive level are quite shallow. But then you asked ...

I: I said that the predictions are not so much about cognitive processes, and you said that to some extent they are, but mostly they're about how people use language?

S: Yes, they're definitely about how people use language. How they use and interpret language.

I: So, I guess we've already talked about some of the kinds of problems that researchers are trying to solve in formal semantics, but I'm wondering if there are any kind of *big problems* that everybody is interested in solving? Or some major questions that everyone in the field is interested in solving? Is there anything like that in formal semantics? In physics for example, everyone is interested in figuring out what happened after the big bang, etc. Are there any big questions like this in formal semantics?

S: I would say yes, there are big questions, but I would also say that not everyone is interested in the big questions, which I find strange actually. For

me the big questions are at the foundational level. Maybe I should give a little bit more background here to show where those questions come from. As I see it formal semantics started with the work of Frege, but the logical framework that Frege developed — classical predicate logic (or he even went a bit further but still...) — the logical framework that he developed is not rich enough for the purposes that formal semantics has now (and has had for a long time). The big question has been at the foundational level how to enrich the logical framework (the formal apparatus) to fit the purposes of the field better. There have been some very big breakthroughs in that regard, but there is also still a lot to do. One huge breakthrough (the first breakthrough) was intentional logic ... the firm establishment of especially semantics ... so the work of Kripke, that has had a huge influence. Then especially Montague's work in the early 1970s — using typed lambda calculi to provide a framework in which the meaning of sentences can really be derived compositionally from the meaning of the words all the way up to the meaning of complex sentences. That's probably seen as *the* major breakthrough in the field. Then there was dynamic semantics in the 1980s and 1990s moving from the classical notion in Frege of meaning as truth conditions to meaning as context change potential; so knowing the meaning of a sentence is knowing how it may change the context in which it is uttered These two developments ... in principle they are just independent — compositionally a la Montague and the dynamic perspective on meaning — but they've also been integrated in the 1990s and 2000s (this is quite recent, but at the same time quite established). But then there are other ways in which this framework is still very limited for the purposes of analyzing natural language. One way in which it is limited is that it is very much focused on informative content, and that results in the fact that it is suitable mostly for the analysis of declarative sentences — sentences whose main function is to provide information — but then there's all kinds of other functions that sentences may have and all kinds of other sentence types like interrogatives that are used to ask questions, or imperatives that are used to give orders, or exclamatives that are used to express some feeling or judgment about some experience, and more. That whole area — further developing the logical framework that we have in such a way that it allows us to analyze that much wider spectrum of types of sentences in language and the multiple functions that they have — that is I think very wide open, and that I would say is the big issue, the big

question at the foundational level. But in everyday practice I would say that most formal semanticists are not focusing that much on this area. The focus is much more on working within the framework that we have and looking at particular phenomena and giving an analysis of the phenomenon in question in the particular framework ... maybe adding here and there little tweaks ... a little patch here and there, but not a fundamental change. You see that also if you look at conferences; I'm quite astonished actually that what you see at conferences is mostly that type of work. Of course, that also needs to be done, because it doesn't make sense to just work on the framework itself. The image that I have at least is that in the 1970s and the 1980s there was much more emphasis on developing solid logical foundations, and now we're in a period where the focus of the field as a whole is much more on concrete linguistic analysis in the framework that we have. My own interests are at least as much at the foundational level ... I do work on both, but I see the work that I do at the foundational level as having much more significance for the field in the long term. I see the other work more as an illustration of the foundational work that I do, i.e., how that work can be put to use.

I: Then I want to move on to ask some specific questions about how modeling works in formal semantics. Logicians are often concerned with modeling in the sense that they try to create formal representations (i.e. models) of some real world phenomena in order to represent them and to reason about them. Can you say what role this notion of modeling plays, if any, in formal semantics?

S: Yes, this is also very much true in formal semantics. I would say that the theories that we build can be seen as models. We construct a model and then reason about the phenomenon via the model.

I: What exactly is the phenomenon then that you are trying to reason about?

S: The phenomenon would be the interpretation or certain entailment patterns or felicity conditions on certain classes of expressions. Those would be the phenomena. Then the model would be a theory of the interpretation of these expressions, etc.

I: To probe a little more on what the phenomenon is ... you say that the phenomenon would be something like entailment patterns, but that's clearly

not something in the world that you can point to and say "here's the thing that we're trying to model." How would you describe it? In some sense the physicist can say "my models are about the physical world," and the economist can say "my models are about transactions in the marketplace", but it just seems a bit more difficult to say exactly what it is that these models are about.

S: So it's about the interpretation (or certain aspects of the interpretation) of expressions by people using the language that these expressions are part of, and you can't really point at them, but you can measure them in the sense that you can ask those people whether they draw a particular conclusion ... if they hear sentence X whether they draw this conclusion or that conclusion ... or whether they find sentence X felicitous in a particular context.

I: So maybe one way of putting it is that you're trying to model peoples' intuitions about language?

S: Yes.

I: Also that you're trying to model the way in which they're actually making use of language?

S: Yes. Not only the meta-intuitions, but also the way they really ... in an experiment you could ask "do you think that this entails that", but you could also just test it without asking whether they think it does, but test it in a more direct way.

I: This leads to my next question. How do you check that the model you're constructing is actually a good model for the thing that you're trying to model?

S: So that happens either by means of these controlled experiments, which can be of various types, but the simplest experiments are truth-value judgements. So, you give someone a certain context, you give them a certain sentence, and you ask whether the sentence is true in the given context. Notice that this is also very much an experiment that fits within this truth-conditional semantics that is about declarative statements that are true or false, but not about interrogatives or imperatives or exclamatives (it just doesn't make sense to do truth value judgments about those). So that's one of the main experimental paradigms. And another is felicity judgments; so you give a certain context

(sometimes you don't even give a context), and then you ask if this sentence is felicitous (either in the given context or per se). Then there are other types of experiments that you could do. You could make people choose between two options ... there is a question and then they get two possible answers ... they choose ... or there is reading time experiments to test predictions that are more about processing speed ... and there are more complicated paradigms, but those are the basic paradigms that are used to test the basic predictions of semantic theories. That said, these experiments are being done more and more, but many predictions are also just tested introspectively, or maybe with a small number of consultants just informally; so, not in a formal experimental setting, but just with four or five people that you ask whether they agree with your own intuitions (they are often called judgments, haha, but yeah, intuitions). That's how theories evolve. Also sometimes you see in a published paper — "This is a prediction and it accords with our own intuitions; we've consulted like five people and they've all agreed (or maybe they didn't all agree - because it's often the case that people have diverging intuitions)" — and then the experimental work is left for another time. So the theories that are constructed usually make lots of predictions, and sometimes those are tested to some extent, but to test them all would be the project of a lifetime, so it's a field in which the theory moves ahead of the experiments. Whereas in other fields, and I've seen to some extent that experiments move ahead of the theory.

I: What are the standards for success or failure in modeling? What distinguishes a successful model from a failed one? So, I guess you've already said that in one sense you could consider a model to be somehow a failure if it doesn't agree with either your intuitions or the experimental data. But are these the only criteria [for success / failure]?

S: No. This is a subtle thing if you think about it deeply, but roughly there is what people call empirical adequacy (which is this — making the right predictions), and there's explanatory adequacy (that's the subtle part actually). I think what's generally understood in the field as being explanatory adequacy is what I mentioned in the beginning about connecting things. So, a theory is more explanatorily adequate the more it connects one phenomenon to another, and it doesn't just explain one isolated phenomenon.

I: So these are the two main criteria for a successful model?

S: Yes.

I: Would you say that formal semantics is a rigorously scientific discipline?

S: Yes. I would say yes. I think that I would say it's as scientific as it gets. I do feel that sometimes people get the impression, from outside or from people peeking over the shed from the neighbors', that they might get the impression that it's not as scientific as other disciplines may seem. That maybe because, as I said, the theory is often ahead of the experimental corroboration of all the predictions that those theories make. As long as those experiments are not done, there is often disagreement about the data; I mean there aren't any data, so that's difficult. Everybody has intuitions about language, and in particular the researchers themselves, so as long as there hasn't been a controlled experiment on some construction, it's everyone against all the others, and there's lots of conflicting intuitions. It can happen that for some time there are theoretical debates that have no empirical grounding, and then at some point an experiment is done to figure out at least who makes the right predictions for those cases. Even then ... with an experiment you only test just a little area ... also there are methodological arguments that could be made against a particular experiment, the way it's been carried out. There is one big debate at the moment about something that people have been arguing back and forth at a theoretical level, and they have had lots of conflicting intuitions, and now people start to do experiments, and you would think that it would all clear up, but now the debate has become to a large extent about the methodology of these experiments. So it's not that somebody did an experiment and said "These are the data, so this theory is correct." No, it's one experiment, and then other people have done another experiment with different results and a slightly different methodology. So the field is still finding out how to better combine empirical and theoretical methods. That's fair to say, but I would say it's as scientific as it gets.

I: Lastly I have a few questions for you about how successful you think formal semantics has been as a discipline. Broadly do you think that, as a discipline, formal semantics so far has been successful in achieving it's goals?

S: Yes, I think in a relatively short period of time and also relatively short manpower (compared to fields like physics it's a small field), I think we've had

some important breakthroughs at the foundational level. We've made a good start at combining theoretical and empirical methods. There are fascinating insights about how things work across languages. So I think a lot of impressive progress has been made, but at the same time I also feel that the field is very much in its infancy and there is a lot to be done.

I: What do you think about the fact that when it comes to applications, the stochastic approaches favored by computational linguists have enjoyed a lot more success generally than formal models?

S: Yes, so I think that's a natural consequence of the fact that formal semantics as such is a longer term project. We're not ready to provide, we can't provide yet, what stochastic methods can provide for technical applications. But I'm confident that at some point formal semantics will be able to provide ... at least it will be able to provide certain insights that stochastic methods cannot provide. Actually at this point there is already in computational linguistics a trend towards at least combining insights from both sides. For instance there is compositional distributed semantics. There are a couple of big projects in Europe ... they happen mostly in computational linguistics labs, so it's computational linguists looking to import insights from formal semantics at this point. I haven't really seen much traffic going the other way — like formal semanticists trying to import ideas from computational linguistics, but I think I see it as a good development that both are making progress. And at this point it's just more practical to use stochastic methods for actual applications, but I foresee that in future as formal semantics makes more progress, the insights that are gained there will also be beneficial for the applications.

I: What do you think is the fundamental difference between these two approaches?

S: Hmm ... [[*long pause*]] ... What's the fundamental difference?

I: Or we can remove the word "fundamental" ...

S: Ok .. What's the difference? So, in stochastic approaches it's not that much about the underlying structure, so it's about learning a grammar or about learning a translation from one language to another from data (so texts) ... you just look at the texts at the surface ... [[*pause*]] ... that's not true

... there are also methods in computational linguistics that are exactly about finding out the structure, so I'm not sure that that's it ... [[*pause*]] ... It's hard to say; I don't have a very good answer.

I: That's fine. One more thing I wanted to ask. There has recently been some criticisms about formal semantics made by Martin Stokhof and Michiel van Lambalgen. I'm wondering if you are familiar with this criticism, and if so how you would react to it?

S: Yes, but I don't have something very sensible to say about that I think. Somehow it just doesn't seem to affect the way that I see myself working.

I: Ok. Why do you think that is?

S: It just doesn't ... I think they make good and helpful distinctions at a meta-level looking at what's happening in the field, but I don't necessarily see it as a criticism, but just as a reflection.

[[Interruption]]

I: What, if anything, could cause you to give up your commitment to formal semantics?

S: Nothing as far as I can see. I don't see something as potentially invalidating the whole enterprise. Because of this modeling nature of it, I see it as a practice that can give us more insight and that gives us more insight ... maybe in an extensional way ... so, the inner workings of language are not necessarily unraveled, but that's not the ambition for doing that. So, if you are not too ambitious then I think that the enterprise is one that is coherent, and I can't see how it could be objected to on principled grounds.

I: Ok. Those are actually all the questions that I wanted to ask you. Is there anything you would like to say before we end the interview?

S: No.

I: Thank you.

B.3 Interview with a formal semanticist

I: Let's start with a brief introduction of yourself. Can you please describe your academic and professional background, and how you got into formal semantics?

S: I started off as a physics student. Then after (what is now called) my bachelor, I did both a masters in mathematics and in philosophy, and actually in the beginning of the master's I discovered that logic was what I wanted to do. Both in the mathematics track and in the philosophy track I did as much logic as I could. But I've always been attracted to the logical analysis of natural language — that started right in the beginning already with an interest in generic statements (I gave my valedictory speech about that). Logical analysis of natural language — that has been my main concern since, but I have a broad interest, so every now and then I did things outside that field. I have done some work in the foundations of mathematics, on interpretability logic — I developed models for that. Some of my work has been more technical, that I think semanticists would never do ... in conditional logic I've done very technical things also ... But roughly it's philosophical logic focussed on the logical analysis of natural language — that's my branch.

I: What would you say was your primary motivation for studying formal semantics?

S: I sort of *got there*. My primary motivation was logic. But then it turned out that to get a hold of logic you have to do formal semantics! I discovered that ... I didn't like developing logical systems just up in the air ... I did do some intentional logic (technical) without any care for applications or whatever, so I can see that there is just a mathematical study of modal logics, etc. ... but I was always attracted to the use of logic in daily practice and to how people reason and stuff like that. In the beginning I also did a lot of argumentation theory, but that at some point I found boring, so I gave up on that. My real love is really in developing new ways of modeling. So at some point things became very philosophical, because we (so it's not just me, but the group here in Amsterdam) thought at some point that looking for truth conditions is not the right way to go ... that you had to do different things if you want to get hold of meaning, and I certainly think that in natural language and in reasoning

a different notion of validity plays a role than, say, the classical standard notion in terms of truth. That led to some more general philosophical (project is a big word, but) papers.

I: Would you describe yourself primarily as a logician or as a linguist or as something else?

S: As a logician ... Last Thursday I said that I think of myself as doing things that are interesting for both linguists, philosophers and people working in artificial intelligence, and that the only thing I have to do is change ... for philosophers you have to use "subject" and for A.I. people "agent" and for linguists you talk about "speaker and hearer" ... so with these changes I can adapt my talks to the audience!

I: OK. So that was about your background and motivation. Now I have a few questions about the field of formal semantics. To begin with, how would you answer the question: what is formal semantics? Suppose you had to describe it to a university student who has never heard of the field?

S: Well I would start by answering the question "what is semantics?" In semantics you try to get hold of how expressions of languages get their meaning ... that's one thing ... well that's enough as a definition for semantics ... of course the answer raises more questions than it answers, but that's the answer! And then *formal* means that you are not ... that you are really ... when you do formal semantics you use mathematical models to do that, and not just ... I see the *formal* more as a heuristic or as a method — something you do to get hold of the meaning. It's part of the methodology and it's not the goal in itself.

I: What are the goals of formal semantics then?

S: Eventually to understand how meaning arises — how expressions of language get their meaning. It's a kind of a miracle in a sense how meaning is built up, for instance, of larger parts by combining in some way the meanings of smaller parts and how that goes.

I: Could you maybe be a little more concrete about what you mean by "to understand?" Because many disciplines share this same goal.

S: To understand the meaning?

I: Yes.

S: Well, if you really understand it, of course, you could build a machine that understands natural language for instance. Because you could then implement in the machine what it is to know the meaning of the expressions of the language. So that would be the ultimate test I think. One of the byproducts ... once you know the meaning ... (actually it's not so much that you want to know that meaning ... you want to know how meaning is built ... it's more a how question than a what question) ... but once you know the meaning you also know the logic. Actually that's what I'm after; I'm not so much interested in meaning. And certainly formal semantics is not the only thing that matters if you're after meaning, because then all kinds of pragmatic things also come in, so you have to do pragmatics too. But for logical purposes semantics should suffice.

I: I was just struck by the fact that previously when you mentioned the goals of formal semantics as perhaps building a machine ...

S: No, that's not the goal. The goal is just giving a description of how meaning arises, but you can test whether you've reached the goal if you can build a machine. I would never build the machine myself, and I would probably think it's a waste of time ... but how do you know ...

I: What are some of the kinds of problems that researchers try to solve in formal semantics? Or what are some of the research questions that they try to answers?

S: If you look at semantics journals and you look at how things grow, then you see that at some point people are interested in, say, anaphoric chaining. Now, for instance, people are very interested in epistemic modalities or in adjectives or in ... In a way you can look at that as people trying to describe certain parts of natural language and how meaning works there, so they go from one topic to the other. But that's not really what's happening you see, because the reason that a certain topic gets interesting is because *the framework* (that is at some point the main framework or the theoretical framework in which people are expressing their theories) doesn't work anymore. For instance, dynamic

semantics was developed because in a static framework you cannot handle anaphors in the right way. That was one of the mistakes I made certainly ... I couldn't care less when this started ... my promoter Hans Kamp who started this discussion, he is a logician too ... and I would think "Hans, what *are* you doing? Who cares about anaphors? That's for accountants. It's not for us." The thing is ... there you see that some things don't work, some ideas that you have about the building blocks and the mechanisms that play a role, and then a certain topic becomes important. So the rise of dynamic semantics started with anaphors, and then epistemic modalities became important there (to show that you really needed to go dynamic) ... that's just an example (that's just because that happened *here*) ...

I: So if I understand correctly, you're saying that there are two kinds of problems that researchers try to deal with: (1) one is trying to give an analysis of some puzzling linguistic phenomena, and (2) the other is to build up the framework that we use for this analysis.

S: Yes. So those things go hand in hand. Yes, but the big issues are "what is the right framework?" and then the normal science is to develop theories within one of these framework. For a long time Montague grammar was the framework, and lots of good work has been done there. Nowadays it's not so much that there is one framework; people are willing to work in several frameworks and sort things out and see how things go.

I: Are there any *big problems* that everybody in the field is interested in solving? Like central questions? For example, maybe in physics one might argue that one of the big questions is to figure out what happened immediately after the big bang, and almost everyone is interested in this question. Would there be anything comparable in formal semantics?

S: I think there is certainly a problem (but that's not so much among the linguistically motivated semanticists, but certainly among the philosophically motivated semanticists like me and most people here) — it's of course propositional attitudes, but nobody dares to touch it. Everybody knows that the possible worlds framework doesn't work for that kind of topic, and people have been trying different things ... but I think that's ... actually even if a student wants to work on that even for a dissertation, I forbid it, because it's way too

difficult! In a sense you cannot say "Oh, that's a very nice problem; let's work on this, and then in four years we'll know how to do it." It's too risky, unless you come in and say "I know how to do it; here's the idea" then you can work on that topic. But I'm certain that everybody wants to solve propositional attitudes, because that's really where all frameworks go wrong.

I: So when would this problem be considered to be solved? Or in general when is a problem considered to be solved?

S: Well here ... well, if it get's the facts right ... but here it's also a question "what are the facts?" But the big problem there is that in all approaches if an agent believes something or wishes something (maybe for wish it's not so bad, but for belief it's bad), a certain proposition, then he believes everything that expresses that proposition. That means that if you say "I believe that A" then you believe that B for every B logically equivalent with A. Then you can weaken the logic there (what logically equivalent means), but however you weaken it, it's still the case that you believe everything that's logically equivalent, and that seems to be wrong. On the other hand, you cannot just make ... you could solve that problem by making a believer relation between an agent and something syntactic a sentence, but then you cannot explain how people can quantify in belief contexts and how I can translate when you say in French "Je crois que ..." ... and I can report that to my friends by translating what you say into English — that works fine and there's nothing wrong if I say "Nikhil believes ... '' in English. So there are puzzles there that make a mess of the topic. Also a problem that is now just fashionable ... I like it, I love that problem but I think it's not so deep as the propositional attitudes — it's the personal taste problem. You can say "that's beautiful" and I can say "that's beautiful" and there's a kind of a subjectivity there which is very difficult to explain. If you work in a truth conditional framework it looks like we have what people call a "faultless disagreement" ... it's so difficult to explain what's going on, because you are almost forced to say that we give a different meaning to the word "beautiful" and you don't want to say that, but the way things are built, you should say that we could just as well use two different words. Because "beautiful according to me" and "beautiful according to you" are like two different words, but that's not what you want. That's also certainly something that concerns the framework ... to solve that problem

will definitely lead to changes in the framework. Now it's a big fight between contextualists and relativists ... those are the two schools there ...

I: To go back to the question ... under what conditions would people agree that the problem has been solved? Is there any such criterion?

S: Well, theories do predictions. Specifically, semantic theories predict which arguments are valid and the conditions under which people will agree with a certain statement and things like that. So these predictions should be right. That's difficult to test by the way, because people also make mistakes.

I: How do you mean?

S: Things that they are prepared to say "oh no, I was wrong." If I make a theory that from *if A then B* it follows that *if B then A*, and I go test this, many people make that mistake. They would say "yes, that's fine." But it's wrong! So at first they will agree, and then you sort of have to convince them that they made a mistake. So it's not that you can easily test these theories. The standard, pretty naive idea is that "we set up a theory; this theory does predictions; it tells you for instance which arguments are valid; then you go and ask people 'do you think these arguments are valid?' and then people have an intuition; their intuition then is the ultimate test." Well, that I find pretty naive, because I don't know what intuitions are here.

I: This actually leads me to some questions specifically about how modeling works in formal semantics. What I mean by that is: logicians are often concerned with modeling, in the sense that they try to create formal representations or models of real world phenomena in order to represent them and to reason about them. Can you describe what role this notion of modeling plays in formal semantics?

S: The way I look at modeling it's like drawing a cartoon. You leave out lots and lots of things and you exaggerate a few other things and you end up with something that strikingly looks like what you want to describe ... that's how you draw a cartoon ... that's how it works. In the mathematical models you leave out lots of things, you make lots of idealizations, but what you want to describe is really something that you could just as well try to describe in ordinary language ... depending on the framework, but let's say

in the standard framework ... what you then try to describe is the conditions under which sentences are true, in general. It depends of course what you are modeling ... that can be quite complex, because the context of utterance may affect the meaning ... of course, you can refer to different situations that will affect whether it's true or false, and so forth.

I: Despite this cartoon like nature of it, we still want to say that this model *is* actually a representation of the phenomenon, and that we can reason about it, and thereby infer things about the phenomenon, etc.

S: Yes, definitely.

I: Typically what is the phenomenon that we're trying to model in formal semantics?

S: Well here's something that's typical. If you say that you want a theory of meaning, then that's way too wide if you look at the practice. Because in practice in formal semantics you don't do what is called lexical meaning (meaning of ordinary words like chair and table); actually it turns out that what we do is we are interested in more structural things like connectives, like tense operators, like modal operators ... like things that we call logical constants ... it's difficult to explain from scratch, without any theory in the background, what a logical constant is ... but I certainly think that formal semantics only works and is only useful for getting hold of the meaning of structural things (what I call structural operations in language). It has no use for lexical semantics.

I: So the thing that you're trying to model fundamentally is meaning and structural aspect of meaning?

S: Yes, structural aspects of meaning. That's right.

I: The predictions that are made by these models, they are also about meaning?

S: Well ... yes ... well, the meaning of some of what I call logical constants is explicitly described in the theory, and the predictions then are about arguments in which these logical constants are used — whether they are valid or not valid, or which sentences will be equivalent, or things like that.

I: So could we say that the predictions are about the use of language and about linguistic behavior?

S: Well, for that you need much more than just this, but it's a natural ingredient for that, it's a basic ingredient for that. But additionally you need lexical semantics (I suppose) and you need pragmatics, and that's really also very important. To understand a dialogue ... truth is not the only thing that matters.

I: Once you have some sort of model of the phenomenon (together with some predictions, etc.), how do you check that the model you have is actually a good model for the thing that you're trying to model?

S: I already said that it's usually checked very naively by looking whether it matches the intuitions of the native speakers. This is actually the easy case (when people at least agree), but there is a more difficult case, because often — but now I'm talking more about philosophical logic than maybe about semantics — often it starts with a problem where intuitions leave us groping. Like a paradox — like the Sorites paradox for instance is a paradox that arises with vague adjectives. As soon as you want to do semantics for adjectives, you have to solve that problem somehow, or think that it shouldn't be solved (that it's fine that it's there), but you have to take a position there. So in that case you can think of the following: that people solve the paradox, but they do this by giving a sort of meaning to adjectives (1) that doesn't fit the meaning they have in ordinary discourse (which you can show) and (2) which makes that totally useless for ordinary discourse. In this particular case, I myself for instance think that gradable adjectives give rise to a paradox, but who cares? It only shows that gradable adjectives only work well in certain contexts, and there are contexts where you should not use them, because they are not the right tools for those contexts and we develop other means and measures to work with [those contexts]. When you have to compare thousands and thousands of things on color, what can you do with your primary color words? Things are bound to go wrong, but that's not bad. So there I would never solve the paradox in that way — by giving a meaning to these color words in such a way that the paradox would not arise. Then you would force the language user to be so precise that they cannot use the words in an ordinary way. Because in mosts contexts there are only a few colors that matter.

I: OK. So I'm still not clear on what are the standards for success and failure on modeling ...

S: If our natural language were a finished thing, then I would say the standard is ... if you think of natural language as a system of conventions ... then you're finished if you've described the conventions that are there correctly. The normativity of a semantic system of logic comes from the conventionality (the fact that language is a system of conventions), but you describe *the conventions*, so in that sense semantics would be descriptive. It is the description of a prescriptive thing. The thing is that that's not a good answer, because it presupposes that language is a nice system of conventions, but it's actually not — it's a lot of systems, and some are only half built and others ... it's a mess. In that sense logicians can also do proposals to sort of help building and expanding language. There's also stuff in natural language that shouldn't be there because it's totally useless. I think of language as something that has been developed by human beings during the course of evolution by trial and error, and so there's parts in it that work well and parts that do not work well, and parts that are very easy to describe and parts that are difficult to describe, and so forth. So it's hard to say when is a theory successful.

I: What role, if any, does empirical testing play in this process of checking whether the theory is successful?

S: Things are changing there, because people now do lots of more extended tests — on the internet, by corpuses, and so forth (I also use the internet, but more for finding problems than for testing). So, I don't know ... empirical ... I think ultimately ... OK, here's the ultimate empirical criterion whether a theory is good (I won't say right): if people are prepared to use their language according to the conventions that the theory describes, and to have themselves corrected by the theory. Yes, I think that's the best I can say. If the theory does predictions about validity ... people reason, but they make mistakes, and then if they are prepared to say "sorry, I was wrong" because the theory says they're wrong and the theory explains why they're wrong ... that's a good test for the theory. So, it's not just testing peoples' intuitions — no, it's the explanatory power of the theory which matters, and the theory should help convince people. So it's not just the predictive power but the explanatory power that matters.

I: Do you think that formal semantics is a rigorously scientific discipline then?

S: Sure.

I: Even though it depends not on testing predictions, but on this explanatory dimension ...

S: Oh, you mean if it's empirical ... well what do you think [about engineering]? If people are prepared to buy something that you as an engineer develop, would you say in that case that it's rigorously scientific? In some cases there is no ... sorry, forget this. Let's put it this way: it's certainly rigorous, and whether it's scientific is something else ... because it's not purely descriptive. But I think you have the same thing in more scientific contexts — as I said with engineering where you develop instruments and stuff. Then of course there is also the case that what we are describing is already a perfect instrument, so in that sense it *is* a rigorous science.

I: Finally I have a couple of questions about formal semantics as a discipline and how successful it's been so far. Broadly do you think that formal semantics as a discipline has been successful thus far in achieving its goals?

S: Yes, I think so. Although I think the ambitions in the beginning and the optimism in the beginning were too high. First the ambitions as to what formal semantics might accomplish, and second that it would be easy to do so. Actually it turns out that it has given us a way to discover how incredibly complicated natural language is, and how rich and so forth. But I think certainly it has given us more insight into how language works, definitely. I wouldn't know what the *ultimate goal* of people were ... we are actually very far from what I said [earlier] that we can build a machine that understands natural language, but I am sure that what formal semantics has to say will play an important role there.

I: So in this context, what do you think about the fact that when it comes to these kinds of applications, the stochastic approaches favored by the computational linguists have enjoyed far more success than these formal approaches.

S: Yes, I know. They have far more success, but the machine doesn't understand anything, so I don't believe in that approach at all ... as it comes to the idea of "let's build a machine that understands natural language."

I: Why ...

S: Because they don't understand it! They might go till, say, 80% and then get stuck ... But it's not so much the reason that it's statistics vs. formal ... the reason that these machines do not understand anything is because they don't have an intentional relation with the world, and that's absolutely necessary. You cannot know what the meaning is of something without having some intentional relation to what you're talking about. It's just symbols ... it's just a Chinese room ...

I: But wouldn't this be true of the models in formal semantics also ... ?

S: Sure. Definitely. If you make a machine that doesn't use stochastic models, but that takes a sentence of natural language and translates it into some logical formalism and then operates on that formalism, it's also wrong. So, that's what I said. To build a machine that understands natural language, the machine should have an intentional relation to the world. But then I think having that, the other things that would [need to] be built into that machine is that it would know the meaning of logical constants and stuff like that, and there I think you couldn't do without the theories that formal semantics has given. I'm not saying that the machine would speak or use or know the logical formalisms that are used in those theories ... that's something from outside — that's a very abstract specification of what the machine should be able to do. How the machine does it in the end I don't know, but that's a different matter ... but I'm certain that it will not be statistics.

I: There have also been some criticisms of formal semantics recently made by Martin Stokhof and Michiel van Lambalgen — that it deals with idealizations, etc. Are you familiar with this criticism?

S: I'm certainly familiar.

I: And how would you react to it?

S: I disagree.

I: Could you elaborate?

S: It's such a scholastic kind of criticism. So yes, I think that formal semantics uses idealizations the same way that mechanics uses idealizations and that

doesn't matter. Think of mechanics where you have point masses and you have this and that ... that's fine. For instance, take the notion of possible worlds — I think that methodologically it's a perfectly fine notion. So there's nothing wrong with it from a methodological point of view, and that's probably what Martin and Michiel would say, but on the other hand I'm certain that in the end the theories will not have possible worlds in them anymore. That's because it's picking the wrong thing ... whereas for point masses ... maybe for some cases you can still use possible worlds just like for some cases you can still use point masses in predicting how billiard balls will roll over and so forth. No, I disagree [with Martin and Michiel] ... I certainly think that the theories as we have them do have some cognitive value (maybe not as much as they will have in the end, but they do have some cognitive value), and I also think that they give a very abstract specification of what's going on ...

I: I think part of the problem they're identifying is that since these models and theories in semantics are about these idealizations, there is no straightforward way to verify or falsify any of the claims made by them, because the data that we have are about actual linguistic use and behavior (not about the idealizations).

S: Yes, sure.

I: Is it not a problem then that there is no straightforward way of verifying or falsifying the claims of semantic models and theories?

S: No ... then they behave as if they come from Mars and look at natural language as Martians would do, but they don't come from Mars and they know natural language themselves. I really disagree. Actually I've had many discussions with them [afterthought: I have had many discussions with Martin only, not with Michiel], so we agree to disagree there. But on the other hand if I see how, for instance, Michiel does semantics, I don't see any difference basically with the way I do it.

I: But I think their criticism was supposed to apply even to their own past work.

S: Yes, but he still does ... I mean ... OK.

I: OK. My final question to you is: what, if anything, could potentially cause you to give up your commitment to formal semantics?

S: There should be something of course ... What are the presuppositions? ... If one of the presuppositions that I am building my theories on ... no, it's not so much the presuppositions of the theories, but the presuppositions of the trade ... I would be prepared to give up a lot ... if people tell me truth conditions do not matter ... I think at the basis of any formal semantic theory ... Let me put it this way: For me semantics starts with the notion of validity, and then it works as follows — if you explain validity in terms of truth, then you're bound to give in your semantics truth conditions. If your validity works in terms of intentional attitudes, then you have to describe how peoples' attitudes change. Semantics starts with the notion of validity ... there are two things on which my trade is built: (1) a notion of validity and (2) the conviction that I can use mathematics to describe what's going on, to describe this notion, or to get a hold of this notion. So I would be prepared to give it up if you can convince me that there is no notion of validity (of logical validity) at all — that's one thing. Then the other thing is that (maybe that's the line of criticism that Martin and Michiel think would follow, but I'm certain that Michiel doesn't think so) if you can convince me that you cannot get hold of this notion of validity by using formal or mathematical means. I think these are the two presuppositions [of my trade].

I: But what about maybe a third presupposition — that this formal notion of validity is in fact operative in natural language?

S: Yes, sure, of course ... I forgot ... that's the way it *should* be put. It should be operative in natural language, definitely (maybe given that I'm retired, I would play with it anyway [[*laughs*]], but no). For the trade it's very important that it should be operative, it should be there, really there in natural language. That's actually one of the reasons why I gave up the idea that it should be truth conditions — that validity should be defined in terms of truth — because I think that's not the notion that is operative in natural language.

I: OK. These are all the questions that I wanted to ask you. Is there anything else you would like to say before we end this interview?

S: No. They were good questions I think. I'm curious to see the end result, and I'm curious to see what the other people say!

I: OK. Thank you!

[[After the interview was over, the subject remarked that they were not too happy with the statements they made about Michiel and Martin's criticism of formal semantics.]]

B.4 Interview with a formal semanticist

I: Let's begin with a brief introduction of yourself. Can you please describe your academic and professional background, and how you got into formal semantics?

S: When I started my studies I started first with communication science (something like this — it was an engineer kind of topic), but then right in the first semester I became interested on the one hand in the mathematics and on the other hand in formal linguistics, so I started these two fields. These were my undergraduate studies (mathematics and formal linguistics) and then in my graduate studies I focused on logic, in particular on the application of logic to linguistics, to semantics. That's basically how I came into this field. Then in my PhD I really specialized in the application of logical tools to certain phenomena in semantics and pragmatics. On the one hand in pragmatics I tried to formalize conversational implicatures, and then in the second part I tried to use non-monotonic reasoning to describe the semantics of conditional sentences. So that's basically my background, and now I'm actually moving on even further — so now I'm shifting into philosophy and to metaphysics more and more. Is this answering your questions?

I: Yes!

S: Good!

I: What would you say was your primary motivation for doing research in formal semantics?

S: I think there were two. One thing was that I was very unhappy with the limited depth of the formalization used in engineering, so I wanted to go deeper and I wanted to ask further "why is this?", "what is the motivation?", "why these formalizations?" ... I really wanted to understand why certain formalisms were used. That was one thing (this was particular for mathematics), and the other thing was linguistics ... I think it started just because it was a lot of fun to do. Formalizing linguistic phenomena was just so much fun — it was like gaming, playing games. That's what got me interested in formal linguistics.

I: Would you describe yourself primarily as a logician or as a linguist or as something else?

S: Logician I think. Not a linguist.

I: Why?

S: Why is that? I think the difference is that linguists are primarily interested in the data (in the data themselves), and I am not. I am primarily interested in the theory — in formalizing the data.

I: Alright. So that's about your background and your motivation. Then I have some questions about the field of formal semantics in general. To begin with, how would you answer the question "what is formal semantics?" Suppose you had to describe it to a university student who has never heard of the discipline?

S: In formal semantics you try to use certain tools to generalize the observations you make in linguistics. In linguistics you have a lot of data (all kinds of data — different languages, different speakers ...), and you try to find a system that underlies these data. In formal semantics you use formal tools to describe these systems, to make them concrete, and to be able to study the underlying system.

I: And what are the kinds of tools that are used for this?

S: All kinds of formal systems that there are ... basically everything ... every framework you have that you can use to build systems and to model stuff, you could basically use [in formal semantics]. There are some that we know that work for certain kinds of phenomena — we know that logic works, we know that game theory works, decision theory is useful — but there might be other tools that we can use for new phenomena that we want to describe. We are just not aware of that connection right now, but basically everything — every formal tool, every formal framework you can model things with — might be useful for formal semantics and formal linguistics.

I: What are the goals of formal semantics as a discipline?

S: I think basically it's a tool. You want to understand the language and you want to understand the system (you think there is a system underlying it), and

this tool helps you to get a clear grasp of the underlying system, in order to really study it, to really understand what the system is, to work with it. It's basically a tool.

I: Can you maybe try to be a little more concrete about what you mean by "to understand?" When would you think that you have achieved this understanding?

S: One important thing is that you are able to make predictions, and that you can test these predictions and they turn out to be correct. So you think that if you can generate, for instance, interpretations for sentences that turn out to be indeed interpretations that speakers get ... or you can predict sentences that speakers accept as well formed sentences of the langugae and things like that ... you can predict limits — for instance, you can predict universals for all languages, this kind of stuff ... then your tools are working. This is one important thing, and then there should also be ... kind of ... that's hard to get ... so you want to understand something, but what exactly does it mean to understand something? Somehow my impression is that a formalism can help, but that might be because of the way I think and the way I understand things that formalizing things helps me to understand them and to get a better grasp of them.

I: So in some sense the goal of the discipline is to understand and to formalize and to make predictions about they way people use natural language?

S: Yes.

I: Can you give some examples of the kinds of problems that researchers try to solve in formal semantics? Or the kinds of research questions that they try to answer?

S: So maybe then I should use something that I worked on ...

I: Or it can just be general ...

S: One of the things I worked on was conversational implicatures. There was this idea that sometimes we generate inferences based on (additionally to the meaning of the expression) the context and certain agreement between the speakers. When I started on this there was a very vague idea of what kind of

inferences this could be — what kinds of / types of inferences you could get. You wanted to get a framework for "these are all the inferences and these are the situations when you get them" and I tried to formalize that, to make exactly these kinds of limit conditions / predictions like "this is all there is, and this is when you get this prediction." What else? It's really to systemize the kinds of observations you make ... to systemize. I also use formalisms sometimes when there's a debate in the field, and some people argue that "because of this and this, this is impossible", and then I try to build a formalism that shows that it is possible. Maybe it's not correct, and maybe this is not how it actually works, but you can't say that it's not possible to have a system with these and these conditions, and that makes these and these predictions. So that's also something I do with formal tools.

I: Then I'm interested to know if there are any big or central problems in the field of formal semantics that everybody might be interested in. What I have in mind is for example in a field like physics, one might argue that almost everybody is interested in figuring out what happened immediately after the big bang—

S: Yes.

I: Are there any big or central questions like this in the field of formal semantics?

S: For me (but maybe that's just me) universal constraints are very interesting — limits to what is possible in language at all, what is possible in understanding. This is interesting. And then there is this link between cognition and language that I think is an important issue that everybody somehow is interested in — what is the link between the way we think and language.

I: Can you elaborate a little bit on this?

S: Not really ... not right now ... I have to think about it more.

I: OK. Then I have some questions specifically about how modeling works in formal semantics. What I mean by that is: logicians are often concerned with modeling, in the sense that they try to create some formal representations or models of some real world phenomena in order to represent them and to reason about them. Can you describe what role if any this notion of modeling plays in formal semantics?

S: Basically it's the same I think — what you just described. There are certain phenomena in language (I don't know, say, implicatures), and you try to model them — you try to build a model in order to make predictions, in order to infer. So basically it's the same idea.

I: And what exactly is the phenomena that you're trying to model?

S: Well, that depends. For me it's a lot about interpretation ... it's about meaning — how we understand things and how we code our intentions in language. But there might be something else depending on the field you work on. But for me it's about interpretation, encoding intentions, etc.

I: So it's about the way people interpret language—

S: Yes, and the way people use language to communicate whatever they want to communicate.

I: More broadly could you say that you're trying to model linguistic usage and linguistic behavior?

S: Yes, but that's very broad. Then it's like everything ... linguistic behavior ... Yes, we use language so that's what it's about ... But linguistic usage is also about the words we use, the sentences we use and stuff like this (and this is not primarily what I am interested in), or where these words come from. For me it's mostly how we now understand these words, and also (for me) the link between understanding and reasoning, thinking.

I: So how do we check whether the models we're building are actually good models for the things that we're trying to model?

S: The question is how to check whether the models really cover the kinds of observations in the world that they are intended to cover, right?

I: I guess that would be ... my question is more general I think ... once you build a model, how do you reflect on whether it's actually a good model or a bad model? So this might be one of the criteria that you use.

S: Basically these models are intended to describe something in reality, and then I have to check whether they really describe the thing in reality that they intend to describe. In my case, it's mostly about whether speakers of the

language do get the kind of interpretation I predict for these sentences. Or even more generally (depending on the theory) about language independence ... so then I should show that they really apply to all languages somehow ... that in all languages you can get these kinds of constructions with these kinds of interpretations.

I: So, basically you're saying that your model or theory will predict a certain kind of interpretation for a certain class of sentences, and you will check whether actually speakers are interpreting the sentences in this way or not?

S: Yes.

I: And if it turns out that they are interpreting the sentences in the way that your theory or model predicts, this is a sign that—

S: An interesting problem that actually came up in a class I was teaching — what happens if one of twenty speakers disagrees? How far is that a problem? When do we really have a problem with a theory? So there is this probabilistic [element]. Most of the times ... right now I think most of the linguistic theories are not intended to make probabilistic predictions saying like "I cover ... '' ... how to explain this? For instance, you also talked to Michiel, and he tries to cover certain inference patterns, and then you have observations like "60% of the people get this inference, 30% of the people get this inference, etc." What kind of theory do we need to describe that? I think that so far we do not have an answer. Even Michiel's theories ... he has like three theories — one theory predicts this, one theory predicts this, and one this — but you still can't predict why we have 60% / 30% and 10%. This is something I also have to solve with my kind of theories. My kind of theories cover like the interpretations that 90% of the people get, but there are still 10% that don't get this interpretation. Do we just want to say that this is the error margin that we have to allow for? Or do we also have to cover this somehow? I don't know, but this is an open issue in linguistics I think that nobody really deals with so far.

I: I guess there are maybe two approaches: one would be that you set some kind of threshold, and you say (if the threshold is 90%) that if more than the threshold number agree with my theory then it is a successful theory—

S: Yes, this is something you could do, then you say of course there is an error margin, and these other people for some reason they just don't generate [interpretations] correctly. But for instance, for the data that Michiel tries to describe this is very hard, because we're taking about a 60% / 30% and 10% distribution. Then you really need something else for these.

I: But in any case you're saying that the fundamental criterion for the success or failure of a model or a theory in semantics is the data?

S: Yes.

I: If it agrees with the data or not?

S: Yes.

I: OK. Can you maybe talk a little bit about what are the kinds of data that you use?

S: That depends also on the topic you work with. So when I worked on implicatures there was not that much empirical material available. There was mostly intuitions of linguists in texts that I could use, and then I tried to cover these intuitions that I found in the scientific texts ... which is not ideal, but this is just what there was. For conditionals (the topic I worked on later), there were corpora that you could use, and there were more systematic studies of the data that I could use. Then I relied on empirical studies of other people — on what (for instance) is possible, on what kind of conditional sentences *are* there, what kinds of constructions are possible. But still, these data are mostly on the form of the language, but there was not (is still not) so much available on the systematic studies of the interpretations people get. You are still limited ... still you rely a lot on what people report in their own scientific papers, which is not ideal.

I: I just thought of this, but I'm curious about the fact that at the beginning of this interview when I was asking you if you would describe yourself as a logician or as a linguist, you said that you would describe yourself as a logician rather than as a linguist, because you are not so concerned with the data, but now it seems that this is also of fundamental importance to your work. How do you reconcile that?

S: Yes, it is very important. I'm aware that it's very important, but I don't want ... it's not ... I'm not interested enough in the data to invest a lot of work in getting the data, and I think that's the big difference between me and the typical linguist, because they [the linguists] really want to gather all this information ... I just want to have it, and then I build my model.

I: And are there any more criteria for determining whether the model is successful or not? Other than the data, maybe there is some criterion that it has to fit into some larger picture that we have, or it has to agree with what we know about other areas...

S: Yes, I think there is. This is where I think the connection with cognition comes into the picture. Somehow our theories for language (for language understanding particularly) also have to fit theories for cognition, because these two are interdependent. This is another way to test the theories I think. They have to fit into a larger picture of what cognition is.

I: OK. Do you think that formal semantics is a rigorously scientific discipline?

S: I think it's on the way. It's still very young. For instance, the data ... how to check our predictions and how to check our models using data ... this is something that really just developed during the last twenty years. At the beginning it was possible to write a whole PhD thesis about the intuitions of the writer and his sister about presuppositions, and that was fine, but now it is no longer acceptable. So you really have to underpin all of your models using serious studies of the data.

I: So your answer is that it's on the way to becoming a rigorously scientific discipline.

S: Yes, it's on the way.

I: OK. Great. So I have some final questions about how successful you think it's been so far as a discipline. So just broadly, do you think that formal semantics thus far has been successful in achieving its goals?

S: I think that formal semantics has been successful in turning semantics as a sub-discipline of linguistics into a serious field of study. So in this way — yes. And also pragmatics — yes. The level of theory we have now is much higher

than we had before we started doing formal pragmatics and formal semantics. So in this way — yes, it has been successful.

I: How about in relation to the goals of the discipline? Originally you said it was something like "to understand and to predict linguistic use and behavior"...

S: Yes, of course we improved here. We have better theories that cover more phenomena. So in this way — yes. I don't know why I have to think about that so hard ... maybe I have some more general concerns about the usefulness of this kind of research in general ... for society and things like that. Maybe that's why I'm hesitating a bit on this.

I: What exactly do you have in mind?

S: Basically the line from what we do to the effects that our research has on society — that's a very long line. It's not a very effective field of research in this way. Is it clear what I mean?

I: Yes, you're just saying that it's a bit far removed from everyday concerns of society and practical applications maybe?

S: Yes, something like this.

I: What do you think about the fact that when it comes to applications, the stochastic approaches favored by computational linguists have generally enjoyed far more success than these formal methods?

S: Yes, that's interesting, because that approach doesn't focus so much on understanding, and it focuses completely on coverage of data. In this respect it's much more successful than we are. While we say the benefit of doing formal modeling is to gain a deeper understanding of the problem, but our coverage is extremely less strong than what stochastics offers. What can I say? I mean ... I see the benefits of the stochastic approach, but there are limits, and I think that sometimes now they even reach these limits. There are limits that they just can't pass because they just don't understand what they are modeling. I hope that maybe this is going to be even more pressing in the future, and then it will be more important again to have models that really understand the phenomena that you try to model.

I: So what you're saying is that they have greater success in terms of coverage and predictions, but the drawback is that they offer less understanding?

S: Yes.

I: And what do you think are the kinds of similarities and differences between what you do in formal semantics and what they do in computational linguistics, either in terms of goals or methods.

S: Yes, it's what I said. The *only* goal for them is coverage — to make correct predictions for as many data as possible. And they are very successful in this respect. This is very clear. So that's why they outran the formal models during the 1990s. On the other hand I would say we focus more ... we want to *understand*. We don't want to just be lucky in guessing the right form, we want to understand why this should be the right form. That's the difference.

I: Another thing is that recently there have been some criticisms made about formal semantics by Martin Stakhof and Michiel van Lambalgen. In particular that it idealizes it's phenomena, etc. Are you familiar with this criticism?

S: A little bit. I haven't read these papers.

I: Then I was going to ask if you have any reaction to it?

S: Well, I think there is something right ... [in] their idea of the idealism of the data. That's true, but this is also linked to the general development of the field — to getting closer to the data and really doing serious data studies before you start modeling. Basically that's all I can say right now, because I haven't read these papers.

I: Then my final question to you is: what if anything could potentially cause you to give up your commitment to formal semantics?

S: Somehow I do ... I'm on the way [to] leaving the field. So why is that? Basically because I really ... I'm not interested enough ... on the one hand I'm not interested enough in the data I think as what is necessary right now in the field. So if you now want to do serious formal semantics you really should get into systematically studying data, and that's just not what drives me. That's one of the things, and the other thing is that I'm still not satisfied

with the ... how to describe it? ... with the level of reflection on the theory building and on the methodology that is done in the field. So I want to get deeper than just focus on the data and just building a system that covers these data ... I want to get deeper. So that's why I'm moving to philosophy more and more I think.

I: OK. And you think that's a way to address these concerns?

S: Yes.

I: And do you think that these concerns are unique to you or are they shared by researchers in the field in general?

S: Well, most of my colleagues I think are ... well, there are a lot of colleagues that are still very successful in the field, and they move more and more into doing data studies as well. So this is really the way to go I think ... even philosophers are moving into this field. So I don't see that many people moving into the other direction, even though generally it's very common that scientists later on in their career develop interests in philosophy, but that doesn't seem to be the case for formal semantics right now. The general movement looks more like going into data oriented modeling ... more and more data oriented modeling ... having more and more sophisticated theories for studying data.

I: OK. Those are all the questions that I wanted to ask you. Is there anything else you would like to say before we end this interview?

S: Yes ... No ... Just that I realized that I'm *really* leaving. I don't reflect that much anymore on formal semantics, because I'm not really working on it now that much anymore. So maybe I'm not that representative of my group.

I: No, but I think this is ... I mean you were for a long time heavily involved in formal semantics research.

S: Yes ... OK. I hope it's useful!

I: Well, thank you very much for your time.

Appendix C

Interviews with the outsiders: computational linguists

This appendix (Appendix C) contains transcriptions of the interviews conducted with the computational linguists: Khalil Sima'an, Ivan Titov, Philip Schulz and Remko Scha. The interviews have been transcribed using the following schema:

> I: Lorem ipsum dolor sit amet, consectetur adipiscing elit. Aliquam ultricies lacinia euismod.
>
> S: Nam tempus risus in dolor rhoncus in interdum enim tincidunt. Donec vel nunc neque.

"I" denotes the interviewer, and "S" denotes the interview subject, i.e., the person who is being interviewed. The interviewer ("I") is in all cases the author of this thesis, while the subject ("S") denotes an anonymous member of the set {Khalil Sima'an, Ivan Titov, Philip Schulz, Remko Scha}.

NB: These interviews are semi-anonymous in the sense that the individual subject of each interview remains undisclosed. Therefore the reader is requested not to identify any particular interview with any particular individual in any future citations of this work.

C.1 Interview with a computational linguist

I: Can we begin with a brief introduction from yourself? Can you describe your academic and professional background, and how you got into computational linguistics?

S: Where do you want me to start? Do you want me to start [with] how I came into computational linguistics? Or where I started? If you want that, I can tell you that I started with a major interest in logic. I studied at this ... not at this institute [ILLC], but it's this university [University of Amsterdam] ... I graduated in theoretical computer science with a strong interest in formal methodology in general, and I had many teachers from the ILLC at the time (I'm talking about Anne Troelstra, Kees Doets, Johan van Benthem and others) and I had developed a strong taste for logic, particularly temporal logics and modal logics (there was a moment [where] I even started a PhD in temporal logics), but after one year I decided to move away. This was for the reason that I met also my mentor (he wasn't officially my supervisor, but I didn't have a supervisor to start with, because I wasn't a PhD student) — Remko Scha. When I met my mentor he puzzled me in some discussions with him about ambiguity, particularly about ambiguity in natural language. Whereas all my computer science colleagues and teachers were talking about "how do we get rid of ambiguity?" he was talking about "how can we really make this choice among possibilities in a way that resembles human behavior rather than ruling them out ahead of time?" He was completely into this distinction between performance and competence. And I came into this field out of this aspect — where I got intrigued by how to deal with this cognitive capacity called understanding and the human capacity to choose between different possibilities. So I came from logic, and I just deviated ... I had to bootstrap the statistics and probability theory and all the things that I needed in this field ... the field was young at the time (1994) ... there weren't books at the time so it was a very young field ... so that's a good point to move away from logic to some new field like that because there isn't much, so you can bootstrap it while the field is developing and develop it yourself!

I: What would you say was your primary motivation for doing computational linguistics?

S: The primary motivation was actually being intrigued by the most difficult question that I faced at that time ... For example, when you build a grammar ... if you start by thinking that I have to clean up the grammar from all ambiguity, you come to natural language, and you realize: I cannot do that ... I have to keep the ambiguity. The question is not to get rid of the ambiguity, but basically how do I solve the issue of having to deal with many possible options? How do I imitate human behavior? Natural language processing to me is not a formal game. There is input - output behavior that is registered from humans ... it's a cognitive problem ... and I'm interested in getting a way to understand how do I best model and get a theoretical understanding of what happens when humans do this. Think about different problems — one of them is ambiguity, the other one is that humans are capable of processing non-gramattical sentences (what my colleagues call non-grammatical) ... we don't throw those out the window. Humans are able to understand them, sometimes they are able to correct them. Humans are able to fill in the blanks. They have expectations about what should happen. They're not just processing machines that are deterministically programmed. They can guess. Give a human the beginning of a sentence and they can finish it for you. Give them a gap in the middle of the sentence and they will fill it in. So humans work with expectations more than anything else — expectations in a probabilistic / statistical sense. Given a certain situation you always have an expectation of the most likely situation ... possibly a distribution of the possibilities. To me the puzzle started with a question of ambiguity and robustness (if I may call it that) — of course, it cannot be done with logic — and it extended towards uncertainty very soon. The idea that we are facing uncertainty all the time in language processing. So my interest was driven by very very basic questions rather than an agenda that I set ahead of time.

I: Would you describe yourself primarily as a linguist or as something else?

S: No. Over the past years I have developed (I didn't have them when I started) ... I'm a computer scientist at the end of the day. I started as a computer scientist (a theoretical computer scientist, a formalist). I developed my interest in linguistics and in general also the same kind of / ways of thinking that are very similar (not the same, because there are major differences, but very similar) to my linguist colleagues. I developed those over the years. I

actually bootstrapped all the linguistics that I needed at the time, and then I contented studying while working and doing research. In fact, I must say that I described linguistics in one of my invited talks once (about the current situation in computational linguistics) as: There is a huge pudding, and there are sprinkles on the pudding. Computational linguistics nowadays is a huge puddings, and there are only few sprinkles (decorations) of linguistics on top. It hardly plays any role. The pudding is the statistical modeling. It is the jelly — it moves, it has the capacity to stay together. And linguistics is the sprinkles on top that basically cannot hold it together, but they decorate it (and they are tasty) and hopefully someday they will do something, but not yet unfortunately.

I: OK. So that was about your background and motivation. Then I have some questions about the field of computational linguistics in general. To begin with, how would you answer the question "what is computational linguistics?" Suppose you had to explain it to a university student who has never heard of this discipline?

S: What is computational linguistics? Computational linguistics in general you can say is the field of modeling language use and producing technology out of it. It encompasses the whole range from taking cognitive aspects where you have phenomena that you observe in language use, and then produce models that can be tested against human standards ... down to the engineering and technological uses of that.

I: What are the goals of this discipline?

S: The goals of computational linguistics ... well, if you ask different colleagues you will get different answers, so I'm going to answer first on my behalf and then in general. The goals in my personal opinion are to produce models that can really process language in certain settings ... in certain settings like when humans translate and summarize — particularly I talk about those, because that's where the aspect of meaning comes back — where humans show input - output behavior, and we are interested in understanding how the input - output behavior comes about. And at the same time producing a model that we can test, not only understand but a model that we can test. As I said before, other colleagues might say that we are really interested in producing technological

artifacts. I don't rule that out, but I see that as the second step. If you ask other colleagues about computational linguistics and what its goals are, it could be that they will answer something else than myself.

I: What are the kinds of problems that researchers try to solve in computational linguistics?

S: There is a whole big range actually. Nowadays, because of the internet and because of lot of other stuff that happened, people confuse what they are doing in natural language processing with computational linguistics. In computational linguistics, if you want to take the narrow sense, you basically look at problems that have to do with input - output behavior of humans and trying to model it with tools that take into account the possible hidden structure between input and output. In many cases like information retrieval, which is also sometimes thrown under computational linguistics (I think that's actually natural language processing which is a broader area) ... in information retrieval there are a whole range of tasks which are not done by humans, but that concern human language, they involve human language. You don't have input - output behavior from both sides, but you have only one side, and the rest is a guess. In terms of how useful it is, it's extremely useful, but in my personal opinion that doesn't belong under computational linguistics per se. Take for example the origins of computational linguistics: machine translation is the starting point for computational linguistics. Syntactic parsing, semantic analysis, semantic modeling ... there is a range of aspects that have to do with what is a historical artifact of how linguistics has evolved over the last 60 years. It's distributing the processing tasks among the subcomponents of the processor: syntax, semantics, morphology, etc. Computational linguistics is busy with those subcomponents. Basically studying them in a computational way, and this always means (nowadays) that you have an input - output examples, and the question is: can you really get the same behavior as humans? But then you have tasks like summarization, paraphrasing, machine translation ... all kinds of tasks that humans on a daily basis do, and we would like to be able to have this human behavior.

I: So these are the main kinds of problems that computational linguists are trying to solve?

S: Yes.

I: And when would such a problem be considered to be solved? You're saying it's basically when the input - output matches what the human would do?

S: Yes. The difficulty here is that you have to realize that language processing happens in so many different settings and situations ... we will consider the problem solved if we can imitate human behavior (in this sense of input - output) in all these settings. And then the questions I'm sure are not done yet, because I'm talking about a black box: input - output. If you look at input - output one step further, you can say I'm interested also ... I know that humans process certain sentences faster than others for example. So you're not interested only in the fact that this is the input and that is the output, but with certain behavior that you consider as part of the output, and there it's processing speed. At the moment we're not worried about that. We just want the input - output. But we might get more interested in certain aspects of the processor, in imitating those as well, because it is interesting scientifically in its cognitive aspects. The problem will be basically solved ... one problem will be solved if you can define it as an input - output mapping with the gold standard being human behavior. But of course that is a moving target, because human behavior shows so many other aspects that at the moment we are not worried about / we are not occupied with, but we will need to look at at some point. Once we solve it we are going to start with the real things, and there are always more real things, so it's a moving target.

I: Is there general agreement among computational linguists about the methods of solving these problems and the criteria for when we say this problem would be solved?

S: Yes. The field has developed over the last 20 years a methodology of testing where you're standing in terms of modeling. You are all the time worried with checking how well you're doing at certain tasks. So there are benchmarks for experiments and for testing where your model is; there are competitions even. And in fact we know all those metrics that are being used for evaluation are actually part of the story, but we know that if we do not adopt together a certain training data and a certain test data that we all agree upon, and certain metrics, then we will not know whether we as a field are moving forward. We

will continue arguing "I'm better than you or you're better than me" without any objective evidence. So the field has adopted these tools with all the difficulties of adopting tools when you start, because we know these metrics are not sufficient for measuring everything. We know that they are automatic metrics, which means that they are not like human judgments. But we can measure correlation, and we do measure correlation with human subjects. These days for example we collect — in machine translation, it's a collaborative work ... multiple teams together — they collect huge human judgments of machine translation output, and basically this is the gold standard, the human judgments. What you want is that if in your system you have your own automatic metric to measure performance of machine translation output, then you can really measure the correlation with human judgments ... to see if you are not accidentally judging but really kind of correlate well with what humans are saying. So this methodology has developed over the last years, and it's a strong methodology for the field, because it's actually helping the field move forward by knowing also what scientific ideas contribute to certain practices that give improvements ... we have to discriminate between the practices and the fact that the field also develops at some point habits, because they are expected to give a measurable benefit. We have to discriminate that from the actual scientific discoveries. Scientific discoveries might be slower. What you see in the field these days is that these measurements are being done, and we know that it helps the total field, but isolated researches have to continue with their own agendas regardless of whether this year this scored top or last position. Because they believe that they have something that is in the long run going to be better. So there is a whole range ... on the one hand there is the community that develops this, and you want to comply with what the committee expects of you, but on the other hand there is the old fashioned scientific, linguistic maybe habits of developing your own ideas and working on them despite not seeing improvements everyday.

I: But in terms of the actual models that you're building. Are there standard criteria for success and failure of these models?

S: There isn't success and failure as in 1 / 0. It's usually relative success or relative failure, relative to what exists as the best practice at the moment. So what you see is people and all their publications ... what they have is a

baseline model. The baseline model is the model that most people would think up today and would use for this problem or the model that is being used by everyone anyway. For example — with one of my PhD students — we are developing models for translation under specific domains. The domain is a latent variable, so we learn to discriminate between domains from the data; we don't stamp it in. The field has developed the habit of working with models that are not sensitive to domain aspects; so they train on this, they test on that, and if they are not aware of the domain difference they will get certain behavior that differs from one domain to another. So if you give it different test sets — one coming from news, the other coming from sports, and so on — it will do well on one and fail on the others. In our model ... to test our new model which is sensitive to these domains and can learn to discriminate, we just simulate the situation: "Ok. This is the best system these days. These are the best models that exist these days." We implement them in the system and we compare to them. So for human evaluations ... we just developed a metric for evaluating the output of machine translation systems. What we do is we compare to the best existing metrics that exist so far ... 10 of them — we have to implement them all and compare to them all to show that we have some benefits. We do that by measuring correlation to human judgments. So we have the data of how the humans have judged it (as the gold standard) and there are the other systems / other metrics that are being used so far (they do the judgments and they correlate that much with human judgment), and here is our new metric and we would like to also show that our metric imitates human behavior with correlation much better than any of those. So it's a tough and slow development in the field, because of these strict rules and strict requirements. You can't just say the model is more elegant. You can say that, but nobody will listen to you basically.

I: But there are basically some strict requirements that your models have to adhere to?

S: Yes, yes. No, no — the models themselves you can think them up as you like, but the testing [has strict rules and requirements]. Once you put them into practice, you have to show that they have certain input - output behavior that is expected to be hopefully better (or at least as good as) what exists. If they are not better, you still can actually argue that it has certain properties

that are attractive, but it shouldn't be so bad in terms of performance that you cannot justify it. You can justify that it has certain nice properties, despite performing slightly good less than the best practice, because for example it promises a great future, a new research agenda, etc. These things are also acceptable in the field.

I: And do you think the field has been successful in achieving its goals?

S: I think we've been astonishingly successful over the last 20 years. I've been in this field 20 years ... I didn't plan this ... when I started I came into this field and the field just started really to tick, and I was too young to claim everything for myself! But basically in 20 years we managed to (I claim) even to change the world. These models that we even thought up here together with other groups, they're basically the vehicle for things like Google Translate, you know. Machine translation was thought to be impossible for the coming 50 years if you would have asked anyone in 1990: "when will we have a machine translation system available for everyone?" Even if it's not great these days, not yet doing everything superb, but it's really amazing that we got there, and it's actually being used. There was a time ... I don't know if you know the history of machine translation ... the history started form the 1950s ... it was the beginning of AI, and big budgets went into that in the USA, and it was a big failure. People were busy with Shannon's ideas, and then I don't know if you know the story, but Chomsky won. And then they tried the Chomsky way, and then after years they said that this is going nowhere ... there was a congress report about this in the USA which said that this is going nowhere and we have to cut down the budgets. And the field of computational linguistics would have died, because it was called the field of automatic translation. Computational linguistics was born out of that moment where machine translation was considered dead. So if you want to measure success, you could say in the last 20 years we managed to produce something that would give evidence to society so that they think "OK. These are not only money losers. They might give something back." If that's the measure of success, OK. If the measure of success is different in the sense that "did we really get more understanding of what's going on in the field?", I claim also that we are in a much better position today than in the past. But we're not yet at a point where we can formulate a

single theory for the whole field. There are small theories, more in the form of models these days rather than one big uniting theory.

I: Then finally I have some questions about the comparison between computational linguistics and formal semantics, which is also popular here at the ILLC. So what do you think are the similarities and differences between these two disciplines, in terms of goals and motivations or methods or the way in which they approach new problems?

S: Similarities?

I: Similarities and differences.

S: I think the differences outweigh the similarities by far ... and this is for various ... it has to do with the subject of study. I mentioned the fact that we are interested in performance, and when I say performance I mean not only looking at the data (staring at it and formulating what we think it should be — the logical formula), but actually really taking care that we have a computational model that imitates human behavior in terms of disambiguation ... in making the right choice given certain circumstances. We care to build a complete model of behavior rather than isolated pieces of behavior, so you can see it that way. What I mean by complete — complete in the sense that it has to have the input - output behavior rather than just naming the possibilities (like here is an input and here are the possible representations of its meaning — we're not interested only in that, we're interested in the next step and that is "which one is going to be chosen by humans? Given a certain setting / certain situation"). So the starting point is already different. We do share with our colleagues the formal background. We do share with them the interest in terms of having representations that are deeper than just surface representations ... it's not only what you see, there is more in between the input and output, there are structured representations in terms of trees and in terms of graphs. So at that abstract level we share a lot. We (computational linguists) tend to consider those representations as representations that can be used to start the learning process from data, where we learn distributions over those representations — probability distributions, such that we can actually weigh the possibilities and also smooth the number of possibilities, because we don't want to have this crisp behavior where we say it's either grammatical or it's non grammatical, but it's

shades of graded grammaticality. So the differences, as I started saying, are bigger than the similarities between the two, but I do see at the moment some convergence (in semantics particularly) towards working on problems that have not been tackled before. Having the role for semantics in machine translation for example. I know that my colleagues are interested in that — not only at the ILLC, but semanticists in general — simply because: They realize that the data we have these days ... for example in machine translation, the huge data we have shows you a sentence and it's translation. These pairs are actually pairs where the meaning has been trapped in between, and it's actually equivalent (approximately equivalent) on both sides. What better evidence can you find from human behavior for how meaning has been preserved than that? So even our colleagues who have been working with logic are getting more interested in that, because of realizing that this data is so rich that it might be a good resource for them to do their own studies.

I: OK. When I asked this same question to formal semanticists, one common response I get is that they say the stochastic models favored by computational linguists, although they have greater coverage of data and greater predictive power, they do not contribute to a deeper understanding of natural language and they don't have much explanatory power. In the words of one formal semanticist, "computational models just don't understand anything." What would you say about that?

S: I would say like ... as a physicist ... I mean, I would just say the same thing that a physicist would answer: Look, you don't know what it means to understand what's going on. When you understand a phenomenon, you actually can predict what's going to happen. Understanding demands from you to predict human behavior, because that's the gold standard. We don't just predict what we like ... in the models that you develop, if you say that you really understand what's going on, then the starting point for understanding is to be able to predict input - output behavior — that's one thing. And second, if you're talking about semantics, I dare claim that none of the algebraic representations that we've been occupied with over the last 40 years or 50 years in semantics are actually any kind of evidence for understanding the phenomenon of human language processing, because it's only accidentally related to that, in the sense that it is formulating certain constraints on what is possible. It

doesn't really say what is the human behavior ... what is the plausible human behavior going to be given a certain situation. Do you see the distinction? It [formal semantics] formulates constraints on what is possible, rather than taking care that here is an input, and predicting what a human is going to do. So that comes back to the issue of disambiguation. So I dare claim that the algebraic representations that we've been busy with—

I: You mean the logical representations?

S: Yes, yes. I talk about representation, but I'm not talking about particulars. It's general. I'm generalizing even beyond just using logic. How we get to these representations I don't care. I'm talking about the final goal for my colleagues, and that is to get to these representations. The representations are just not sufficient. They are not evidence of understanding, and I think they're basically completely understanding what you like to understand. It's in the eyes of the beholder in this case, rather than really saying anything about human processing. You want to understand? You have to do like physics does: You have a phenomenon in nature, and you are modeling that phenomenon. To model that phenomenon you need to do measurements, and the measurements you have to do not by choice (choosing one measurement over an other) but really you measure whatever. You run an experiment, you repeat it many times ... many, many times ... the same experiments under the same circumstances. You isolate certain effects, and you take what you put in and what you got out of the experiment as measurements. And you are basically taking care to model these two — inputs and outputs. That's in my opinion really science, and also kind of understanding, because then you might be able to take a theory that you have in your head and show that it is predicting experiments. That shows understanding or non-understanding.

I: So do you think that these two [formal semantics and computational linguistics] are somehow competing disciplines?

S: Well, at the moment there is no competition — we basically won! [[*laughs*]] I'm joking; there are no winners — it's not a football game. But really I think computational linguistics has developed in such a way that these days we are at a point where we (or a majority of my colleagues) do not believe we need a linguist at all. So it's called computational linguistics, but it's an accident. It's

actually not computational linguistics, it's computational-linguistics in one go with a dash in between. It's one name. That's how the field now looks at itself — detached from linguistics. There is approximately no relation these days like in the 1970s. We are not people who are taking linguistics and making it computational. Some people still think that that's what we're doing, but that's not happening anymore. That's not what the field is. The field has developed to be a field of it's own. It's standing on its own feet. It's a strong field. It's a field that is one of the strongest at the moment in terms of credit ...

I: Strongest in linguistics?

S: No, in general. I dare say that if you look at the USA and Europe, it's one of the strongest growing fields in terms of attracting budgets, interests, etc. simply because we have shown over the past 20 years when we got rid of the linguistics infusion. How do you call that thing in the hospital where you are sick and they feed you by the line? ... we got rid of the linguistics line and we went on our own. That's the point where we had to survive, and in fact the field has grown and outgrown its mother fields of linguistics and AI. We are, I dare say, bigger than those even — more successful. All companies these days like IBM, Google, etc. — the biggest ones — they have teams called computational linguistics teams. They have adopted computational linguistics. In the past these companies would say "A linguist? What do I do with a linguist? Why should I hire a linguist?" Companies are not a gold standard for me for science, but this is the result of the last 20 years. We have shown like physics that we can also produce models that show predictions. These companies happily take these models and put them into practice and they benefit from them. We get no money, but at least there is recognition that this field is like physics. It's giving models that perform [with] physical phenomena and are useful at the same time. We have to reflect on what is understanding at the end of the day. What does it mean to understand a phenomenon? I am sure that philosophy has put a lot of effort into that, and I know some of the basic literature. I think really this is a contrast between the fields in general ... this is a traditional contrast ... even in physics itself there was a contrast at some point between the stochastic modeling people and the traditional physics people. This contrast I think never disappeared. But the ones who've been setting the agenda are the stochastic

people for the last 100 years. It's not going away. People thought it's going to go away, but it's not going away. It's not an accident.

I: Then my final question to you is: what, if anything, would cause you to give up your commitment to computational linguistics?

S: Well, I don't foresee anything to be honest. I love what I do. I enjoy it so much, and even with my PhD students ... I don't think ever ... I see one thing ... Look at it like this: Computational linguistics looks at human processing at a higher level (in terms of understanding) than linguistics. Linguistics has set itself at a stage where it says language is different from anything else. Language has its special aspects. Language processing has its special aspects. But, nothing is sure that language processing is different than visual processing, different than other modes of processing that humans do on a daily basis. Even performing tasks where uncertainty plays a big role. It's not clear that language is too much different from those. So, I dare say that we've gone up to a level where we care about human capacity to process under uncertainty in general. And I may constrain myself to visual and language processing ... so, if at all I get interested in something else, it will be because I know that I can do more. Not because I'm fed up with computational linguistics or I've had enough of it. Because I know that with the same tools we can change other fields and with the kind of experience we have these days.

I: OK. Those are all the questions that I wanted to ask you. Is there anything else you would like to add before we end this interview?

S: No.

C.2 Interview with a computational linguist

I: So let's begin with a brief introduction from yourself. Can you describe your academic and professional background, and how you got into computational linguistics?

S: Yes. My original background is actually in applied mathematics and computer science. That's the degree I got in Russia. I did my undergrad and masters in applied mathematics and informatics, and my masters thesis was on computational complexity theory. But I think you might be aware ... it was something like 10 years ago or 12 years ago ... the situation at this time in Russia was pretty bad in academia, so I was starting to look at this point for jobs locally. I had been working for a while for a software development company and then realized I'm not terribly interested in this on the one side, and on the other side I realized I don't feel terribly comfortable about working in theoretical computer science, because I had a feeling I'm losing relevance to real problems. It might be just specifics of Russian school in this area, because often you even don't try to think how this problem relates to real life ... you just think there is this theorem ... for this set of problems we have some properties like computational complexity classes or approximation complexity classes, and then you try to say "OK, if I add a few conditions to this problem, how would it change? Would you be able to come up with an efficient algorithm?" but you're not even encouraged to think how relevant these extra conditions are to anything in real life. So basically I got into computational linguistics, because some of the things I had been doing for my work had to do with processing language (obviously very basic things), and I was also trying to study how to do this properly more from an academic perspective. Then I see a job announcement and I just applied ... it was a little bit not terribly serious how I did it, because I didn't look for different positions; I just saw one announcement, I got a job, and I went with it. Then I started work in computational linguistics and my work was on syntax. My thesis was on syntactic parsing with latent variable models. At that time it was still not as terribly popular as these days. Then at some point in your thesis you realize that you want to do something else as well, so I went for an internship in Google and started working on non-parametric Bayesian models, and we tried to look

into something like sentiment analysis problems. Then, as many people in the community, I got interested in problems relevant to computational semantics. Recently my research is mostly on some form of computational semantics.

I: What would you say was your primary motivation for doing computational linguistics?

S: I guess that if you are not really that interested in cognitive aspects ... and I guess my motivation generally has to do with building systems (so an engineering motivation), but I also don't want to create small hacks for this. I am trying to have nice models which are useful in the long run in practice. I'm ok with having some models which might be not tractable these days, algorithms might be still not quite efficient yet, etc. but I still want to have practical applications in mind and not become completely disconnected from this.

I: Would you describe yourself primarily as a linguist? Or as a computer scientist? Or as something else?

S: I am certainly not a linguist, because I have no formal linguistic training. I was attending a few classes on linguistics when I was starting to do NLP more recently, but basically I am a computer scientist.

I: So that was about your background and motivation. Then I have some questions about the field of computational linguistics in general. To begin with, how would you answer the question "what is computational linguistics?" Suppose you had to explain it to a university student who has never heard of the discipline.

S: I think it's a very tricky question, because many people in our community perceive it very differently, but to me it's a study of models of language (computational models of language) or development of methods for natural language processing. So it can be either you're trying to use computational methodologies to answer some linguistic or cognitive problems (that's one side of this), or you can try to develop models which are applicable for practical problems (and then they are of course slightly different than the ones you may be developing from a cognitive perspective), and [or] you can be somewhere in the middle (e.g. prior knowledge in cognition helps building applications, or you believe

that linguistic prior knowledge should be in some way helpful). I guess most of us to some degree would agree with this — it has to do with either questions about language, or problems which have to do with language and use computational methods. Of course much of this in our community is about statistical modeling these days, but I guess that's just because it works.

I: What are the goals of this discipline? You mentioned that maybe there are two different kinds of goals — one is about answering questions in linguistics and one is about applications. Can you elaborate on this a little bit?

S: I guess there are many sub-goals I think, but to me the main goal is going into the direction of natural language understanding, and I think many people in the community think about problems in this direction, but of course that's not the only direction even in the more engineering oriented sub community (because you can think about natural language generation, etc.). Some people in the community are more cognitive ... people, for example, think about problems of language acquisition, understanding how children learn languages, etc. Also about how linguistic structures in different languages are related ... and there are many questions you can ask from a linguistic perspective ... basically understanding how language works.

I: What are some of the kinds of problems that researchers try to solve in computational linguistics?

S: I guess, to be more specific, we can think of syntactic parsing. It's still a challenge, and from this perspective, the fact that much of the research has been done on English is a little bit unfortunate, because English is a bit of an idiosyncratic language. So there are many things we don't know how to do yet with syntactic parsing, and certainly syntax seems like an important step for doing semantics and for doing many things properly. Then there are some more practical but still challenging applications like machine translation, summarization, etc. They bring in many interesting problems from statistical modeling perspective, from linguistic perspective, from algorithmic matters as well. So that's really kind of the important end applications I would say. Then there are questions about natural language understanding and how to proceed with this, and there is less agreement on how to go on there, because some people look into very expressive formalisms and logical representations, but

then they are restricted usually to very simple data sets and somewhat not so interesting sub problems. But in other cases, you can think of, for example, things like frame semantics, semantic role labeling as a type of semantic representations. Then of course, on the more linguistic side, you can think of discourse processing, because much of natural language processing (until a few years ago, and still) looks only sentence by sentence, and that seems like a very wrong assumption. There are really many problems, so different sub communities work on different problems. I guess for computational linguistics, because still much of the community has engineering goals in mind, doing completely theoretical research (and not showing that it is relevant to practice) is not really appreciated. Usually to get a paper at one of our top conferences you need to show not only some theoretical framework which can cover a few examples or some very limited phenomena ... you need to show that this actually is relevant to some kind of at least semi-real applications. So that's how it works in computational linguistics these days.

I: In terms of these applications, what exactly are the kinds of applications that you are trying to apply computational linguistics to?

S: I mean machine translation in itself is an application. Then of course question answering, but there are different types of question answering applications. But there are some kinds of applications which are more like tests rather than real applications ... so for example textual entailment is one application where you have a passage and you have another passage, and you're trying to test / predict if one entails another. There are few problems like this for testing, for example semantic similarity ... it's not really the end task ... but you can have a human judgment on this and you can see how well your model mimics this human judgment ... and intuitively you can hypothesize that this information should be useful for applications, but you're not really building a real application. The thing is not all of us have huge groups so that they can really focus ... not all of us really want to do engineering to get for example question answering ... it's a tricky task, because there are very many engineering aspects which are not so much interesting from an academic point of view. But some people collaborate with industry — so I'm trying to collaborate with Google on some things ...

I: I'm interested to know — when you have a problem that you're working on in computational linguistics — is there some kind of standard conceptual framework that is used for solving problems?

S: I guess it depends what you mean by that — let me answer and if it's not what you want you'll help me. Generally we are really focused on the ways we can evaluate our techniques. So we don't want to pick some examples and try to see that we cover phenomena. It's usually the set up that you have some representative evaluation test set ... and for your problem, maybe annotated by human experts (by linguists) ... take syntactic parsing for example: You have some standard tree bank, you have your data set for training your model (and maybe some held out for development so that you don't test too much on the final test set). And then how you usually proceed with this ... you should have a very basic model as a baseline for your approach and then it kind of gets a little bit incremental. So you have some theory of how to improve this baseline, you test this theory, you look into what kind of things don't quite work out as planned, you refine your model. Or at some point you might abandon because you realize your idea was completely rubbish (it doesn't have any relation to what happens in real data). It's kind of cycles of this. Usually you have some general ideas of what hasn't been done quite right in the past, and in some cases it's useful to look at what kinds of mistakes previous methods make and see if you can come up with some model which seems to be solving these issues. So for example for basic syntactic parsers which don't use any words in the grammar, you can see that they are not very good — for example it doesn't do prepositional attachment problems, and then you start thinking how would I go about it? (This example is from twenty years ago). Basically I guess the difference is that you really need to show that your changes to the model have some effect on a representative sample of the data. Generally you are not supposed to choose examples with a phenomenon you like, and then show that you are doing something reasonable on that. Generally you start with the data set.

I: Would you say that there is general agreement on this procedure among computational linguists?

S: Specifically which data sets to use? You mean exactly or what is the general procedure?

I: In terms of the general procedure.

S: In terms of the general procedure it seems that there is a general agreement. Some people don't quite agree, but I think majority of the community would agree on this. Of course, the question is that you might be looking at some of end problems and you're really very practical, but sometimes you are interested in some subset of problems which you can really identify and then you look only in the subset of the data (you look at parsing questions only for example or parsing only wh- questions), so you can just look into that data set and you don't care about anything else; it's still quite ok. In some cases you might just look into something which is not connected to the practice, so you leave it to someone else to solve this problem — you want someone else to fill this gap, but you think this problem is meaningful in itself. I think that's also reasonable (if the problem is meaningful). But in general I think the data driven view is what prevails.

I: What would you say are some outstanding or innovative or pioneering achievements in this field?

S: I guess it depends on which level you judge, but of course from application point of view the big success is in machine translation. I would say there were big successes in syntactic parsing — so now we can fairly reliably and accurately (and really fast) parse many languages. Of course parsing is not the end application, and many applications still don't use syntax so much, but I guess it just needs some time. On the application side I think that's the thing. Then of course there is little bit of movement now in the direction of semi-supervised learning, because people realized that for interesting problems (for learning semantic representation for example) it's a reasonable way ... we need to learn from data ... this is something that is really taking place these days. So people start looking at very huge data sets and you see that you can perform a lot better on many problems. For machine translation it has been known for ages, but for some problems it's happening now ... semantics I think is one of them, because it's so tied to lexical information, so you can't really hope to have enough annotation to do it well.

I: In terms of the models that you build in computational linguistics, how do you actually check to see whether it's a good model for the thing you're trying to model?

S: The standard way is to evaluate it on the data and to see how well this model works. Of course, I think you should also try to apply simple models as you can. You don't want to include many hacks, etc., otherwise you have no idea why it works, how to apply to different languages, etc. But I wouldn't say that all the people in the community necessarily follow this, but I think many good people do.

I: Are there some uniform standards for success and failure in modeling?

S: For good and bad reasons, in few areas there are standard successes. For example for syntax (for English and for a specific type of syntax — constituent syntax) there is a very very standard benchmark. There is a split into training, development and testing sets, and there is a specific evaluation methods to be used. I think it was really helpful for a number of years, because you can really see ... comparing to previous work you know that the reason is not that they used different data or that this guy just didn't manage to replicate your parser correctly. You can just look it up in the paper and see. But of course there are some problems with this approach. It is stuck in a bit of a local minima to some degree, because reviewers in the top journals and the top conferences wouldn't accept your paper if you don't get state of the art results, if you don't improve on this core. I think it's even more the case for machine translation, but probably Khalil told you more about it, because that's his area. So in a way it's good, but of course it makes it hard to propose something (I would say radically new but) really quite different from what has been done in the past. That's the problem right now. For more well established applications (syntax, machine translation) it is the case. For semantics, since there is no very good agreement on what we should do there and what kind of formalism we should use, things are a little bit more messy, because different groups might come up with their own data set (and they might have a good reason to criticize other set ups), so it's hard to really track success ... but I think it's just a period ... it's going to converge to maybe not one evaluation, but at least to a few standard evaluation set ups that people stick to.

I: Overall do you think computational linguistics has been successful in achieving its goals so far?

S: We have a number of real applications. We can see that many things we developed in one form or another are used in industry, so that certainly shows there is some significant success in these applications. But ... of course there is a bit of a but here, because in many cases the models which work almost as well are really dumb models. And some of the things we often develop (I'm thinking of linguistic representations) just give us a little bit of improvement over these dumb models. Then in this sense, some of the things we develop don't have that much impact yet, but I think this happens for many other fields as well. To reach some performance is easy, but for improvement on this you have to take more small steps forward. But yes, as I said, I think there are obvious successes. One success in machine translation. There is also certainly success in question answering, information extraction, etc. When you type your Google query there is quite a bit of NLP going on there on this side also.

I: So that was about computational linguistics, and now finally I have some questions about the relation between computational linguistics and formal semantics. To begin with, what do you think broadly are the similarities and differences between computational linguistics and formal semantics? In terms of their goals or methods?

S: I think there are similarities in many ways ... I work on probabilistic models for semantics, so we also think about language understanding. But in general, unlike formal semantics we are maybe slightly more pragmatic. Because we know that we are not yet able to do complex things ... so we will (at least from my perspective certainly) go to the stages where we will need more complex semantic representations, but we are still not capable of doing even much simpler analysis like semantic frames for examples — that's a basic problem. We are a little bit more engineering oriented and data driven, so we see that we are not solving many simpler problems, so it seems premature to us to solve some very complex issues since you can't really solve them anyway probably. But I guess we are making some progress, and we are going to (hopefully) converge a little bit more with what's going on in formal semantics. When it's going to happen and how it's going to happen I have little idea, but still if you see what was happening 10 years ago — semantics wasn't a big thing

in computational linguistics at all because people were saying many things are just too hard and it just doesn't make sense to do semantics ... let's sort out syntax, let's sort out morphology and other simpler things first. So I think we are now at least looking at similar problems.

I: What about in terms of approach? The way in which you approach a new problem. How do you think a formal semanticist would conceptualize and think about a new problem as opposed to a computational linguist?

S: I think many computational linguists — at least if you really think about computational linguists, because in our community there are some linguists (not so much computational) that are a minority these days — in general we often try to have models which are not quite right, but which have nice computational properties and cover the most important aspects. If for example I would be to model some linguistic phenomena like alternations of verbs, I would think a little bit about this and I would think what are the most important aspects, but I wouldn't necessarily try to create a model which will be super plausible from a linguistic point of view. I would think about algorithmic constraints from a learning perspective that maybe there are too many parameters, too many things you need to estimate, etc. In this point of view we usually start from simple things and go into more complex things step by step. That's the approach which is more appreciated in our community. If you see that capturing some linguistic (or specifically semantic) properties doesn't help (even if such a model might be more correct from a linguistic point of view), you just remove this component and you don't care about it. So I guess that's the point. I can basically think of myself also as an applied machine learning person, so for me I always keep in mind properties of the problem from a learning perspective. That certainly may be less relevant to formal semanticists. But recently I am also trying to collaborate, since there are many points of interaction between formal semantics and statistical people.

I: Like what?

S: So right now I'm collaborating with a team doing natural language generation. Natural language generation is I think much more connected with formal semantics. In the standard one they use very little statistics, they basically just

consider logical form. But there are many issues ... for example if you're trying to generate a dialogue or in our case instructions in a virtual environment, there are many questions of how you generate it so it's easy for a human to understand and you can model the interpretation of this utterance by human beings, potential confusions, etc. ... the models we developed in computational linguistics are super helpful, but at the same time the starting point is usually formal semantic representation, because that's what the system wants to verbalize in this case. So I think natural language generation is one of the cases where the synergy is obvious right now and we can do many interesting things together.

I: When I asked this question to formal semanticists, they often say that the stochastic models favored by computational linguists, although they have greater coverage of data and greater predictive power, that they don't contribute to a deeper understanding of natural language and that they don't have much explanatory power. In the words of one formal semanticist "computational models just don't understand anything." What would you say about that?

S: I don't think it's necessary that to perform human-like inferences you need to have logical semantic representations. I'm not sure that understanding conceived of as writing some specific logical formalisms is necessarily the only way to go about it. I think it should be quite useful, but there are certain limitations. Also in a way, stochastic vs. formal is maybe the wrong antagonism, because of course you can combine the two, but I agree that in the computational linguistics community we often have models which are not necessarily very interpretable, so it [the model] learns something and it's sometimes hard to understand what specifically it learned. In recent trends with neural networks, etc. it's getting even worse in a way, so it's getting even harder to understand than it was [earlier] in supervised models, linear models and non-parametric bayesian models, where at least you can do some analysis and get some intuition of what it learned. Even so it's not going to be perfect. It's not going to be axiomatic and it's not going to be in any way very formal, but at least you get some understanding. Of course that's an issue, and if you use a lot of data then there are many surprises. I talk often to industry people to get some data from them on a very large scale ... you get some clusters of phrases and

clusters of words ... sometimes they seem to make no sense to you, but they really work very well. It's a bit of a puzzle. I kind of agree that interpretability is an issue, and it's also an issue for us, because if you can understand what your model is doing, you are better at refining your model. But there are some partial substitutes for this. For a statistical model for example, you can sample the data, you can see what kind of errors the model is making ... so we have some other ways to figure out these issues. But the main goal is slightly different than formal semantics. What does it mean to understand the language? For us it's more like getting a model which is able to understand the language and which is able to make inferences similar to humans or at least useful for applications, and we don't care about some small parameters filled in specific matrices.

I: Do you think that these two disciplines — computational linguistics and formal semantics — are in any sense competing disciplines?

S: It's hard to say. I guess to some degree they are, but it depends what you mean by competing. They have different research questions of course, so that's a good thing and a bad thing. The bad thing is that it discourages us from collaboration in a way which I think is the way to go, but since we care about different outcomes, research is a bit tricky. From an applied point of view ... to be honest, my view now is that you can't do much without statistical modeling in some form. In statistical NLP, if we use some form of formal semantics, it's mostly too simple for formal semanticists to be interested in. For example, I would highly suspect that a question answering system like an Android or iPhone [uses interesting formal semantics] — they probably do use formal semantic representations of questions (because they need to map it to APIs, commands, etc.), but certainly not something too complicated to make it interesting from a research point of view in formal semantics. But I guess there are some areas where there is a bit more competition. If you think of generation of dialogue systems and that's where there is scientific overlap and collaboration, and maybe a bit of competition. But in machine translation nowadays no one is dreaming of doing translations through logical semantics. Formal semanticists probably have no idea how to do it, so it's only going to happen in statistical modeling. If you think about extracting knowledge from a noisy web, that's probably also the way to go.

I: My final question to you is: what, if anything, could possibly cause you to give up your commitment to computational linguistics?

S: You mean switching fields?

I: I mean is there anything that could make you think that "maybe this is the wrong approach" or something like that?

S: The term itself and the field itself of computational linguistics is changing so much, so it's more likely that computational linguistics will change in such a way that it will become something completely different. For me it would be very frustrating if we learn that some kind of more interesting linguistic representations end up being completely useless for real applications for processing language — that of course would be quite frustrating. I'm interested in something which is still rooted in some way in linguistics, inspired by linguistics. If something super flat (an unstructured model) using simple ngram combinations and maybe applying some kind of neural network on top of this works better than anything, that might be slightly frustrating! But I don't think that's going to happen first of all, but other than that I think we have too many interesting problems to abandon this field. I think there are many exciting things which are going to happen.

I: These are all the questions that I wanted to ask you. Is there anything else you would like to say before we end this interview?

S: I don't think so. I think in general you should talk more to formal semanticists about the fact that they don't often come to our conferences, etc. Maybe that's not right, but maybe you can blame us for not learning enough formal semantics to be able to meaningfully speak to them. That might be another side of this.

I: Alright. Thank you very much for your time.

C.3 Interview with a computational linguist

I: Let's start with a brief introduction from yourself. Can you describe your academic and professional background and how you got interested in computational linguistics?

S: Sure. I started out as a pure linguist, and I also did some formal semantics during that time. But I actually also had a big interest in phonetics, which is probably the most empirically well founded sub-field of linguistics, because in phonetics since the beginning of the 20th century they've been conducting experiments, which only starts happening now in other branches of linguistics. While I was doing my master's, we actually had a group at my university which is one of the best groups in psycholinguistics ... and also the evaluation of psycholinguistics involves a whole lot of statistics. I've always been interested in the experimental part, but I never could really evaluate my experimental results because I didn't know the statistics. When I came to Amsterdam I started doing machine learning — machine learning is essentially just a fancy name for applied statistics. So I did that, and I saw that it was very well connected with my interest in language, so step by step I went into computational linguistics. I should say that I really think computational linguistics is just a particular methodology within linguistics.

I: What would you say was your primary motivation for doing computational linguistics?

S: It has a clear methodology and it has certain research standards that you can follow, so it's much easier to evaluate your results in computational linguistics (just as any kind of empirically driven science) against what other people are doing, against what has been found out before ... maybe also to falsify findings that other people have claimed. There is this expression "armchair linguistics" — it's where people just sit on their armchairs and think about stuff, and that makes it really hard to falsify them. I also experienced that doing something on a computer forces you to be very explicit about your theories. Whereas if you're just doing formal semantics, what people do sometimes is to just introduce a context set, and then they say "this is what happens in the context set" and they never explain where it comes from or where they get it from. You can't

do that if you have to implement your stuff — then you really have to be very explicit about your theory. That's a great advantage, and that's what motivated me to enter this field.

I: So basically you think it's a good way of doing linguistics?

S: Yes.

I: Would you describe yourself primarily as a linguist or as something else?

S: I would describe myself as a linguist. I'm currently working on machine translation. That would probably still be linguistics, but also to a certain extent just engineering. But I see myself still as a linguist.

I: So that was about your background and motivation. Then I have a few questions about the field of computational linguistics. Just to begin with, how would you answer the question "what is computational linguistics?" Suppose you had to describe it to a university student who has never heard of this discipline.

S: There are two views on that. One view is that it's parsing the Penn Treebank. The Penn Treebank is one of the standard corpora — it's from the beginning of the 1990s, and that's what people have focused on for almost a decade exclusively ... many results in computational linguistics didn't have much ecological validity obviously, because they were just focusing on this one corpus. That obviously caused a lot of criticism. In the beginning of the 2000s with the advent of the internet and large quantities of user generated data, we also had an explosion of text data. That's when computational linguistics really took off. Also with the help of Google, because Google obviously has a large interest in developing different technologies. To describe it to someone nowadays, I would say that it's basically the study of text mostly ... although now it's also moving into the direction of spoken language as well ... Mark Steedman (a famous computational linguist) has a nice analogy where he says that linguists are usually focusing on very peculiar phenomena that you hardly ever observe, whereas computational linguistics by its very nature (because it's working with data) has to focus on more or less frequently observable phenomena and to draw generalizations from those observations.

I: What are the goals of computational linguistics?

S: I would say in principle it's the same as in linguistics ... there is the linguistics part and there is the engineering part. The linguistics part I would say is really the same: you want to build an understanding of how language works both in the human mind / human brain and also between people (how do we actually communicate)? On the engineering side it may be that you're just interested in machine translation (for example), or if you're Google you just want to get good search results for queries, but you should also keep in mind that those two things are interconnected. You can't do good statistical analysis on large data without having the right technology, and the technology that you're trying to engineer will be pretty poor if you don't put in any linguistic knowledge.

I: So you're saying that there is this scientific goal and this engineering goal, and you're saying this engineering goal requires that you have the proper science to back it up?

S: Yes.

I: But then does this engineering part also feed back into the scientific goal?

S: It definitely does. As I said, if you want to do good statistical analysis, you need the statistical tools. But also in other ways — for example, in semantics you also have this discussion of "what is meaning?" In machine translation we are working with parallel texts, so we have an English sentence and also it's French or German (or whatever) translation. Through working with this kind of textual data, people actually came up with the idea that we don't know what meaning is, but to a certain extent we can represent meaning as the translation of a word into another language, or we can represent meaning just as vectors of neighboring words. Those are things that people would probably not have thought of if they hadn't had the data to actually come up with these ideas. So I think that the engineering part is also fueling the discovery on the scientific side.

I: What are the kinds of problems that researchers try to solve in computational linguistics?

S: All kinds of problems. There are these challenges in different subfields of computational linguistics where people say "I have problem here" and you can

help me solve it, and the winner will get the reputation and maybe also a monetary prize. What I'm saying is that new problems are popping up every year essentially. People are formulating new problems every year, and other problems may become obsolete or maybe be replaced by these new problems. So it's really hard to say what kind of problems—

I: Maybe you could give some examples?

S: OK. Examples would be: Anything that has to do with textual search (so that would be your standard Google example — you type in some key words and then you want to get all the webpages that are somehow related to the keywords ... that not only contain the keywords but that may also contain different descriptions of the keywords); machine translation obviously; there is a field that is currently emerging called computer assisted language learning (basically if you want to learn a second language you want a computer program to help you and to react to your individual needs and modify the lectures accordingly). What else is there? What people are doing now is actually very interesting grounded learning ... they are trying to infer textual description from data of other modalities. For example you want to generate descriptions of YouTube videos automatically — you just want to feed in the video into your computer program and the program should give you a summary of what is going on in the video — stuff like that. Coming back to linguistics, there is also very research oriented stuff like conducting your experiment with humans and then building a statistical model from that. I would also count that into computational linguistics.

I: Typically when would one of these problems be considered to be solved? When you build a program that can achieve what you're trying to do?

S: Never, because none of these programs will ever reach 100% accuracy, it's just that at some point people are satisfied with their results. For example there is this task of stemming (which basically just means that you cut off the inflection from English words), and the standardly used program is the "Porter stemmer" which dates back to the 1980s I think. It's still not perfect, but it's good enough for our purposes. Basically the question at some point becomes "what do you want to focus your efforts on?" If something is good enough so that the errors it introduces won't have too bad a repercussion on systems

built on that system, then you simply say "the problem is not solved, but we can live with this solution."

I: Then I have some questions about how modeling works in computational linguistics? Logicians are often concerned with modeling — in the sense that they try to create formal representations (i.e. models) of some real world phenomenon in order to represent it and to reason about it. Can you describe what role, if any, this notion of modeling plays in computational linguistics?

S: It's extremely important ... I mean, it's the central notion of computational linguistics. I would say there are two major kinds of models. First of all there are statistical models where you take data and you extract so-called features from this data. Feature just means that you look at the specifics of the data (for a word it might be something like how many characters it has, what are its ending characters, whether it contains capitalization, etc.), and then you build statistical models from these features. You basically try to come up with representations of texts that are richer than the text itself, and from that you build statistical models. The second kind of models I would say are ontologies, which come in two flavors: human constructed and nowadays they also try to construct them automatically. A standard example is relational data, where essentially a relation is a verb and then you see which kind of entities you can plug into the arguments of that verb. A classic example nowadays is Wikipedia's relation abstraction — you have something like [Napoleon, place of birth, Corsica] and you want to extract that automatically. Another resource that has been quite popular in the 1990s but is not that much in use anymore is WordNet (there is also GermaNet for German and for many other languages, but WordNet is the original one for English), where people really tried to construct a graph for English. They sat down, collected words and created so-called synsets (sets of synonyms) and then they have relations like homonymy, hypernymy, etc. These things are usually not probabilistic, although the automatically induced knowledge bases actually are. This is the second kind of model, where you basically try to build a knowledge graph over your vocabulary.

I: So in both of these cases, would you agree that the phenomenon you're trying to model is somehow the use of language?

S: In the broadest sense, yes.

I: And in a more specific sense?

S: I would say that the statistical models are often more specialized than ontologies. Ontologies are really broad purpose models that you never use by themselves, but you always put them into some other engine that does whatever task you want it to do. Whereas statistical models are often really geared towards specific tasks and can't be used for others.

I: Specific regions of language use then?

S: Yes, exactly.

I: These models also have some predictive function?

S: Yes. Statistical models anyway, because the entire goal of machine learning is to predict future data. Ontologies — they don't really have predictive power, but they do allow you to discover relations between words that you might not have thought of before, because you can basically traverse the graph and see that "Hey! These two words are connected by hypernym two levels up" and this is something that you might not have thought about before. So you can still use it for discovery, but whether it has that much predictive power I'm not sure.

I: Both these types of models — are they of equal importance?

S: No, no. Statistical models are by far the dominant ones.

I: Then do you mind if for the rest of this interview we focus the discussion on the statistical models (unless you feel there is some reason not to). Because otherwise for every question you would have to say "this is the answer for statistical models, and this is the answer for ontologies."

S: Yes, sure.

I: What are the predictions supposed to be about? They're about language use or about particular regions of language use?

S: Say you set yourself a task ... let me stick with machine translation for example. You train your model on a particular set of data, and your goal is

obviously to translate sentences that you haven't seen in your training data before. So the prediction that you're trying to make is: given a new input sentence, you're trying to predict a translation of it. Another example is textual retrieval, where you train on a particular kind of ... actually there's another nice application which is very popular nowadays which is sentiment analysis — for example, users give comments about a product on Amazon, and you're trying to automatically classify whether this is a good product or not based on this user information.

I: How do you actually check that the model that you've built is a good model for the thing you're trying to model?

S: There are different ways of doing that. Ideally you want to use human evaluation. So you have a test set, you make predictions on the test set, and then you want humans to check how well your predictions match what they would have thought. But obviously that's very expensive and also very time consuming, so that's usually what you do at specific steps of your modeling process — mostly at the end when you want to check if your model is really working well, or when you want to compare it to other models. For intermediate steps you usually try to create so-called metrics, which are basically just automatic evaluation templates that allow you to ... usually they don't allow you to say how good your model is overall, but they allow you to compare two models against each other ... so you can say model A is better than model B. That still doesn't mean that model A is a good model, but it allows you to make relative comparisons.

I: What are the standards for success and failure in modeling? Obviously one thing you've said is that we use these checks, and if it fails one of these checks then there's something wrong with the model. But is this the only criteria for success and failure in modeling?

S: So there was this statistician George Box and he said "all models are wrong, but some are still useful." So, success would probably be defined as a useful model, and success really has to do with what point of time you're proposing your model at. Models that are totally obsolete nowadays were big breakthroughs ten years ago. It's an incremental process where you keep improving your models and you keep building on ideas that other people have proposed.

In that sense it's hard to quantify success scientifically. On the scientific side it really just means that you're outperforming all currently existing models. Success on the commercial side ...

I: But what does it mean for one model to outperform another model?

S: Either on a generally acknowledged metric (where people think that this metric really makes sense) — it's relatively better on this metric than all other models. Or, again, through human evaluation.

I: OK.

S: If you're a business, then obviously your standard of success would be ... if you're using the software yourself, then do you get the results that you want to a satisfying extent? If you're a service provider, then is your client happy?

I: In terms of the human evaluation, what you're trying to test is whether the predictions of your model agree with human intuitions?

S: There are two ways of testing. You can either have your model produce a prediction, and you can have your humans produce a prediction, and then you just assume that the humans are always right, and you measure the overlap between the model's prediction and the human prediction. Or you can explicitly ask humans "is this prediction sensible given the input and given the task?"

I: Do you think that computational linguistics is a rigorously scientific discipline?

S: As rigorously scientific as any empirical science. You'll never have 100% certainty, but it is rigorous in that it is a very vivid community with lots of people working on the same topics, lots of people producing similar systems ... if you want to publish a paper in a very well acknowledged journal or conference you have to compare your system to other systems, and there are usually also standard data sets. Your model can improve by just using better data or data that's more suitable to the task, so there are also some standard data sets that everybody is using. So in that sense I would say that it's rigorously scientific, because you compare your approach to other people's approaches on the same data. That's what you do standardly. Obviously, as I said, new problems are

being invented every year, and then mostly the people who spotted the problem first are also the ones to propose a data set.

I: Do you think that computational linguistics as a discipline has been successful in achieving its goals?

S: I think it has actually been getting increasingly successful. As I said, with the advent of the internet on a large scale it's actually becoming more and more relevant, because all kinds of information that we have currently stored is mostly stored in textual form, so I think it hasn't reached its ceiling yet. It's in fact getting better and better, and more and more relevant, because of the internet mostly. It's also moving towards spoken language now — you have these applications like Siri on the iPhone which are pretty impressive to most people, so I think it's a success story that has only just begun.

I: But I mean ... the question was if you think computational linguistics has been successful in achieving its goals, and when I asked you about the goals you mentioned two distinct goals — one scientific and one engineering. But in the answer about the successes you've made reference mostly to the engineering part of it.

S: I think scientifically you can never ... there is no ultimate goal in science. You just keep progressing, progressing, progressing, and you should always be aware of the fact that you're always wrong — there is no perfect model — but you're just trying to get less and less wrong. I think in that way, it is achieving success, but I don't think that there are any goals to achieve in any science.

I: Those were the questions that I had about computational linguistics as a field. Then I want to ask you some questions about how it compares to formal semantics. What do you think are the similarities and differences between computational linguistics and formal semantics? Maybe we can start with the goals. What do you think are the similarities and differences in the goals of these two disciplines?

S: That's hard to say, because officially the goal of formal semantics is also to explain language and how language works in communication. In terms of these abstract goals, they should be pretty much equivalent. I think it's mostly a difference in how they are trying to achieve these goals.

I: So, in terms of goals you think they share similar goals?

S: Except maybe for the fact that computational linguistics is just broader, because it captures all branches of linguistics and not just semantics, but yes.

I: How about in terms of the actual methodologies of these two disciplines? What do you think are the similarities and differences between them?

S: As I already said, the predominant models in computational linguistics are statistical models. Formal semantics is mostly centered around logic. I already mentioned evaluation. There are clear evaluation standards in computational linguistics. Those standards just don't exist in formal semantics. Usually there are no standard data sets that everybody agrees upon. So, basically every author is picking his own example sentences. Sometimes they are the same, but mostly it's just random sentences that the authors themselves made up.

I: So far you've said that one difference is that you rely on statistical models whereas they rely on logical models, and another difference is in terms of data — you have concrete criteria for evaluation and concretely accepted data sets, whereas over there there is not much agreement on the criteria for evaluation—

S: It's also not that much of a worry [in formal semantics]. In computational linguistics, there are people who are really investing their whole research time into developing metrics for evaluation. Such an effort is not even being made in formal semantics, so I think that to a large extent people are not aware of the problem of not having standard evaluation metrics.

I: The way you're describing it — it's as if both disciplines have almost the same methodology, but that one discipline is more successful in executing that methodology?

S: Why do you think that it's the same methodology?

I: Because they are both in some sense trying to build a model and then use this model to predict things about language use. So in that sense I guess they are following a similar methodology, but what you're saying is that the formal semanticists are not executing this methodology in a rigorous way.

S: Well, at least they're not evaluating it in a rigorous way. Their implementation is arguably more or less ... there are also people who say that formal

semantics is not really formal, because they sneak in these pragmatic things that are not really implemented in their formal system ... I guess implementation wise they are definitely closer to computational linguistics than evaluation wise.

I: OK. What do you think are the similarities and differences in terms of the kinds of explanations that computational linguistics offers versus formal semantics?

S: I think that again has to do with the data. Explanations in formal semantics are mostly based on abstract concepts like predicates, frames, models etc. This probably has to do with the fact that they are not using that much data. Whereas, as I already explained, in statistical modeling you are trying to extract features from the data. What explanatory factors you use depends on your own ingenuity in designing features, but also on the data. So I think in that sense the explanations that computational linguistics potentially has to offer have more ecological validity, because they are basically derived from the data itself.

I: Do you view these two disciplines as sort of competing ways of analyzing the same phenomenon?

S: Yes. Mostly people just view them as completely distinct and not having anything to do with each other, but actually my advisor (he is a psycholinguist, so he would talk about semantics and psycholinguistics, but he was also doing computational psycholinguistics, so he would build statistical models of the experimental data he got) would always say that we are in competition for explanatory power. Just because we are using different formalism — on the one side logic and on the other side statistics — doesn't mean we are doing different things. We are doing the same thing and we are in competition for explanatory power. It's basically up to your standard of evaluation to say that one is better than the other, but they are in competition definitely.

I: But you don't think that there is room for collaboration?

S: Oh, yes there is, and there should be way more collaboration. People in computational linguistics recently got interested in formal semantics again. If you have formal semantics, or if you use any kind of logical (or logic like) formalization, then it's much easier to connect databases to textual data. The

problem however is that those formalisms need to be much more flexible than a rigid logical system. So you kind of have to be on the middle ground between pure logic and—

[[Interruption]]

I: We were talking about whether formal semantics and computational linguistics are competing disciplines, and then we moved on to talking about possible collaborations between the two.

S: Yes, I think that collaboration in any field should be more and more enforced. There currently is some interest in formal representations of textual meaning or language meaning in general, because it makes it easier for you to interact with databases or with any knowledge source that stores relational data. People in computational linguistics are kind of looking for semantic representations as well, but those have to be way more flexible than what you have in formal semantics. You are looking for a middle ground between rigid formal systems and systems that are flexible enough to capture textual data ... if you haven't seen a word before, you still want to be able to capture it somehow. I do think that both disciplines could benefit from each other, but it requires some open mindedness and some will to change one's perspective from both sides.

I: OK. Do you think that formal semantics is a rigorously scientific discipline?

S: No.

I: Why or why not?

S: Well, we talked about this whole evaluation stuff. Rigorous methodology basically establishes itself through rigid evaluation, and as long as that's not taking place you can't really say that you're working in a scientific way, because you're not making your results comparable to others.

I: Do you think that formal semantics has been successful in achieving its goals?

S: Well, if you stick with seeing the goals of formal semantics as being the same as computational linguistics, namely explaining how language works and how linguistic communication works, then no, definitely not. Because they have hardly offered any conclusive results that they could also verify.

I: Then I have here with me ... so I've been conducting these interviews with formal semanticists over the past few weeks, and in my interviews with them I asked them also to comment on the relation between formal semantics and computational linguistics. I've collected some salient points from their comments, and I would like to present them to you to see if you have any comments on that.

S: Sure.

I: First of all, one of the points that often comes up is that formal semanticists often say that these stochastic models favored by computational linguists, although they have greater coverage of data and greater predictive power, that they do not contribute to a deeper understanding of natural language, and they don't have much explanatory power. In the words of one formal semanticist "computational models just don't understand anything."

S: So neither do logical models. A model never understands anything. A model just gives you correlations between factors in your data, and you have to interpret the model or the output of the model. The models by definition don't understand anything, and the claim that stochastic models haven't contributed to our understanding of natural language is just plain false, because basically all modern research that is being done in psycholinguistics or any kind of behavioral linguistics where you examine people, where you conduct experiments, where you gather data ... is done in some sort of statistical evaluation, and statistical evaluation includes building stochastic models. So this claim is plain wrong.

I: But I think part of the motivation behind this claim might be the feeling that in stochastic processes and models you are only explaining the surface data, whereas the formal semanticists are looking for insight into the deeper structural features of language, and they feel that somehow their approach gives them more insight into these deeper structural aspects of meaning, which is something that they don't see in the stochastic approach.

S: There are two things you can do. First of all really just look at the surface, at the data you have, and that's what I would call the real scientific approach, because you can only evaluate what you see. You can only evaluate observations. This is what is being done in psycholinguistics, etc. In computational

linguistics proper, we also have for example parsing models, which try to induce latent structures over sentences or over phrases. But the thing is that this doesn't mean that these models are right or that the latent structure you assume is right — it just means that it hasn't so far been falsified by your data. So, I think this is the main difference between these two approaches. You can always put forth theories about latent structures, but you can only falsify them — you can never absolutely verify them. Whereas if you only look at the surface data, the data you've gathered as such, then you can actually verify something.

I: Do you think that this points to an important difference between the goals / motivations / aspirations of these two disciplines — that what they are aiming for in a certain sense might be different? In formal semantics, they really are aiming to get a deeper insight into these structural features of meaning, whereas this is not something that computational linguists are aiming for even?

S: Computational linguists *are* definitely aiming for it. There is semantic role labeling, there is semantic labeling, etc. — all of these things exist. I'm just saying that it's only theories. It's only proposals. We shouldn't think that this is something that really exists in natural language. It's just a way of conceptualizing natural language, and this way may (given a certain task) be appropriate or not. But latent structures never *exist* per se — this is something we should always keep in mind. Whereas our observations — they really exist, because we have gathered them from the real world.

I: OK. So I think the debate is really not about the existence or non-existence of these latent structures. It's just "do they have explanatory power?" and "do they have insight?" I think in formal semantics the researchers feel that on both counts yes, they offer both insight as well as explanatory power.

S: Well, then I can only say that most of the insight that's currently gained and most of the explanatory power is given by stochastic models. That also has to do with data coverage, because you can't claim that you have high explanatory power when you can't account for most of the data. Data coverage and explanatory power to some extent go hand in hand.

I: OK. Another point that's often made is that computational models have some sort of threshold when it comes to accuracy and data coverage. They

may get to, say, 90% or 95% accuracy, but there is the idea that they can never get the last 10% or the last 5%, and that some kind of formal, structural postulates are needed to get this full accuracy.

S: So let's put it like this: there's no single model in the entire field of natural language investigation that would ever have reached 100% accuracy. Such a model just doesn't exist — neither in formal semantics, nor in stochastic approaches nor anywhere else. So, it's just a claim. Theoretically, maybe even stochastic models could reach 100% accuracy. Theoretically maybe logical models could as well, but nobody has shown such models to exist, so it's a mere stipulation. Actually modern stochastic models that have been developed over the course of say 20 years (say, for example in parsing) reach accuracies of more than 95%. They'll probably never get to 100%, but I think this point is valid for logical models as well as for stochastic models — that up until now there just is no model that can reach 100%, and we don't even know whether such a model can ever be built or not.

I: But it doesn't somehow invalidate the stochastic approach?

S: It doesn't at all. Because as I said, we're in competition for explanatory power, and if someone comes up with a model that delivers more explanatory power — be it a stochastic model or a logical model — then this model is better. As of now, there are no better models than stochastic models. Maybe in ten years this is going to be different, but as of now this just doesn't exist.

I: Another thing [that is said] is that computational models are only applicable in very specialized contexts, and they don't have the kind of general applicability that formal models seek to have.

S: Again, that's just a claim. To my knowledge there are more problems nowadays that are being attacked with stochastic models than with logical models. And you have the great freedom in stochastic modeling that all you need really is probability theory, and you build on top of probability theory. Whereas in logical modeling, you often have a sense that people come up with entirely new calculi just to tackle one specific problem. I don't know whether you've heard about donkey anaphora and discourse representation theory — discourse representation theory was really only developed to deal with the specific case of

donkey anaphora, which in stochastic modeling just doesn't happen. Obviously we devise specific model structures for specific problems, but the underlying formalism is always probability theory, and nothing is changing there. From that point of view, I would even claim that stochastic models are more general. Again, that's just a claim, but the fact is that nowadays more problems are being solved with stochastic models than logical models.

I: Researchers in formal semantics also claim that their research is of such a nature that in principle it can be used in building computer systems that can understand and process natural language, although that's not explicitly what they're doing. Do you think that this is a valid assessment of what they're doing?

S: I guess that in principle you could use just about anything to build a computer system. It's just that somebody has to do it. Up until now nobody has done it, so I don't know whether it's feasible at all. I mean, the question is always that you can implement a solution to a problem in many different ways, but you're looking for the most general and most efficient implementation. Currently that's always done with probability theory, and I don't think there's going to be much of a change in coming years. In principle, sure, everything is possible, but as a matter of fact that's not what's happening.

I: But you don't see these semantic theories making contributions towards the engineering dimension?

S: Oh, they have, they have. For example, there is categorial grammar which has been developed by Mark Steedman who is a linguist and a computational linguist at the same time. Categorial grammar is basically a formalism that ties syntax and semantics together, and he has had some nice results with that. So, I'm not saying that there is no contribution from formal semantics at all. I'm just saying that when you're looking for a solution, you should probably look for the simplest solution possible, and oftentimes that solution is not formal semantics (although it possibly could be).

I: Another claim is that there is already in computational linguistics a trend towards importing insights from formal semantics. One example cited was distributional semantics. Is this actually a trend in computational linguistics?

S: That's definitely a trend. Although distributional methods have been around in computational linguistics since the 1980s essentially — so way before distributional semantics was around. Distributional semantics is really just an extension of the distributional method (that has been around since the beginning days of computational linguistics) towards semantics. I wouldn't say that this is inherently semantic ... actually it was invented in the context of computational linguistics and computer science, so formal semanticists actually had no say in developing this particular approach.

I: Another thing is that formal semantics sees itself as being a longer term process in a certain sense, but in the end it may be able to provide insights into language that computational methods cannot provide. Do you think that there are certain insights that are inaccessible to computational methods which formal methods are more suited towards?

S: That sounds pretty religious to me. So, there are people who are waiting for the final salvation which will give them the ultimate insight into language at some point in the distant future. As I said, all of these things like "logical methods can achieve more than stochastic methods" or the other way around, they are just claims. Claims that can either be proven or disproven ... but the thing is that whenever people claim something like this (that in the distant future there might be something that is inaccessible to stochastic methods) they also never give concrete examples, which basically means that they themselves can't think of any examples and are just making some random claim. It's hard to argue with someone who is just claiming something out of the blue.

I: My last question to you is: what (if anything) could potentially cause you to give up your commitment to computational linguistics?

S: I'm not actually committed to computational linguistics. I am committed to stringent evaluation, and this mostly involves statistical evaluation. Because statistics is extremely complex (especially on large data sets) you just need computers, because you just can't do it by hand. There is too much data. If anybody ever comes up with a formalism that is better than probability theory, that is better than statistics, I would be happy to adopt that formalism. But I think that no matter what formalism you use, since we have so much data to

deal with, you will always have to resort to computational methods, because you can't just do it by hand.

I: Sorry, maybe the question was wrongly posed at the level of computation, but I guess more broadly: what (if anything) could cause you to give up your commitment to statistical linguistics?

S: Oh. Well, as I said, if there is any better evaluation framework ... anything that yields more insight than statistical methods, I would be happy to adopt that framework.

I: And how would we determine whether a putative method is in fact better or not?

S: That's a very hard question, because I can't think of ay better method. Basically statistics is the go-to formalism that you use whenever you have any quantitative data. The problem is that even if you measure correlation with human judgments, you measure this correlation using statistical tools. Even if your original system was a logical system, and then you get predictions from your logical system and from human judgment, you will still use statistics to measure their correlation. What would be the criterion to replace statistics? I don't know really, because it seems that any evaluation criterion is again based on statistics. You might come up with methods that work better than stochastic models, but you would assess their effect through statistics. You would compare them to stochastic models and see whether there is any improvement in accounting for the data or not.

I: Are there maybe some key presuppositions that your discipline is committed to?

S: As such I would say no. There is one cool thing about statistical modeling: any statistical model makes assumptions, has presuppositions in the way the model is built—

I: Not the presuppositions of any particular model, but the presuppositions of the discipline as a whole. For example, in formal semantics one of the key presuppositions is that formal logic is operative in the way in which human beings make use of natural language. There would be no comparable presupposition for your field?

S: There is nothing like that. As I said, any model that you build has its own presuppositions, but if you resort to the simplest model possible then you're really just looking at the data and measuring correlation between factors. That's it. So there's no inherent presupposition.

I: OK. Those are all the questions that I wanted to ask you. Is there anything else that you would like to say before we end this interview?

S: Not much. No.

I: Alright. Thank you for your time!

C.4 Interview with a computational linguist

I: Let's begin with a brief introduction from yourself. Could you please describe your academic and professional background, and how you got into computational linguistics?

S: Yes. I started out studying physics, physical engineering actually, at the University of Eindhoven. I was during the time already interested in many different things, including (especially) perception ... auditory perception was in some sense my focus in the last few years of my studies. I spent a lot of time at the Institute for Perception Research in Eindhoven, so I did work on pitch perception and things like that. I did a lot of signal processing. Also computer science things — I did formal languages and automata theory as part of the computer science. That was sort of my background when I finished [my studies]. Then I got a job as an engineer at the research division of the Philips computer factory. So I got in there on the basis of my physics degree. They like physical engineers there — that's the kind of people they hire. They were starting a project on artificial intelligence and there were ideas about perhaps doing something in the direction of natural language question answering, natural language interfaces for database systems, etc. There was a small group of people (all basically physicists and engineers) and we started this project on natural language database interfaces. So that's how I started, in some sense out of the blue, with a little bit of background in formal language theory, a little bit of background in logic (I had taken one logic course — predicate calculus — at the University), and no [background in] linguistics. That's how I started with that project. So I started to get up to date about how to approach this issue about natural language database interfaces, so I studied Chomskyan linguistics and versions of that, and formal semantics type things, which were at the time practiced in computational linguistics already. There were a few attempts at natural language question answering, which did something like logical semantics, but in a very sloppy way in very messy systems, but things in that direction existed. I just took some courses in summer schools when I first started ... the most influential course that I went to was about systems for lambda calculus. That seemed very useful ... that was independently of Montague (I hadn't heard of Montague at that time). We set out to build a

natural language database interface which was based on the idea of translating natural language sentences to logical formulas, and then translating the logical formulas into different logical formulas, so that in the end we could evaluate the logical formulas with respect to a particular given database. That's the project that we started out doing in the early 1970s and finishing in the late 1970s. We worked more or less from scratch in trying to create rules for mapping these natural language sentences to logical formulas. So we invented many things from scratch, and while we were doing that I came across Montague grammar and found out that there were people in Amsterdam working on that, so I started going to the colloquium of the group in Amsterdam (the people that you already interviewed: Martin, Jeroen, etc.). That's how the connection with that group started.

I: Then how did you move into the more computational side of things?

S: This was computational from the beginning. We actually built a system where you could type in a question, and then the system would translate the question into a logical formula and do all kinds of things, and then look at the database and come up with an answer. This was about a database with marketing information, so you could just type in "how many Philips computers did Shell buy in 1958?" and then it would give you an answer.

I: What would you say was your primary motivation for doing computational linguistics?

S: Well, my motivation was in some sense related to cognitive science. I had an interest in seeing whether it would be possible to make mathematical models of language in that way, with an emphasis on the interpretation of language, and to see whether one could build such models and implement them on a computer. The practical application was not so important in my own mind, but that was a nice excuse to be able to do it.

I: Would you describe yourself primarily as a linguist or as something else?

S: Like I said, primarily I was a physicist ... but in the way I look at language (compared to certain logicians) I would say yes, I would be more of a linguist in the sense of being really strongly empirically motivated about what the reality of language is, independently of preconceptions that one might have about it

or independent of logical ideals that one might have about it. But then I would be a funny kind of linguist, because I would operate completely independently of linguistic traditions, and also not in terms of having preconceptions about linguistic universals or what languages have in common, or things like that. So it's very strongly empirical. We say here we have this language, and we know these people use these sentences ... they mean something by that and they use that to convey information ... so we build a system that does that. The computational aspect puts a strong emphasis on making everything completely formal right from the beginning. We are really interested in how to make that formal so that we can put it on a computer.

I: So that was about your background and motivation, and then I have a few questions about computational linguistics as a discipline. To begin with, how would you answer the question: "what is computational linguistics?" Suppose you had to explain it to a university student who has never heard of the field.

S: Well, I would say it's a very broad area of activities. It's not necessarily one field in a very strong sense of that word. It's a broad field, and it includes many subfields. There can be an emphasis on putting syntactic theories on the computer, say writing parsing systems based on certain assumptions of what grammar looks like. It can be motivated by the desire for practical applications like natural language database systems, or machine translation, or information retrieval. It can be a cognitive science type motivation, that you try to build little models of certain processes of sentence processing and things like that. So it's a very diverse field which is defined say as the computational dimension of linguistic theory, or defined by computational applications which involve language, or as a computational version of cognitive science theories.

I: What are some of the kinds of problems that researchers are trying to solve in this field?

S: That again is quite diverse. It can overlap with problems that people try to solve in theoretical formal semantics or in theoretical syntax, but there can also be specific computational problems. For instance, if you do statistical parsing — you have sentence analysis programs which are based on statistical information and there are specific problems about how to use really large amounts of information and how to make it efficient; how to do the right kind of statistics

which gives you the best kinds of results ... that would be a very specific computational problem. So there are many specific problems.

I: Broadly is there a unifying set of goals for the discipline? Earlier you mentioned that one of the goals is to make connections with what we know about cognition and cognitive models, and you said some of the goals have to do with application (like natural language databases, etc.) ... so is it just that there are different goals for different areas, or do you see any unifying set of goals for the discipline as such?

S: I really see different goals for different sub-areas. I don't see a big unification. In some sense the field is quite fragmented. There are very specific research traditions about statistical parsing or about machine translation, and they are separate. You could think that of course machine translation is inherently related to statistical parsing, because how are you going to translate a sentence if you don't even know how to parse it? If you really want to do it right, parsing has to be part of the machine translation. If you want to do the parsing right, you have to bring in the statistical information ... so you could easily think of putting it all together, but in practice these fields get quite specialized. If there are certain machine translation methods which are useful, which are successful and don't use parsing, but if they get better results, they determine the nature of the competition in machine translation. So what you get currently is some really different research traditions which are strongly competitive in terms of reaching certain specific quantifiable goals with whatever methods work.

I: Then in terms of the methods of computational linguistics, I'm curious specifically about how modeling works there. Logicians often employ modeling, in the sense that they try to create formal representations of some real world phenomena in order to represent it and to reason about it. Can you describe what role, if any, this notion of modeling plays in computational linguistics?

S: Well ... if you take the old natural language database interface systems for instance ... the way we worked on those is strongly related to this modeling notion, because there you have a situation where somebody has a database — it's structured in a certain way, it represents certain information and one can talk about that information in certain ways in English, and how are you going to bridge that? One way to think about that is to think about what is the

domain that is represented in this database and abstract away from certain details of how it's represented and to think about what this database is really about ... to devise a logical language where you say: These are the core notions. It really involves these domains of entities and these relations between them. This is what is really represented in that database. And you're going to make an elegant model of that, and then you have a language which is a nice interface between English on the one hand and the database on the other. You have an idea about the kind of language that you want to map your English to and the kind of rules for that. That's like the old fashioned AI, where the notion of modeling micro-domains was very important. If you build an expert system about a particular domain, you really think about "what's going on in this domain?", "what are the relevant notions?", "what are the relevant logical relations between them?", etc.

I: And the domain is always some aspect of natural language?

S: Well, in this case the domain was computers and who buys them and who uses them and who sells them. That's what I mean by the domain. Then there's another aspect of modeling, namely how language works, which is very similar to what you've seen in formal semantics. We try to think of what's the machinery ... there's a reason why we can put these words together so that they together end up meaning what they mean ... so what's the machinery behind that? That is in some sense a semantic interpretation of the syntactic operations with which we put words together. That's how you do formal semantics. The way I worked on those problems, that's strongly and directly related to these modeling notions you talk about. In a lot of modern computational linguistics that's in fact different. People work on statistical parsing or on machine translation as if it's purely a matter of statistics. They sort of skip the modeling. So that's become a somewhat different kind of field. That part of computational linguistics is in fact not very relevant for formal semantics.

I: In that area of computational linguistics, this notion of modeling is not in play you think?

S: No.

I: But in a certain sense they are also trying to build representations of how natural language works and make predictions about this, etc. Could one not

argue that they are also in a certain sense trying to model natural language? They are making stochastic models.

S: Yes. If you insist on that, of course the answer is yes. I just want to emphasize that the nature of these models is very different, and there is a tendency to make the models very shallow and to have an assumption that the statistics will do all the work. That's a possibility which emerges as soon as you embrace statistics. If you have enough data, you can always be hopeful about what the statistics will account for without needing any explicit modeling, and that's been very much a tendency in today's computational linguistics.

I: What exactly is the alternative to that?

S: It's more like old fashioned artificial intelligence and old fashioned formal semantics. Where you explicitly build systems where you say these are the rules ... you try to craft the rules so that they give the right results ... you really handcraft them. Then you can still add statistics to that, but the statistics would not carry so much weight.

I: So in either of these two cases, the models that we build, how do we check that it's actually a good model for the thing that we're trying to model?

S: Say, for the natural language database interface systems that I used to work on several decades ago ... there the check would be whether somebody could actually use the system so that the language which is modeled in the system is some meaningful subset of language that people actually use, so that people can actually work with it. That's sort of the practical way of evaluating it. In modern computational linguistics, people very often have very narrowly defined goals. Say in statistical parsing you have a given corpus and you have a given training set with annotated sentences, and you have a test set which you're not allowed to look at. Then you have a parsing system which you can train on the training set, and then the evaluation is percentage accuracy on the test set, where you try to get the right results for the test set. It's very formal, quantitative and empirical.

I: So this is the basic criterion for success and failure in modeling. If it makes bad predictions with respect to the test set, then it's not a good model.

S: Yes, that's how people deal with it.

I: Is that the only criterion? Or do we have any other criteria also? Maybe the model should agree with some other things that we know?

S: The thing is people tend to focus very much on just the performance on the test set. It's very difficult to publish a paper which contains new ideas if it doesn't give better performance on the standard test sets. I don't think that that's a very good situation, but that is the situation.

I: Why do you think that that's not such a good situation?

S: Well, because it sort of treats the test set as a goal in and of itself, as if there are no other intellectual goals anymore. The questions of whether something makes sense from a cognitive point of view, or a linguistic point of view, or even from a practical point of view — those questions aren't on the table anymore if you focus completely narrowly on the test set behavior.

I: Do you think that computational linguistics is a rigorously scientific discipline?

S: No ... because ... certain versions of it are more practical where the goal isn't even scientific. You just try to make something that works, and if it works it works. And then there are things which are not so practical, like some of these contests where people try to get good accuracy on the test set. That has the looks of objective science, but in some sense the parameters of the situation are completely arbitrarily defined. You try to get good accuracy in assigning trees on the basis of the training set to the test set. Something comes out, but what does it mean? The truth is that it's extremely non-scientific in the sense that it's completely not clear how it generalizes to what this means about anything beyond what you measure. So it becomes very narrow in the way in which a lot of experimental psychology is very narrow. You measure something, but because of the lack of a theory, and because the experiment was not part of a bigger theoretical enterprise, you can measure these things (if you put people in the lab and if you give them these stimuli they do that, and then you measure that they do that), but in some sense it's completely unscientific in the sense that it completely lacks the dimension of generalization. You just measure some things, and you may measure them very carefully, but then all you know is this very specific thing. That's a little bit going on with today's

statistical computational linguistics. We have these techniques, and they work in these specific test situations, but we don't really know what that means.

I: Do you think that computational linguistics has been successful in achieving its goals?

S: No. Perhaps it could have been ... like the classical natural language database interface systems — they work to a certain extent, so that could have been successful if that would have been an ongoing line of research. But that line of research hasn't gone on, because it is a lot of work to model particular domains in a very careful way, and then all you get is the capability of talking about only that domain. It's a lot of work with limited results. What has in fact been successful is information retrieval (like Google and systems like that). That's all based on techniques from the 1960s, and that has been very successful. It's very limited in what it can do, but for practical purposes it's very useful. Machine translation is still pretty bad — there, the old approaches of handcrafted grammars have been taken over now by statistical methods, but what one can do with such completely automatic methods is extremely limited. So all the really interesting challenges have not been answered yet. Although one could imagine that it could be done, but it still would require a lot of work.

I: Then lastly I have some questions on the relation between computational linguistics and formal semantics. First of all, what do you think are the similarities and differences between computational linguistics and formal semantics in terms of their goals and motivations?

S: Then it depends what part of computational linguistics you talk about of course, because it's very broad. Some parts have motivations which have nothing to do with formal semantics, because they're purely practical or they're only about syntax ... I think for the problems I worked on ... if you still want to work on natural language interface systems for instance ... then the goals are very similar ... the way I look at it, from a computational point of view it's very good to work with formal mathematical methods and then it becomes basically one field, the way I look at it.

I: But if you look at the people who are working on the more statistical side of things, then you think they are not very related in terms of their goals and motivations?

S: No. Then I think it's completely different. Although this is a shame to a certain extent, because one could imagine, say, if you look at the work in formal semantics, you can certainly see that people tend to work with truth-conditions where everything is very strictly defined ... you could imagine that perhaps it would be nice to be able to put in a statistical component of some sort in a formal semantic system. One could see that that might make sense, so one could certainly imagine that there would be interactions between these fields, but I don't see much of that happening. In fact they are very different. They are very very different enterprises.

I: The way I would have understood it is that they seem to share similar goals, but they have become very different in terms of methodology and approach. But you are saying that they don't even share similar goals?

S: Yes. Because, say in machine translation, one might imagine that some kind of analysis where you get the meaning right would be an important part of what you're trying to do. But that's not how they do it. They just try to map from one language to the other, and then see whether they get away with it. So if you look at how the systems work, you cannot point to this component where you say "OK, now we've decided what the meaning is, and we can see whether that's right or not." That component isn't there. Or if you look at information retrieval ... you could imagine that you have a query and you try to have a more refined notion of getting information. You have non-trivial queries where you could say that I'm first going to look at whether I understood your query correctly. But nobody is working on that. Nobody is working on a notion of queries where they entertain the notion of whether they got it correct or not. Everybody is just looking at the query as a bunch of words and they try to get the information, and then hope they get it right, and then decide based on percentage accuracy whether an approach is better than another. People go completely only for the practical end results and don't think of the systems in terms of particular components where the notion of meaning plays any role for instance ... they just drop that.

I: Would you say that these two are competing disciplines in some sense?

S: In terms of practical applications, certainly [they are competing], with the statistical ones winning, because it's much less work. Intellectually perhaps

not. Certainly to the extent that we do logic and formal semantics with a philosophical background, those goals are still valid, and they are not undermined by whatever people do with statistics. That's valid, and I think that work goes on. Then perhaps the most interesting question is whether they are competing approaches if we talk about possible future theories of cognition. There it seems obvious that we need some kind of synthesis ... that we do need both approaches, but in a way where they are not competitive, but where they are integrated. That is what we are going to need for cognition. There also I think it's not so clear. We may hope that in some future version of cognitive science it will actually all come together. We can still hope that.

I: But as it stands today, do you think that perhaps these are two different scientific paradigms?

S: Yes. Certainly they do fit the definition of what one calls a paradigm. They're just different ways of operating and different ways of looking at the problem to begin with.

I: Can you elaborate on this a little bit? What do you think is the difference between the way in which the formal semanticist and the computational linguist look at problems?

S: Well, it ties in very much with what you were asking before about modeling. Whether you are really inclined to make a formal model of something, versus, that you can think you can just gather a lot of superficial data and do statistics. If you take the extreme versions of these two approaches, they are extremely different. In that sense they are different paradigms, which doesn't mean that they can't be combined, but they are really different sets of assumptions and they are different mindsets in terms of which one can work. One can see that difference. It's also the case with statistics ... statistics is a very live branch of mathematics which is very active also. People are getting increasingly clever about increasingly complex statistical things and about their computational implementation — about how to do really complex statistical computations efficiently. There is science there, but the science that's happening there really is about the statistics, and it's not about the content matter.

I: OK. So we have these two seemingly competing paradigms, and you have expressed dissatisfaction with the statistical computational approach, and you

don't think it's rigorously scientific and that it hasn't been successful in achieving its goals. But what about the formal semantics paradigm? Do you think that paradigm has been successful in achieving its goals?

S: Like I said, I think it could have been ... in terms of interface systems, it could have happened, but it hasn't. Just as a matter of practical contingencies, it hasn't happened. That's on the practical side. Then I think formal semantics as such (in the logical, philosophical, linguistic context) has been successful in helping develop ideas about what's going on with language. We also run into the limitations of that. We see increasingly in papers in formal semantics, that people talk about judgments involving pragmatic dimensions of the interpretations where it becomes increasingly difficult to really follow what the predictions of particular theories are. Too much hinges on assumptions of aspects of the pragmatics that we haven't formalized. So we run into the limits of that ... not all aspects of language can be understood so well with these techniques. In some sense, I would also say that's also a success, but that really raises the question: where do we go from here? If we really want to understand things very well, we have to get more formal theories about pragmatics, about real world contexts, integrate statistical stuff, etc. So clearly that approach has its limitations, and that's becoming clear.

I: Do you think that formal semantics is a rigorously scientific discipline?

S: Well, I would in any case drop the word rigorous. Because that would be emphasizing very strong objective qualities that it [formal semantics] doesn't quite have. If you don't put "rigorously" and you just say "scientific" one could probably say yes. It just is the case that, say in physics, it's easier to agree on what the evidence is and it's easier to separate the data from the theory, so it's easier to construct objectivity and to show that it has objective validity. In formal semantics it is more a matter of having to make certain assumptions, and in terms of those assumptions you make your observations, etc. It is much harder to really make a point for the objective validity of what one claims, and in that sense it's just more difficult than physics. It requires more good taste. You have to really know what intuitions to take more seriously than other ones, and what counterexamples to take seriously and what counterexamples to dismiss. If you don't know how to do that, you can't do formal semantics, and for somebody who doesn't understand such things, you can't explain to

them why it's objective in any sense. But I think among people who have the right mindset and who are willing to get into it, it has a certain degree of objectivity, but only a certain degree. So it isn't physics, but it also isn't completely unscientific.

I: Recently there have been some internal criticisms of the discipline made by Martin Stokhof and Michiel van Lambalgen. I'm wondering if you're familiar with this criticism?

S: I have read things by them, but perhaps you should summarize it for me ...

I: I think their main argument is that it deals too much with idealizations of the natural phenomenon that it's trying to model, and therefore it becomes very difficult to verify or falsify any of the claims made by formal semantics, and therefore its scientific status is rather problematic.

S: Yes, that ties in with what I was just saying, and it emphasizes very much this difficulty. Yes, it emphasizes that we can't really pin it down ... we can't *really* show what the objectivity is. So if they point that out, that is correct. Given certain idealizations, if we embrace those idealizations, we end up agreeing on a certain set of phenomena. The real question is how meaningful is that for our understanding of language as it actually exists? Is this only an artifact? Are we only inventing something which has some remote connection with language, but is really a branch of mathematics (just uglier than ordinary mathematics) and it has nothing to do with actual language? Or is it a meaningful idealization? Is it meaningful as a perspective on language at large, including all the phenomena that are not covered? There one can have different feelings, but I think that in some sense they are exaggerating a little bit, because I can certainly see that a lot of insights that we got in terms of these idealizations are very useful in terms of thinking about actual language and how you use it — what things mean that people actually say — and it really is applicable to that. I think perhaps that what Martin and Michiel are saying is a little bit (now this may sound nasty but I don't mean it that way) ... that it is philosophically naive. [The reasoning seems to be] that because it isn't right, therefore it's wrong. That because you can point out this lack of objectivity, because you can point out this problem, because you can show that it cannot be shown that this thing is objectively valid, that therefore it follows immediately that we have a big

problem. The only thing that follows is that it isn't completely obvious what the nature of the meaningfulness is. It does follow that it *can* be questioned. It does not immediately follow that the whole thing is meaningless and wrong, because that holds for almost any theory. If it really comes down to it we don't know, and if you are really honest about it, about most theories we know that they are false. Every empirical theory that we seriously work on we know is false, because we know the limitations, and not only in terms of things we haven't looked at, but we know realms where we know there are things it doesn't account for. We know that it only works to a certain extent with certain limitations, but that's not a reason to give up. That's how science works. So the fact that there are problems about objectivity is not necessarily a big deal. What you have to show is that it really isn't meaningful. If you just insist that theories must be completely right and proven, then you might as well just give up before you start.

I: My final question, coming back to the question of these two paradigms — the formal approach and the computational / statistical approach — where do you see the future of these disciplines? And where do your sympathies lie?

S: Well, like I said, I think ultimately the big challenge is in cognitive theories. One thing that one could say about philosophical theories (theories which have their context in a philosophical realm) is that in some sense the idealizations could be a bit arbitrary, in the sense that their justifications are only philosophical. In terms of practical applications, again, it is not an intellectual challenge, because there with big data you can get away with doing a lot without any theorizing or modeling or whatever. So if we have an interest ... let me just speak for myself — my interest as I said started with being intrigued by language as a cognitive phenomenon, and I think there is still a big challenge there. We have made big progress in the last half century in terms of understanding certain formal things, which only work in the context of certain idealizations, but which are nevertheless really interesting. We know a lot of such things. We also know a little bit about how to put statistics into the picture. I've been actively involved in launching the statistics stuff, and I've only been disappointed in that it became this independent paradigm which didn't have any connection with the modeling anymore. That was never my intention, but I believe in the importance of the statistics. I've written manifestos explaining just that,

but what I believe in is the synthesis between the statistics and the modeling approach, and I do believe that to get a viable picture of what's going on cognitively we will need that synthesis. I don't see any reason why that will not happen sometime in the future. These things do not intrinsically exclude each other. So if cognitive science survives as an intellectual, academic agenda ... for all I know, we may end up doing that at some point. So that's my vision. It really must come together. That's been my vision for a long time now, and it seems like it takes forever, but it might still be possible.

I: OK. These are all the questions that I wanted to ask you. Is there anything else that you would like to say before we end this interview?

S: No, I think you raised all the basic things, so I don't have any afterthoughts right away.

I: Thank you very much for your time.

S: OK. That was my pleasure!

Bibliography

Association for the Philosophy of Mathematical Practice. About, May 2014. URL `http://institucional.us.es/apmp/index_about.htm`.

Robert Audi, editor. *The Cambridge Dictionary of Philosophy*. Cambridge University Press, 1999.

Alexander Bird. Thomas kuhn. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Fall 2013 edition, 2013. URL `http://plato.stanford.edu/archives/fall2013/entries/thomas-kuhn/`.

Noam Chomsky. *Aspects of the Theory of Syntax*, volume 16. The Mit Press, 1965.

Paul Cozby and Scott Bates. *Methods in Behavioral Research*. McGraw-Hill Education, 2011. ISBN 9780078035159.

Catarina Dutilh Novaes. What is practice-based philosophy of logic? a case-study: uses of formal languages in logic, 2009. Paper presented at workshop: Practice-based Philosophy of Logic and Mathematics.

Catarina Dutilh Novaes. Towards a practice-based philosophy of logic: Formal languages as a case study. *Philosophia scientiae*, 16(1):71–102, 2012.

Uwe Flick. *Designing qualitative research*, volume 1 of *The SAGE Qualitative Research Kit*. Sage, 2009a.

Uwe Flick, editor. *The SAGE Qualitative Research Kit: Collection*. Qualitative Research Kit. SAGE Publications, 2009b. ISBN 9780761949749.

John D Heyl. Paradigms in social science. *Society*, 12(5):61–67, 1975.

Leon Horsten. Philosophy of mathematics. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Spring 2014 edition, 2014. URL `http://plato.stanford.edu/archives/spr2014/entries/philosophy-mathematics/`.

Institute for Logic, Language and Computation. Master's programme logic: Critical reflection. Technical report, Faculty of Science, University of Amsterdam, 2013.

Institute for Logic, Language and Computation. Logic and language, May 2014a. URL `http://www.illc.uva.nl/Research/Programmes/lola.php`.

Institute for Logic, Language and Computation. Scientific mission, May 2014b. URL `http://www.illc.uva.nl/AbouttheILLC/About/mission.html`.

Margaret King. Epilogue: on the relation between computational linguistics and formal semantics. In Michael Rosner and Roderick Johnson, editors, *Computational linguistics and formal semantics*. Cambridge University Press, 1992.

Joshua Knobe and Shaun Nichols. An experimental philosophy manifesto. *Experimental philosophy*, pages 3–14, 2008.

Thomas Kuhn. Logic of discovery or psychology of research. In Imre Lakatos and Alan Musgrave, editors, *Criticism and the Growth of Knowledge*, pages 1 – 23. Cambridge University Press, 1970a.

Thomas S. Kuhn. *The structure of scientific revolutions*. University of Chicago press, 1970b.

Steinar Kvale. *Doing interviews*, volume 2 of *The SAGE Qualitative Research Kit*. Sage, 2009.

Brendan Larvor. Review of paolo mancosu, ed. the philosophy of mathematical practice. oxford: Oxford university press, 2008. *Philosophia Mathematica*, 18 (3):350–360, 2010.

Brendan Larvor. What are mathematical cultures? (unpublished typescript). 2014.

Benedikt Löwe. Formalization: methodology and concrete examples. Presentation at lecture series at the ILLC: Logic, Language and Computation, 2012.

Benedikt Löwe and Thomas Müller, editors. *Philosophy of Mathematics: Sociological Aspects and Mathematical Practice*. College Publications, London, 2010.

Paolo Mancosu, editor. *The Philosophy of Mathematical Practice*. Oxford University Press, 2008.

Richard Montague. Universal grammar. *Theoria*, 36(3):373–398, 1970.

Barbara H. Partee. Formal semantics: Origins, issues, early impact. *The Baltic International Yearbook of Cognition, Logic and Communication*, 6(1): 13, 2011.

Barbara H. Partee. The amsterdam three and their part in a glorious history. Paper presented at "The dynamic, inquisitive, and visionary life of phi, ?-phi, and diamond-phi: A festschrift for Jeroen Groenendijk, Martin Stokhof and Frank Veltman, 2012.

W Keith Percival. The applicability of kuhn's paradigms to the history of linguistics. *Language*, pages 285–294, 1976.

Michael Rosner and Roderick Johnson, editors. *Computational linguistics and formal semantics*. Cambridge University Press, 1992.

Wilfrid S. Sellars. Philosophy and the scientific image of man. In Robert Colodny, editor, *Science, Perception, and Reality*, pages 35–78. Humanities Press/Ridgeview, 1963.

David Silverman and Amir Marvasti. *Doing qualtative research: A comprehensive guide*. Sage, 2009.

Keith Stenning and Michiel van Lambalgen. *Human reasoning and cognitive science*. MIT Press, 2008.

Martin Stokhof and Michiel van Lambalgen. Abstractions and idealisations: The construction of modern linguistics. *Theoretical linguistics*, 37(1-2):1–26, 2011a.

Martin Stokhof and Michiel van Lambalgen. Comments - to - comments. *Theoretical Linguistics*, 37(1-2):79–94, 2011b.

Johan van Benthem. Writing history: history versus practice of logic. Paper presented at workshop: Practice-based Philosophy of Logic and Mathematics, 2009.

Bart Van Kerkhove, editor. *New Perspectives on Mathematical Practices: Essays in Philosophy and History of Mathematics*. World Scientific, 2008.

Bart Van Kerkhove and Jean Paul Van Bendegem. Editorial preface to a selection of papers presented at the 'perspectives on mathematical practices conference 2002'. *Logique et Analyse*, 45(179-180):231–233, 2002.

Bart Van Kerkhove and Jean Paul Van Bendegem, editors. *Perspectives on mathematical practices: bringing together philosophy of mathematics, sociology of mathematics, and mathematics education*. Springer, 2007.