Univerzita Hradec Králové

Filozofická fakulta

Katedra filosofie a společenských věd

Language and Meaning in Community

Disertační práce

Autor: Matej Drobňák

Studijní program: P6101 Filozofie

Studijní obor: Filozofie

Forma studia: prezenční

Vedoucí práce: Prof. RNDr. Jaroslav Peregrin CSc.

Hradec Králové, 2018

Univerzita Hradec Králové
Filozofická fakulta

Zadání disertační práce

Autor:	**Mgr. Matej Drobňák**
Studium:	F14DP0001
Studijní program:	P6101 Filozofie
Studijní obor:	Filozofie
Název disertační práce:	**Jazyk a význam v společenství**
Název disertační práce AJ:	Language and Meaning in Community

Cíl, metody, literatura, předpoklady:

Práce má za cíl tematicky obsáhnout filozoficky relevantní aspekty vysvětlení vztahu mezi jazykem a jazykovým společenstvím. Vzájemná provázanost jazyka a společenství může znít triviálně. Jazyk je přece prostředkem komunikace mezi lidmi, a tudíž není ani pochyb o jejich vzájemných vztazích. Intuitivně se nabízí zejména představa jazyka, který je sdílený v rámci nějakého společenství všemi jeho členy. Tenhle intuitivní pohled se však začne komplikovat, pokud se snažíme o to, aby vysvětlení fungování jazyka ve společenství dokázalo obsáhnout i vysvětlení postavení jednotlivce vůči společenství a společnému jazyku. Z pohledu jednotlivce, který se snaží osvojit si jazyk, se "přidružení" k nějaké jazykové skupině nezdá být triviální záležitostí a snaha o vysvětlení způsobu jakým si jednotlivec jazyk osvojuje, vedla ve filozofii k otevření několika kontroverzních diskusí. Tato práce bude tematicky sledovat hlavní problémy, které souvisí se snahou o vysvětlení vztahu mezi jednotlivcem, společenstvím a jazykem. Jelikož jde o komplexní problém, pro vytvoření uceleného pohledu na vztah jazyka a společenství bude v práci třeba sledovat několik základních otázek/problémů: zejména a) problém sdílení jazyka ve vztahu k vysvětlení komunikace a b) otázku vymezení hranic jazyka společenství, která úzce souvisí s otázkou determinace významu.

Borg Emma (2004) Minimal Semantics. Borg Emma (2012) Pursuing Meaning. Brandom Robert (1994) Making It Explicit: Reasoning, Representing, and Discursive Commitment. Brandom Robert (2000 Articulating Reasons: An Introduction to Inferentialism. Davidson Donald (1984) Communication and Convention. Davidson Donald (1986) A Nice Derangement of Epitaphs. Davidson Donald (1994) The Social Aspect of Language. Dummett Michael (1986) A Nice Derangement of Epitaphs: Some Comments on Davidson and Hacking. Dummett Michael (1994) The Social Aspect of Language: Reply to Davidson. Dummett Michael (1996) The Seas of Language. Peregrin Jaroslav (2006) Meaning as an Inferential Role. Peregrin Jaroslav (2009) Inferentialism and the Compositionality of Meaning. Peregrin Jaroslav (2014) Inferentialism: Why Rules Matter. Peregrin Jaroslav (2017) Is inferentialism circular? Quine Van (1960) Word and Object. Quine Van (1969) Ontological Relativity. Recanati François (2004) Literal Meaning. Recanati François (2010) Truth-Conditional Pragmatics.

Garantující pracoviště:	Katedra filosofie a společenských věd, Filozofická fakulta
Vedoucí práce:	prof. RNDr. Jaroslav Peregrin, CSc.
Datum zadání závěrečné práce:	26.4.2017

Prohlášení

Prohlašuji, že jsem tuto disertační práci vypracoval samostatně pod vedením svého školitele Prof. RNDr. Jaroslava Peregrina CSc. a uvedl jsem všechny použité prameny a literaturu.

Declaration

I hereby declare that this thesis is my own work, elaborated under the supervision of my supervisor Prof. RNDr. Jaroslav Peregrin CSc., and that I make reference to all the sources and literature used in this thesis.

V Hradci Králové dne

Podpis

Annotation

Drobňák Matej. 2018. *Language and Meaning in Community*. Hradec Králové: Philosophical Faculty, University of Hradec Králové, pp. 152, PhD. Thesis.

In this thesis, I focus on philosophical explanations of natural languages and how they relate to explanations of other phenomena related to natural languages, such as successful communication and linguistic communities. Many philosophers take for granted the idea that speakers within a linguistic community share a language. However, it is also widely accepted that the notion of shared language is not sufficient to explain what makes communication successful – successful communication requires speakers to be able to take into account the influence of context on the content of utterances (including the intentions of speakers). This opens up interesting questions of a) how shared languages are employed in communication and b) how shared languages can be demarcated from within the linguistic production of speakers.

In the first part of the thesis, I criticize the currently predominant explanation of the relation between shared languages and communication, which is built on a Gricean semantic-pragmatic distinction. In the second part of the thesis, I discuss normative inferentialism and I try to show how inferentialism can be understood as a general framework for the explanation of the relation between natural languages, communication, and linguistic communities.

Keywords: Shared language, Communication, Inferentialism, Pragmatics

Acknowledgments

This PhD. thesis is a result of seven years of research and writing, starting at the Catholic University in Ružomberok and finishing here at the University of Hradec Králové. There are many people who have influenced my views during those years and to whom I am grateful for their help and guidance. Eugen Zeleňák, Daniel Cohnitz, Chris Gauker, and Gary Kemp are among those who helped me at various stages of the process, either by expanding my knowledge about metasemantics, and philosophy of language in general, or by giving me valuable feedback and various kinds of critical comments on some chapters included in this thesis.

I am especially grateful to my supervisor Jaroslav Peregrin, who carefully read and discussed with me each word of the thesis. He has had the biggest influence on shaping my views on language and philosophy and, ultimately, on my becoming who I am – a proponent of inferentialism. I would also like to give special thanks to Pavol Labuda for his endless support and trust at times when it was not clear at all that I might become a philosopher one day.

I would like to thank Ladislav Koreň and the whole Philosophical Faculty at the University of Hradec Králové for creating an ideal academic environment for my work on the thesis and Simon Gill for his help with proofreading. Last but not least, I would like to give thanks to my family, and especially to my wife Zuzka, for their patience and support and to apologize for talking too much about philosophy – sorry, it will not get any better.

Contents

1. Introduction

A word is dead
When it is said,
Some say.
I say it just
Begins to live
That day.
(Emily Dickinson)

With a little simplification, we can say that the topic of this thesis belongs to the field of metasemantics, as broadly construed. As the prefix 'meta' indicates, I am not interested in creating and developing theories of meaning, or semantic models of languages. Rather, I am interested in meta-questions about the nature of meaning and language. I understand that talking about the "nature of meaning and language" may sound a little mysterious, so in the next few pages I will try to explain what I think is hidden under this mysterious label.

First and foremost, it is important to emphasize that I am particularly interested in natural languages. There is no doubt that formal semantics, the aim of which is to build semantic models of formal languages, has proved to be a fruitful field of study throughout the last century. Formal languages, as opposed to

natural languages, are easier to study as their properties are well-defined and the meanings of expressions are considered to be stable and determinate. I do not doubt that studying formal languages that may have nothing in common with natural languages can bring interesting results. However, in this thesis I am more interested in a specific trait of semantics which understands formal languages as a gateway to natural languages. Natural languages are rather complex structures with many specific properties. Formal languages, on the other hand, can be defined in such ways that they systematically overlook some properties and highlight others. Because of that, focusing on formal languages sounds like a reasonable method for delivering an exhaustive description of natural languages step by step – by studying one property at a time.

For example, formal languages were initially defined in such a way that the meanings of expressions were considered to be context-invariant, i.e. fixed regardless of the context in which they appeared. But because the interest of formal semanticists started to shift towards natural languages a few decades ago, we can now find many attempts (e.g. Kaplan 1978; Lewis 1970, 1980; Stanley and Szabo 2000) to incorporate context-dependent properties of natural languages (e.g. indexical expressions, temporal expressions, or context-dependent features studied by dynamic logic) into formal languages.

The most general question that led me to write this thesis was whether the method of gradual expansion of formal languages can deliver an accurate depiction of natural languages. Answering this question requires a careful reconsideration of how we should approach natural languages and what makes the

depiction accurate, hence the "meta" character of my thesis. Laying my cards on the table, I do not engage with the question explicitly throughout the whole thesis. Rather, the chapters included in this thesis reflect the development of my views over the span of my research.

In particular, the question of how we should approach the study of natural languages is scrutinized from three different perspectives: A) whether the notion of language, as understood within formal semantics, can depict natural languages accurately; B) whether the depiction of meaning as stable and determinate can depict semantic aspects of natural languages accurately, and C) what kind of specification of meaning ensures the most accurate and fruitful study of semantic aspects of natural languages. Strictly speaking, only B) actually engages with a topic within the scope of metasemantics. A) and C) are more likely to belong to something that we might call philosophy of linguistics. They provide a reconsideration of what theoretical tools semanticists, philosophers of language, or linguists (should) rely on for the study of natural languages.

During my research, I had to face one unexpected challenge. Answers to A), B), and C) are often taken for granted as part of a theoretical background and, unfortunately, detailed discussions that compare different proposals are not very common. Because of that, I still find somewhat problematic to classify my thesis as belonging to one clearly delimited subfield of study/tradition, or as following a single discussion within some field. This is, from what I know, rather unusual for PhD. theses and it caused the thesis to be distinctive in some respects. Notably, I do not start the thesis by presenting one single historical background to the topic.

This is not to say that there is none, but I could not get rid of the feeling that the more I attempted to follow one historically grounded discussion, the more I was moving away from the topic of the thesis. What is more, the content of the last chapter of this thesis aims to give a theoretical background to empirical research within the framework of inferentialism. While inferentialism is a well-established doctrine by itself, there has been no attempt to link it to empirical research so far and so it is not possible to present a historical background to the topic. Following one's own path is always a risky endeavour. I hope that the originality of the thesis partially excuses this unusual feature.

In any case, even though I do not present a historical background to the topic of the thesis in general, each chapter included in this thesis incorporates a careful presentation of the state of play for the particular topic discussed and thus links my research and views to well-established and historically grounded academic discussions. In addition, I will try to present a general context of academic discussions that served as a background for this thesis in the rest of this introduction.

The thesis can be divided thematically into two parts and each part consists of two chapters. In the first part, I discuss the explanatory role of the notion of shared language. In the second part, I discuss specific aspects of inferentialism related to the nature and specification of meaning. Both parts are a result of my long-lasting interest in natural languages and, especially, in how we use or implement natural languages in communication with other speakers. On the one hand, there is the idea of natural languages as stable and rather static

structures which evolve over centuries. On the other hand, there is the idea of communication as a creative activity in which reaching understanding often lies at the end of a winding and mysterious path. Both these ideas are intuitively acceptable and deeply rooted in how we think about language. However, the answer to the question of how these two ideas relate to each other is not so self-evident and it has caused several controversies within philosophy of language. This thesis represents my attempt to discuss different proposals and to find the best philosophical doctrine for the reconciliation of the idea of language as a stable structure and the idea of communication as a creative activity.

The first two chapters are focused on the notion of shared language and on the role that the notion plays in different philosophical explanations of how natural languages work. I think that many philosophers, as well as lay persons, take the idea that we share languages for granted. I admit that the idea of shared language is intuitively appealing and it is strongly supported by our everyday experience – especially by the fact we are mostly able to communicate with other people smoothly and misunderstandings are rather rare (at least when compared to the number of cases of successful communication). The simplest explanation of success in communication is that we share a language in a sense that we have all learned the same lexicon and syntactic rules and we use them in communication. If I understand all the expressions exactly as you do, then it is no surprise that we understand each other and our communication is successful. Such a relation between language and communication is, I believe, intuitively appealing for lay persons, but it has been significantly criticized at least since Davidson (1986). In

this paper, Davidson argues that the notion of shared language is neither necessary, nor sufficient to explain what makes communication successful as success in communication is a result of our ability to follow the intentions of speakers.

Reading Davidson's paper inspired me to conduct research on doctrines that do not rely on the notion of shared language at all – not only with respect to the explanation of communication, but with respect to the explanation of any aspect of natural languages. Surprisingly, I discovered that alternatives are rather sparse and not well-received within philosophy of language. Meaning holism (Rapaport 2000, 2003; Pollock 2014) is probably the most prominent doctrine that aims to explain how natural languages work without any dependence on the notion of shared language. Meaning holists take the notion of idiolects as fundamental. From this point of view, there is no such thing as a shared language of a community and successful communication is a result of our ability to detect and reinterpret specific features of the idiolects of other speakers. Even though meaning holism is probably the most prominent alternative to the doctrines that rely on the notion of shared language, it is still often considered to be "counterintuitive", with the label "counterintuitive" covering specific problems with which this view struggles. For example, if there is no language shared among members of some community, then it is not clear how language acquisition can be possible (or even what should be acquired).

Another alternative can be found in Ludlow (2014) and his notion of microlanguages. The main idea of this position is that there is no such thing as a

shared language of a community. Rather, speakers share microlanguages. It is hypothesized that microlanguages are shared locally and they are specifically devised by particular speakers for the purposes of particular conversations. Ludlow's project can be understood as an elaboration and a continuation of Davidson (1986). However, so far he has only been able to present a sketch of a full-fledged theory and its viability is a matter of future development.

Considering the deficiencies and problems of alternative explanations of natural languages, I decided to focus on standard and well-established doctrines in the subsequent research. More specifically, I focused on the explanatory roles that the notion of idiolects and the notion of shared language play in those doctrines. The question whether the notion of shared language should have primacy over the notion of idiolects (or vice versa) turned out to be especially fruitful with regard to Quine (1960, 1992) and his views on natural languages. On the one hand, Quine often talks about language as a deeply social phenomenon, but, on the other hand, when he defines stimulus meaning he refuses to rely on the notion of shared language and he sticks to the notion of idiolects. This is an interesting clash in his views, and rather unusual within philosophy of language. As I mentioned earlier, the questions discussed in this thesis are often taken for granted and philosophers do not usually discuss them explicitly – and this applies to Quine as well. Because of that I tried to retrace the reasons that might have led Quine to his preference for the notion of idiolects and to clarify his position in detail.

As I try to show in the first chapter included in this thesis, his preference for idiolects has roots in his thesis about the public nature of meaning and

indeterminacy with regard to natural languages. According to Quine, behavioural data about the overt linguistic behaviour of speakers is not sufficient to fully determine the meaning of an expression. If it is not possible to fully determine the meaning of an expression, then it makes no sense to ask whether two speakers ascribe the same meanings to the same expressions. What is more, if we consider the sameness of ascribed meanings to be among the relevant criteria for language identity, and it is not possible to decide whether two speakers ascribe the same meanings to the same expressions, then it is not possible to decide whether two speakers share the same language or not. Even though Quine has never used this argument explicitly, I believe that there are some clues in his writings that suggest that this was the reason for his preference for the notion of idiolects and for abandoning the notion of shared language. The content of this chapter has already been published as a research paper in the journal *Philosophia* under the name "Quine on Shared Language and Linguistic Communities".

Reading Davidson (1986), while bearing the reason for Quine's preference for the notion of idiolects in mind, inspired me also to think about other explanatory roles that the notion of shared language plays in contemporary philosophy of language. The view that the notion of shared language is not sufficient for the explanation of successful communication is predominant nowadays and it is accepted even by the most enthusiastic proponents of shared languages. However, if communication is not a simple application of natural languages, then the need for the explanation of the relation between shared language and communication arises. As far as I can see, there are two strong

approaches that attempt to resolve this issue currently under discussion. Lewis (1979) proposed an account of communication based on the notion of scorekeeping in language games. According to this view, when a hearer interprets a speaker, she has to keep track of what sentences mean in a shared language and what sentences mean according to a speaker (allowing that these two things often overlap). This view was later adopted by inferentialists, especially by Brandom (1994, 2000). The second and more influential approach that aims to explain the relation between shared language and communication stems from the work of Grice (1957, 1961). Distinctions between semantic and pragmatic aspects of content, and between linguistic meaning and the speaker's meaning, dominate current philosophy of language.

In the second chapter included in this thesis, I discuss the currently predominant version of the semantic-pragmatic distinction called minimal semantics and advocated by Borg (2004, 2012). Borg explains success in communication as a two-step process. First, we semantically process a sentence and so we understand its linguistic (literal) meaning. Second, pragmatic influence operates and modifies the meaning of a sentence to the final content of an utterance. Interestingly, the first part – semantic processing – is considered to be congruent for all speakers and so it opens a way for delimiting the shared language of a community. Simply put, we all share a language with regard to semantic processing. According to the two-step model, the notion of shared language is necessary, though not sufficient for the explanation of what makes

communication successful. I consider this to be a legitimate explanation of successful communication (though not necessarily the correct one).

However, I believe that the reason that led Quine to his rejection of the notion of shared language, namely the indeterminate nature of meaning, poses a serious challenge for any proponent of the notion of shared language and I do not see how minimal semantics can avoid it. The challenge is simple: if it is not possible to determine whether two speakers ascribe the same meanings to the same expressions, then how can we decide whether they share the same language or not? In the second chapter included in this thesis, I try to spell out this challenge not only in terms of Quine's indeterminacy thesis, but also in terms of Waismann's thesis of open texture and other related sceptical arguments in semantics and I discuss how they problematize the notion of shared language as understood within the framework of minimal semantics.

As the topics of particular chapters reflect my views at different stages of my research, there are many assumptions hidden in the background that might not be clear from reading particular chapters of the thesis alone. Because of that, my discussion of inferentialism in the second half of the thesis may look like a sudden change of topic. In the rest of this introduction, I would like to show why this is not true and to clarify the reasons for my inclination towards inferentialism. In short, the inclination results from my conviction that inferentialism can deliver a persuasive response to Quine's thesis about the indeterminate nature of meaning and the subsequent argument that problematizes the notion of shared language.

It is necessary to understand Quine's reasons for the indeterminacy of meaning to be able to understand why I think inferentialism avoids the problem. According to Quine, our linguistic competence is a result of language learning in which we start with learning observation sentences. The learning process requires several components – a hearer (a learner), competent speakers, and a situation in which the "lecture" takes place. A hearer learns an observation sentence by observing: A) what sentence a competent speaker uses; B) in what situations, and C) how other speakers react to this use. The corrective behaviour of other speakers plays an important role in this process as it helps a hearer to single out situations in which it is acceptable to use a sentence. In such situations, a hearer is looking for a correlation between uttering a sentence and the presence of a specific type of stimulus; standardly caused by the presence of a specific object in the visual field (Quine uses rabbits as an example). External stimuli caused by perceived objects, which must be present at the time when a sentence is uttered, are then understood by Quine as constituents of a set that represents the stimulus meaning of a sentence.

The reason why Quine believes that the meaning of the majority of sentences is indeterminate lies in the fact that it is not possible to single out such a set of stimuli for all the sentences within a language. For example, the corrective behaviour of speakers with regard to abstract or standing sentences does not depend on the presence of any specific types of objects at the time when a sentence is uttered. This leads Quine to the conclusion that it is not possible to fully determine the meaning of sentences beyond observation sentences and this

leads to the conclusion that it is not possible to decide whether two speakers share a language or not.

I believe that the solution to avoiding Quine's challenge lies in changing the requirements for external stimuli that must be singled out as the stimulus meaning. The change required for avoiding indeterminacy consists in redirecting the focus of a hearer. If a hearer wants to learn a sentence beyond observation sentences, her aim must be to find a correlation between the occurrence of a sentence and the occurrence of other sentences. Co-occurring sentences, or rather the stimuli caused by their being uttered, can then be understood as external stimuli that must be singled out as the meaning of a sentence. With a little simplification we can say that if we want to understand observation sentences, we have to focus on rabbits; if we want to understand standing sentences we have to focus on other sentences that have been uttered. In such a case, the corrective behaviour of other speakers plays basically the same role as it plays in the case of observation sentences – it helps a hearer to decide which correlations are acceptable and which are not, i.e. it helps a hearer to single out a set of sentences that a competent speaker must be disposed to accept in order to understand a sentence. This brings us directly to the vicinity of the inferential view on meaning. According to inferentialists, the meaning of a sentence is determined by the inferential relations of the sentence to other sentences within a language. If this is so, then, roughly speaking, the meaning of a sentence can be specified as a set of sets of sentences from which the sentence can be inferred and which can be inferred from the sentence.

In other words, it is not completely true that there are no stimuli present in the case of non-observation sentences, as Quine believed. Hearers can find correlations between a sentence that is uttered and other sentences that are, or at least can be, appropriately uttered before or after the sentence. This is not exactly what Quine meant by external stimuli. But the stimuli caused by hearing other sentences are external stimuli anyway. The conviction that inferentialism can avoid Quine's challenge influenced my subsequent research and it explains and vindicates what might look like a sudden change in the topic discussed in the second half of the thesis.

Even though I believe that inferentialism offers the best solution for the challenge posed by indeterminacy, there are still many other challenges that inferentialism has to face on its own. In the third chapter included in this thesis, I discuss Peregrin (2014b) and how his reliance on widely shared corrective behaviour can be used as a criterion for deciding which inferences are meaning-constitutive and which are not. The objection that inferentialism cannot provide such a distinction has famously been raised by Fodor and Lepore (2001, 2007) and it is still one of the general problems that prevent a wider acceptance of inferentialism.

As Peregrin emphasizes, our acceptance of inferential relations between sentences is shaped by other speakers from whom we learn a language. The acceptance of some inferences has a specific status as it is required as a proof of a basic linguistic competence.

My suggestion, presented in the third chapter of this thesis, is to use Peregrin's idea to back a criterion of meaning-constitutiveness – only those inferences that are widely accepted by almost all the speakers within a community can be considered as meaning-constitutive. I admit that, if we are talking about natural languages, the boundary between widely accepted and less than widely accepted inferences can often be rather vague and it can easily turn out that meaning-constitutiveness is more a matter of degree than a matter of clear-cut distinction. This may seem to be a weakness of the criterion, but, quite the opposite, I find it to be its strength. The inferential criterion of meaning-constitutiveness is naturalistic in nature. It makes inferentialism accountable to actual data about natural languages and the linguistic practices of actual speakers. The fact that the criterion raises doubts about clear boundaries of meaning is not a deficiency of inferentialism; it is an important insight into the nature of meaning in natural languages. The content of this chapter has already been published as a research paper in the journal *Organon F* under the name "Meaning-constitutive Inferences".

Surprisingly, there has been no attempt to bring inferentialism closer to empirical research, despite its naturalistic character and inclinations towards natural languages and the linguistic practices of speakers. The idea of an empirically grounded philosophy of language/semantics (e.g. Machery et al. 2004, 2009) is becoming more and more popular and it may easily become a rule in the near future. I welcome such a tendency. Ultimately, inclinations towards natural sciences have been at the heart of analytic philosophy since its beginnings – even

though the view on what the relation between the natural sciences and philosophy should look like has changed radically.

I believe that inferentialism has huge potential with regard to empirical research in semantics. The last chapter included in this thesis is an attempt to present a theoretical background for future empirical research based on inferentialism. My aim is to show how the framework of inferentialism could be used to explain several phenomena related to natural languages (ambiguity, free pragmatic enrichment, and conversational implicatures) and to show how the inferential specification of meaning can serve as a platform for gathering empirical data about semantic aspects of natural languages. My proposal is to understand the difference between the inferential potential and the inferential significance of a sentence as the difference between the meaning of a sentence and the content of an utterance. This allows inferentialism to accommodate the influence of context on content that is typical of natural languages and their use in communication. The fact that the meaning of a sentence is specified as a set of sets of sentences that can be inferred from the sentence or from which the sentence can be inferred makes inferentialism easily testable. If we want to learn something about semantic aspects of natural languages, we only need to focus on the commitments and entitlements that are created by the utterances of particular sentences in particular conversations. It is an open question whether the project of inferentialism on empirical grounds will be executed one day. There are many details about theoretical explanations and proper methodology to be discussed. However, I believe that inferentialism, with its intrinsic inclination towards

natural languages, is a doctrine with a great potential and that the project can be successful.

With regard to this wider project, the value of this thesis does not lie in presenting definite and eternal answers, but in establishing a strong background for future development. It presents the reasons why inferentialism represents the best alternative for the reconciliation of the idea of natural languages as shared and the idea of communication as a creative activity. What is more, it shows directions for a further elaboration of inferentialism as a general empirically-grounded framework for natural languages. I leave it to the reader to decide whether this is a virtue or a vice of this thesis.

2. Quine on Shared Language and Linguistic Communities

If we want to give a general account on what natural languages are, how they are formed by speakers and how communication works we have basically two options: to take either the notion of shared language or the notion of idiolect as fundamental. We can include Lewis (1969, 1973) or normative inferentialists such as Brandom (1994, 2000) and Peregrin (2014b) among those who take the notion of shared language as fundamental. On the other side of the spectrum, we can find people like Chomsky (1986, 1992, 1995) or Davidson's views on communication (1986, 1994) and some of his followers[1].

If we look at the first group, a natural language is standardly understood as a set of syntactic rules and meaningful expressions.[2] It is taken for granted that such a language is accessible for speakers – many speakers can acquire meanings in such a way that one expression has the same meaning for all of them.[3] If different speakers acquired the same language, or at least a sufficiently large part of it, and so they partake in the same set of syntactic rules and meaningful expressions, we can say that they share a language.[4] The question of how it is

[1] E.g. Ludlow (2014) or Glüer (2013), though their interpretations of Davidson differ.

[2] Even though there is no general agreement on how to specify meanings (Lewis prefer possible-world semantics, Brandom is an inferentialist...), no one doubts that there is something like a semantic aspect related to our linguistic production.

[3] In a similar way to how one sentence can express the same proposition for different speakers or in different occurrences.

[4] Notice that partaking in a language does not require that a speaker is able to reflect what meanings she associates with expressions of the language. However, this does not mean that it is not possible to decide whether two speakers possess the same language. For example, according to

possible that a shared language emerged within a community is then answered differently – as an outcome of rational decisions of speakers to follow the majority or as an outcome of interlocking normative attitudes of speakers. The biggest advantage of this approach is that the notion of shared language has a high explanatory value. We can rely on it if we want to explain language learning, successful communication or it can serve as a criterion which delimits the boundaries of linguistic communities. Naïve communitarism is then an extreme position which takes the notion of shared language for granted in order to explain all the aforementioned phenomena:

A) a linguistic community is a group of speakers who share the same language;
B) the outcome of the process of language learning is to acquire the shared language;
C) understanding in communication is achieved because a speaker and a hearer use the same (shared) language.

The biggest weakness of this approach is that it tends to oversimplify. In praxis, philosophers who take the notion of shared language as fundamental usually do not accept all the explanations above. Especially C) seems to be an oversimplification[5] and we can find advanced accounts on communication in Lewis (1979) and Brandom (1994). On the other hand, we can find also authors

Peregrin (2014b), we can decide whether two speakers possess the same language by tracking their corrective behaviour.

[5] See Davidson (1986, 1994) for the discussion.

(Dummett 1986, 1994; Hornsby 2008; Weiss 2010) who argue that the notion of shared language is inevitable for the explanation of communication.

Authors who take the notion of idiolect as fundamental are sparse and in general, this approach has not been well-elaborated so far. In some sense, we can say that Chomsky is one of them because he believes that the notion of shared language is redundant.[6] However, he still believes that speakers harmonize as a result of evolutionary predispositions and language learning.[7] This makes his position to be close to a middle ground. A more radical proponent of this approach is Ludlow (2014). He takes Davidson's mantra that there is no such thing as a language seriously and argues that there are only conversational microlanguages and that meanings are dynamic; i.e. they can change from conversation to conversation, from speaker to speaker. The biggest advantage of this approach is that its proponents aim to give empirically well-founded conception without relying on abstract entities (such as, for example, propositions can be). The extreme position in this case would be quite close to Ludlow's view and we can call it linguistic individualism. A view according to which:

X) there is no such thing as a shared language of community; language sharing is only local – with regards to the number of people as well as the number of expressions shared;

Y) the outcome of the process of language learning is an enhanced ability of speakers to adjust their communicational strategies for particular conversations/hearers;

[6] See Chomsky (1986, 1992, 1995).

[7] Probably, Chomsky would not have a problem to accept B) from the list above.

Z) understanding in communication is achieved when a hearer is able to identify speakers intentions.

In this chapter, I would like to discuss whether Quine stands closer to naïve communitarism or to linguistic individualism. The idea that language is primarily social is at the heart of his views. As he famously states: “Language is a social art” (Quine 1960: ix) and he continues to talk about language as “ours” and “social” and about “shared conditions” which lead to a language at many places in his writings. Numerous references may create the impression that his view is close to some kind of naïve communitarism.

This way of interpreting Quine is sometimes explicitly and often implicitly suggested by commentators. I must admit that it is a tempting interpretation, especially if we look at the role which Quine ascribes to the corrective behaviour of speakers in the process of language teaching. However, I do not believe that any kind of naïve communitarism can be legitimately ascribed to Quine. If we look at his views on language learning in detail, we can conclude that his way of talking about natural languages as shared is no more than a provisional simplification. On the contrary, I believe that Quine is much closer to linguistic individualism and that he even made explicit commitments to it.

Quine did not discuss how he understands linguistic communities and language sharing in detail. However, I believe that there is an interpretation which is coherent with his more general views on language (as formulated e.g. in Word and Object) and which was explicitly endorsed by Quine in his latest writings.

Therefore, the main aim of this chapter is to clarify the role of linguistic communities in Quine's views on language and to clarify his view on language sharing. This aim ultimately reduces to answering the question "What does Quine mean when he says that language is a social art?"

In subchapters 2.1 and 2.2, I will present motivations which might support naïve communitarism as well as examples of how naïve communitarism occurs in the literature about Quine. In subchapters 2.3 and 2.4, I will argue for what I see as a more accurate interpretation of his views.

2.1 Corrective Behaviour

"E pluribus unum" is the name of the second chapter of Word and Object and it means "One from many". Even at this early stage of Quine's work, we can find a reference to sociality – an individual is only a part of a bigger community; one from many. Other references to the social dimension of language can be found mainly in those passages in which Quine discusses how language is taught. As he states: "Language is socially inculcated and controlled" (Quine 1969a: 81). The community of speakers of a language plays an important role as speakers are responsible for the socialization of a newcomer (let this be a child or a linguist from the famous example). Speakers repeatedly manipulate a newcomer until she conforms; that is, she is able to use well-formed sentences in expected circumstances, or, as Quine states, until she "ends up using the word to suit"

(Quine 1960: 7).[8] To compel a newcomer to conform, speakers use corrective behaviour. Corrective behaviour includes negative as well as positive reactions on the part of a community – punishments and rewards – towards the linguistic use of a newcomer. The influence of a linguistic community and its corrective behaviour as a tool of language teaching is so strong that Quine even calls it "the objective pull".

> The objective pull will regiment all the responses still as 'red', by activating myriad corrective cues. These corrective cues are used unconsciously, such is the perfection of our socialization. (Quine 1960: 8)

A linguistic community can, in particular, push a newcomer to learn which associations between a word/sentence and stimuli are acceptable and which are not, or, roughly speaking, in which situations it is acceptable to use a word/sentence.

The authority which Quine attributes to communities and their corrective behaviour seems to be absolute. Communities seem to act as uniform entities and the corrective apparatus that they use seems to sustain the conformity of speakers and, as its consequence, sharing of a language by speakers as well. Therefore, mastering a language seems to be the same as mastering a sufficient subset of the same set of syntactic rules and meaningful expressions as the rest of the

[8] Notice how vague the phrase originally used by Quine is. This is the first hint that his view is not as straightforward as it may seem.

community possess. So if we look at Quine's view on language teaching, it seems that he is a proponent of theses A) and B).

2.1.1 Communities and Natural Languages

Another way in which the sociality of language appears in Quine's writings is rather implicit. It is in the *way* that he talks about natural languages and linguistic communities. Quine often uses formulations such as 'native language' or 'English', e.g. in relation to translation. For example, he talks about the English language compared to a native language when discussing the conditions under which a sentence in one language can be translated into the other one:

> Suppose both translations, along with these accommodations in each case, accord equally well with all observable behavior on the part of speakers of the remote language and speakers of English... (Quine 1969c: 29)

And several pages later: "Insofar as the native sentences and the thus associated English ones seem to match up…" (Quine 1969c: 33). Such formulations suggest that there is *English* and there is *the native language*, as two sets of syntactic rules and meaningful expressions, which sentences can or cannot match up.

Elsewhere, he talks about English speaking communities: "Even we who grew up together and learned English at the same knee..." (Quine 1969e: 5) or

> Implicitly the learner of "yellow" is working inductively toward a general law of English verbal behavior, though a law that he will never try to state; he is working up to where he can in general judge when an English speaker would assent to "yellow" and when not. (Quine 1969b: 125)

Such formulations may give the impression that there are speakers who mastered English and those speakers who succeeded in mastering the language form a community. The task for a newcomer is then to find out what *the language* is and to become a member of the community by possessing the same language as others do. In other words, all the quotes above give the impression that Quine might be a proponent of theses A) and B) of naïve communitarism.

2.2 Commentators

Naïve communitarism is sometimes explicitly ascribed to Quine and often implicitly indicated by many philosophers. An example of the first case is Mark Risjord:

> Quine had no qualms about writing about linguistic communities. Indeed, he gave them the causal powers to mold an individual's speech dispositions. For Quine, an individual's ability to speak and communicate is ultimately a set of dispositions. These vary among individuals, but our social interaction shapes them into similar outward forms. In Word and Object, he invoked the

> image of two topiaries; while they share the same outward form, their configurations of branches and leaves are different (Quine 1960: 8). The *uniformity among speakers' total dispositional set gives language its "objective pull" away from idiosyncratic associations* and toward a public realm of words and objects. Quine uses "the community" in a naturalistic way. It causes an individual to develop speech dispositions that are similar to the rest of the community, thereby *creating a common language*. Thus for Quine, the model of a linguistic community is exactly the idealized form of culture found in mid-century anthropology. It was static, stable, and exercised control over the individuals. (Risjord 2007: 208, emphasis added)

In this quote, Risjord appeals to the role of community in language teaching in a similar way to that in which it is presented in subchapter 2.1.1. The social pressure generated by a community leads to "creating a common language", i.e. to creating one language which can be shared by all members of a community.

Another example of an explicit ascription of the thesis A) of naïve communitarism to Quine can be found in Stoutland (2000). According to Stoutland, Quine's approach to meaning is normative and the normative aspect lies in the uniformity of speakers. For each expression of a language there is the correct meaning. To become a member of a linguistic community requires learning the correct meanings for expressions of a language:

> On this view, the normativity necessary for linguistic meaning and competence derives entirely from the fact that speakers of a language come

> to respond uniformly to certain sentences, so that those who conform use the sentences correctly, those who deviate use them incorrectly. These sentences are such that *speakers become competent in the language by acquiring the same dispositions* to assent or dissent to them, so that a speaker who does not share those dispositions does not understand what the sentences mean. (Stoutland 2000: 184–185, emphasis added)

An example of an implicit acceptance of a similar view can be found e.g. in Hylton's book, aptly called 'Quine':

> You and I have an understanding of English in common, but our behaviour is not the same. It is not just that your actions are yours and mine are mine; it is not at all obvious that your behaviour and mine can be coordinated so as to make it clear how each amounts to the same thing – an understanding of English. (Hylton 2007: 103–104)

In this quote, Hylton talks about the outcome of language teaching (coordination of behaviour) as leading to understanding of "the same thing" – the natural language, namely English. Such a formulation gives the impression that Quine believes that there is one set of syntactic rules and meaningful expressions which is shared by speakers and so that he is a proponent of theses A) and B) from the list above.[9]

[9] Hylton seems to believe that this is the outcome of language learning even though he is aware that it is hard to explain how it is possible solely on behaviouristic grounds. To be fair, in the later comments in the book Hylton suggests that the issue is not so clear as it may seem.

In the rest of this chapter, I will try to show that Quine's view is much closer to linguistic individualism as it may now seem because he explicitly committed himself to idiolects as fundamental for the explanation of natural languages. Clarifying this issue is important for exegetical reasons – to avoid possible misleading interpretations. But most significantly, the issue is important with regards to attempts to find alternative theoretical frameworks for natural languages. If I am right, Quine's work can be seen as one of the first attempts to deliver a complex and coherent view on natural languages and related phenomena without relying on the notion of shared language.

2.3 Clarificatory Remarks

Before we move further, two points deserve clarification, firstly, what it means 'to share a language', and secondly, Quine's negative attitude towards meaning.

Sharing a language can be trivial in some sense. If we define a language as a set of well-formed graphemes/phonemes organized into units corresponding to our words and sentences – without focusing on their meaningfulness – then we can decide if two speakers share a language solely on the basis of the sounds and inscriptions they make. But there is a reason why we are interested in sharing. The notion of shared language is supposed to play some role in explanations. It can play a role e.g. in the explanation of how language learning works[10] and it is hard

[10] In general, the list may be longer. The notion of shared language can play an important role in the explanation of communication, translation and so on.

to explain how speakers learn a language without distinguishing meaningful utterances. So for our purposes, language sharing should be based on a semantic comparison. To say that two speakers share a language is to say that they possess sufficiently big subsets of the same set of *meaningful* expressions and syntactic rules. In some sense, we can say that different speakers share a language if they associate (explicitly or implicitly) the same meanings with the same expressions of a language.

This brings us directly to the second remark. In general, Quine refuses to talk about meaning as it is only a relic of a mentalistic view on language. Of course, it does not mean that there are no semantic phenomena but his insistence on naturalism[11] leads him to behaviourism and this forces him to see semantic phenomena as intrinsically behavioural:

> In psychology one may or may not be a behaviorist, but in linguistics one has no choice. Each of us learns his language by observing other people's verbal behavior and having his own faltering verbal behavior observed and reinforced or corrected by others… There is nothing in linguistic meaning beyond what is to be gleaned from overt behavior in observable circumstances. (Quine 1992: 37–38)

And verbal behaviour is nothing more than an overt demonstration of speech dispositions:

[11] Quine sees his inclination towards behaviourism as "simply a proposal to approach semantical matters in the empirical spirit of natural science" (Quine 1970b, 8).

> For naturalism the question whether two expressions are alike or unlike in meaning has no determinate answer, known or unknown, except insofar as the answer is settled in principle by people's speech dispositions, known or unknown. (Quine 1969c: 29)

Speech dispositions mentioned above are dispositions to assent or to dissent to sentences when prompted by stimuli.[12] Dispositions, which link our response and some set of stimuli, are an outcome of language learning and they can be identified as specific neural configurations or processes of the brain.[13] This is because a disposition is, in some sense, a mental state and, as Quine believes, mental states are neural states of the brain.

> The deepest explanation, the physiological, would analyse these dispositions in explicit terms of nerve impulses and other anatomically and chemically organic processes. (Quine 1975: 255).[14]

Unfortunately, such a view leads to a problem. If we specify meanings as dispositions and dispositions as neural processes of a brain, then to say that two speakers associate the same meaning with an expression would require that two speakers have the same neural processes. If this is so, then, strictly speaking, no

[12] "Query and assent, query and dissent – here is the solvent that reduces understanding to verbal disposition." (Quine 1975: 249)

[13] Stimuli are also understood physiologically as sets of triggered receptors. "The stimulation that he undergoes at any moment is the set of receptors triggered at that moment." (Quine 1979: 5) "If querying the sentence elicits assent from the given speaker on one occasion it will elicit assent likewise on any other occasion when the same total set of receptors is triggered; and similarly for dissent." (Quine 1981: 25)

[14] See also Quine (1985: 323) and Quine (1989: 347).

two people can share the meaning of an expression because no two people have identical neural configurations.[15] It all starts with the subjectivity of stimulations:

> The stimulation that he undergoes at any moment is the set of receptors triggered at that moment. This makes good sense of sameness and difference of stimulation of that person from moment to moment. It does not make sense of sameness of stimulation of two persons, since two persons do not share the same receptors. (Quine 1979: 5)

Quine did not fully realize this problem in Word and Object, where he seems to incline to sharing distal causes (stimulatory situations).[16] But it is explicitly discussed in Quine (1969d: 157) and he suggested an answer by taking one step back from neural configurations to sets of stimuli in Quine (1974) and later on in Quine (1995). In other words, he decided to specify meanings by relying on the notion of *perceptual similarity*. The idea is simple – we cannot have the same neural processes, but we can have the same standards of perceptual similarity. The perceptual similarity of stimuli is testable. We can find out which stimuli are perceptually similar for a speaker by varying the external conditions and demanding reactions. If two speakers react in the same way under the same conditions, then we can say that they share standards of perceptual similarity. Moreover, in his later writings, Quine accepted an evolutionary explanation of

[15] For a detailed discussion of this point see Kemp (2017).

[16] The idea of shared distal causes is typical of Davidson (1991a, 1991b). Quine decided to abandon such a view in favour of proximal causes shortly after 1960. In Pursuit of Truth Quine summarizes: "...so I could place the stimulus out where Davidson does without finessing any reification on the subject's part. But I am put off by the vagueness of shared situations" (Quine 1992: 42).

why we (more or less) share standards of perceptual similarity. It is because we are evolutionarily predisposed to consider the same stimuli as perceptually similar.[17]

So if we specify meanings as subjective processes of brains then, strictly speaking, sharing meanings is not possible because no two people have exactly the same neural configurations. The question is whether shared standards of perceptual similarity could do the trick instead. Whether there is a specification of meanings, on the basis of shared standards of perceptual similarity, which would enable two speakers to associate the same meanings with expressions of a language, i.e. to share a language.

2.4 Quine on Learning, Sharing and Communities

We did not state anything that contradicts naïve communitarism so far. The situation starts to be less favourable towards it if we focus on how Quine presents the position of an individual in learning a language at length. In some cases, an individual acquired a sentence if she is able to single out an acceptable set of sensory stimuli which must be undergone at the moment of uttering a sentence. The task is facilitated by the fact that an individual is able to compare her standards of perceptual similarity with other speakers and by "the objective pull". The members of a community control the progress of an individual by checking if she singles out acceptable sets of stimuli.

[17] Quine calls it a "pre-established harmony" of perception. See Quine (1996, 1997, 1999).

A set of stimuli can play a role of meaning for some sentences; more precisely, it can play the role of stimulus meaning. Since we are able to compare our perceptual standards, we are able to compare stimulus meanings as well and so it is possible that the whole community of speakers adjust their behaviour in such a way that they associate the same stimulus meanings with the same sentences; i.e. that speakers share meanings of sentences with stimulus meanings.

But Quine is also very explicit about the fact that even if stimulus meaning is comparable, it is not enough to guarantee the general comparability of all the meanings among speakers. "Sameness of stimulus meaning is no appreciable approximation to the general notion of synonymy to which semantics has aspired." (Quine 1970a: 223–224). Why is it so? It is important to bear in mind that Quine talks about stimulus meaning only in relation to observation sentences. Observation sentences are specific in two ways – they prompt assent/dissent in the presence of appropriate stimuli and assent/dissent to such sentences depends on the presence of stimuli *only*.

> Occasion sentences, as against standing sentences, are sentences such as 'Gavagai', 'Red', 'It hurts', 'His face is dirty', which command assent or dissent only *if queried after an appropriate prompting stimulation*. (Quine 1960: 35–36, emphasis added)

> Occasion sentences whose stimulus meanings *vary none under the influence of collateral information* may naturally be called observation sentences, and

> their stimulus meanings may without fear of contradiction be said to do full justice to their meanings. (Quine 1960: 42, emphasis added)

The problem is that, despite the fact that observation sentences play an important role in the initial stages of language learning, they form only a small group from the perspective of the whole language. A vast majority of sentences lacks sufficient observation character for a specification of stimulus meaning because there simply are no stable and specifiable sets of immediately perceived stimuli associated with these sentences.[18] Among other things, this means that the strategy of language learning as presented above is not generally applicable. This is why Quine talks about "primitive level" when he discusses how we acquire observation sentences:

> What we need to notice is that all language learning *at this primitive level* is directed only to the learning of what may be called observation terms, or, more properly, observation sentences. (Quine 1975: 246, emphasis added)

The lack of determinate content influences how an individual proceeds in language learning as she moves to more abstract (less observational) parts of a language.[19] The move from purely observation sentences requires an individual to

[18] The decreasing behavioural evidence can be observed from the perspective of a linguist who tries to come up with a systemic way to achieve a segmentation of sentences. Any attempt to specify a determinate set of analytical hypotheses (i.e. a system of hypotheses about the segmentation of sentences) is underdetermined by the behavioural evidence as well: "...the most notable thing about the analytical hypotheses is that they exceed anything implicit in any native's dispositions to speech behavior" (Quine 1960: 70).

[19] In this respect Quine often talks about non-observation occasion sentences.

use her creativity and to try to communicate with others without being able to specify a determinate set of stimuli which they all associate with non-observation sentences. In fact, there are no determinate sets to be specified and there is no guaranteed, general way for speakers to learn non-observation sentences. "There is guesswork here, and more extravagant guesswork to follow" (Quine 1992: 45).

This is more problematic as it may seem because if non-observation sentences lacks determinate content, then the notion of sharing a meaning (of a sentence) is problematic as well. If we talk about observation sentences, sharing their meanings requires specification and comparison of sets of stimuli which must be perceived at the time of hearing/pronouncing a sentence. This requirement guarantees the possibility that different speakers adjust their behaviour in such a way that they share a language. However, as a consequence of the lack of determinate content of non-observation sentences, the notion of sharing a meaning does not make sense anymore. It is not clear what a speaker has to "accomplish" or acquire to say that she shares a meaning with someone else, i.e. that they associate the same meaning with a sentence.[20] If this is so, then the notion of shared language is problematic as well – if the majority of expressions of a language lacks determinate content, then it is hard to imagine how different speakers could adjust their behaviour to share a language. In other words, if the majority of expressions lacks determinate content, then it makes no sense to say that two speakers associate *the same meanings* with expressions.

[20] In a similar way, the lack of determinate content problematizes the notion of understanding as well. As Quine states: "Standing sentences can be queried too, but the stimulatory situation at the time of querying them will usually have no bearing on the verdict. … I do not know how, in general, in terms of behavioural dispositions, to approximate to the notion of understanding at all, when the sentences understood are standing sentences" (Quine 1975: 249–250).

I believe that Quine was aware that the lack of determinate content problematizes the intuitive notion of shared language because if non-observation sentences lack determinate content then there is no easy explanation of what sharing a language might amount to. Unfortunately, he did not discuss this consequence explicitly and for some reason he turned his attention to the indeterminacy of translation at this point. However, even if he continued discussing the indeterminacy of translation, what is important is the reason why translation is indeterminate according to him – the publicly available language lacks determinate content and so it does not offer enough evidence to settle one correct manual of translation.[21]

An indirect evidence that Quine was aware that the lack of determinate content of non-observation sentences threatens comparability of their meanings (and so the notion of shared language) is that he decided to eliminate the notion of shared language from his explanation of what governs speakers in their language learning. According to Quine, the only clue which encourages a learner that she is "on the same path" with her more experienced fellows is the fluency of communication and this is a standard which does not have to rely on language sharing in any way.[22]

The move from observation sentences to the less observational part of language is also the point at which "the objective pull" loses its grounds. As the behavioural evidence decreases, the corrective interventions of members of a

[21] The point that the lack of determinate content is more important than the indeterminacy of analytical hypotheses is mentioned only seldom. An exception can be found in Føllesdal (2008: 2), who talks about the indeterminacy of translation as a consequence of more fundamental and more important ideas about the public nature of language.

[22] See subchapter 2.4.2 for a longer discussion of this point.

community start to be indeterminate as well. At some point, corrective cues are not able to lead a learner to specify a determinate set of stimuli associated with a particular non-observation sentence because there is no determinate set to be reached. Because of that the corrective behaviour of speakers is not able to sustain uniformity among members of a community and so it is not able to guarantee that all the members of a community share a language.

To sum up, some of Quine's statements about the role of communities in language teaching might indicate that he is a proponent of naïve communitarism. But if we take seriously his assumption that non-observation sentences lack determinate content, we can conclude that he does not believe that two speakers can associate the same meanings with non-observation sentences. As its consequence, it is not clear what sharing a language among members of a linguistic community might amount to.

2.4.1 Linguistic Theories and Actual Practice

If Quine does not believe that two speakers can share a language, then the question arises as to how to understand his statements about natural languages and linguistic communities. An answer to this question can be found in the difference between linguistic theories and actual practice. Surely, finding a theory (and it does not matter if we talk about semantics, grammar or other aspects of natural languages) which is able to represent a language used by some speakers is possible. But we must be very careful in stating what the theory really amounts to.

Quine (1970a) distinguishes between fitting and guiding the linguistic behaviour of speakers. The fact that a theory fits the behaviour of speakers does not mean that those speakers are consciously or unconsciously guided by the theory in their linguistic practice. In fact, as his approach to proxy-functions indicates, we can easily create many candidate semantic theories which fit linguistic dispositions of speakers equally well and Quine seems to endorse this view since his early writings.

However, a linguist does not have to find the "real" theory used by speakers. As Quine believes, such an effort may even turn out to be impossible because of the lack of determinate content:

> [A theory] imputes to the natives an unconscious preference for one system of rules over another, equally unconscious, which is extensionally equivalent to it. Are the unconscious rules the same, even, from one native speaker to the next? (Quine 1970a: 387)[23]

All a linguist needs is to outline *some* theory that fits the behaviour of speakers well enough and the criteria which are important for shaping scientific theories may be far from the criteria which are important for shaping our own linguistic behaviour. A scientific theory should be simple and systematic so that it can be easily understood and applied according to a small number of rules. This may force a linguist to overlook irregularities in the actual practice of particular

[23] Quine addresses this worry specifically to Chomsky, but it can be used generally for any theory which aspires to be the theory used by some speakers.

speakers. In general, a linguist abstracts from the actual practice. The result of this abstraction is an idealization of a language. Quine mentions natural languages (such as English) mostly in the context of linguistic theories being produced and so he actually talks about scientific idealizations. This is, for example, the case of translation, mentioned in subchapter 2.1. The whole Gavagai example is presented from the perspective of a linguist who is responsible for producing *some* systematic manual of translation.

However, he can talk about natural languages as scientific idealizations and at the same time believe that there really are no communities which members share those languages. The distinction between the real linguistic practice and the scientific practice of producing theories is often hard to see in Quine's views but it is of the utmost importance in our current context. The actual linguistic practice does not need to correspond perfectly to the scientific idealizations and speakers can communicate smoothly without being able to tell if they share a language or not.

2.4.2 Outcomes of Language Learning and Flexible Communities

Moreover, nothing in Quine's views depends on the notion of shared language or communities in which members share a language. On the contrary, he often points out that his definitions of semantic phenomena, such as stimulus meaning or synonymy, are first defined with respect to a particular individual and only then does he look for a possibility of defining those terms with respect to communities.

This is a point in which Quine makes an explicit commitment to the approach which takes the notion of idiolects as fundamental:

> *Cognitive equivalence for the individual, however, is the prior notion conceptually*, that is, in respect of criterion. Two occasion sentences are equivalent for him if he is disposed, on every occasion of query, to give them matching verdicts or, on doubtful occasions, no verdict. The *summation over society comes afterward*. (Quine 1979: 4, emphasis added)

Actually, the quote above can be read positively or negatively. The positive (superficial) reading would state that Quine believed that speakers form communities in which they share a language because he explicitly claims that the summation over community is possible (even though it is conceptually posterior). A negative reading focuses on how broad the summation can be. Clearly, Quine believed that some comparison of dispositions between speakers and a subsequent summation over community is possible, but in this quote he talks about occasion sentences, i.e. sentences which are close to the ideal of (comparable) stimulus meaning. This quote gives us no reason to believe that Quine thought that we can decide whether any two sentences are cognitively equivalent. In fact, he is very explicit that this is not possible for non-observation (standing) sentences:

> I feel that the relation of cognitive equivalence is in good shape so far as occasion sentences are concerned. … There remain, of course, the other sentences – the standing sentences. (Quine 1979: 5)[24]

If the summation over community is possible for occasion sentences only, then it is a poor one and it is very far from reaching the extent needed for the whole natural languages (such as English) and communities (such as English speakers) as they are usually understood.

However, this is not a problem. Quine does not need to assume that members of a community share the same language even when he talks about language teaching (where the appeal to sharing seems to be the strongest). The role of community in language acquisition depends on how we define the outcome of language learning and it is a mistake to say about Quine that the outcome of language learning is to learn the shared language. For Quine, there is no final point to be reached and the outcome is less ambitious: it is to produce a language which is applicable in smooth communication with most of the people, most of the time. Quine, with regard to the consequences of the thesis of the indeterminacy of translation, explicitly mentions smooth communication as the only indicator of successful manuals of translation:

> Our radical translator would put his developing manual of translation continually to use, and go on revising it in the light of his successes and

[24] A similar idea is presented in Pursuit of Truth: "Unlike observation sentences, most utterances resist correlation with concurrent stimulations" (Quine 1992: 45).

> failures of communication. The successes consist – to repeat – in successful negotiation and smooth conversation. (Quine 1992: 46–47)

Later on he adds that his discussion of translation also applies to particular languages:

> I have directed my indeterminacy thesis on a radically exotic language for the sake of plausibility, but in principle it applies even to the home language. (Quine 1992: 48)

A few pages later, Quine explains how lexicography can be understood in such a way that it fits the actual needs of speakers. This is possible, if lexicographers focus on what is important for speakers in communication and language learning – on smoothness of communication:

> Lexicography has no need of synonymy, we saw, and it has no need of a sharp distinction between understanding and misunderstanding either. … He [a lexicographer] does what he can, within a limited compass, to adjust the reader's verbal behavior to that of the community as a whole, or of some preferred quarter of it. The adjustment is a matter of degree, and a vague one: a *matter of fluency and effectiveness of dialogue*. (Quine 1992: 59, emphasis added)

In praxis, if smooth communication is the outcome of language learning, then communication can be seen as a never-ending process of language acquisition. During conversations with different people we run across problems and we try to retrace where they originates (if necessary or if we are in a mood for such a venture). This might lead to an adjustment of our dispositions and to an adjustment of our communicational strategies. What is important is that any such adjustment is basically unwarranted and accepted only provisionally until we encounter other problems in communication with other speakers. "What we objectively have is just … an evolving set of dispositions to be prompted by stimulations to assent to or dissent from sentences." (Quine 1960: 38–39)

On the basis of such a view on language learning, the idea of uniform communities is unnecessary and untenable. What we need is a flexible definition of linguistic communities: a linguistic community is a group of people who can smoothly communicate most of the time. How we understand "smoothly" and "most of the time" is a matter of practical decision and it can be adjusted to reflect our particular needs. In one of his latest writings, Quine explicitly endorses flexible communities in response to Stoutland's remark that stimulus meaning in not enough to guarantee the correctness of language use:

> Vagueness of the limits of the intended linguistic community does matter to the notion of observation sentence, and I deal with it by taking the community as a parameter, varying from a broad society to a clutch of specialists according to our purpose. Vagueness is emphatic and deliberate in

> my standard of adequacy of translation: just smooth dialogue and successful negotiation. (Quine 2000: 422)

Someone speaking Scottish English can count as an English speaker for an American tourist in Russia (where only approx. 5% of people can speak English), but it does not need to count as such for an anthropologist studying cultural and linguistic differences on the border of Scotland and Northern England. Such a view of language learning and linguistic communities allows Quine to talk about linguistic communities without being forced to accept that their members have to share a language. No doubt, a community influences an individual in language learning, but the contributions of its particular members do not need to be fully uniform and they do not need to lead an individual to a shared language of community.

2.5 Conclusions

I agree that the way Quine talks about linguistic communities and natural languages may give the impression that he believed that linguistic communities are groups of people who share a language. The role which Quine ascribes to corrective behaviour in the initial phases of language learning may support such an interpretation as well. But if we look at the process of language learning beyond observation sentences, we can conclude that Quine did not believe that

speakers can share a language because of the lack of determinate content of non-observation sentences.

In this respect, Quine is most probably the first proponent of linguistic individualism. The main thesis of linguistic individualism, as I understand the term, is that the shared language of a community is a myth – it is a scientific abstraction, an oversimplification of how natural languages really work and what they look like. The only things we can find are sets of dispositions (idiolects) of particular speakers which are constantly adjusted under the influence of interaction with other speakers without any guarantee of sharing semantic contents. There is no final point to be reached, no shared language to achieve. As the interaction goes on, the dispositions of all the speakers are always evolving.

The thesis has a radical impact on metasemantics and theory of communication. The notion of idiolect is fundamental and anything we can say about semantic phenomena should be relativized to a particular speaker at a particular time. The way to preserve the notion of a linguistic community is to focus on the fluency of communication. Two speakers are members of the same linguistic community if they are able to communicate smoothly. How broad a linguistic community is depends on how liberal we are in defining smooth communication. The boundary of a linguistic community is a parameter that is adjustable according to our current purposes.

At the end, I return to the question posed at the beginning of this chapter. What does Quine mean when he says that language is a social art? In my opinion, Quine does not want to say that language is one piece of art created by many

speakers. It is more a skill – a skill of an individual to "survive" linguistically among other speakers. A skill which is so complicated to master and which rests on such poor foundations that its acquisition must be an art.

3. Do We Share a Language? Communitarism and Its Challenges

We often take the notion of shared language for granted. We talk about English speaking countries or German grammar and at the New Year we resolve to improve our Spanish or to learn yet another language. To see a natural language as something what we can (and do) share with other speakers is very intuitive for lay persons as well as for philosophers.

In this chapter, I use the label 'communitarism' for those philosophical approaches to natural languages which take the notion of shared language as fundamental. For example, Lewis (1969), Brandom (1994, 2000) and Peregrin (2014a, 2014b) argue that a shared language is an outcome of some intra-group processes and Borg (2004) argues that sharing of natural languages is an outcome of innate language modules.[25] In general, proponents of communitarism believe that natural languages are shared in a sense that members of a community share one set of meaningful expressions and syntactic rules and that speakers play an important role in maintaining shared languages by using them in communication.

However, as Davidson (1986, 1994) noted, if we look at how communication works, we can notice that the way how we use sentences or expressions is not always in accordance with their pre-learnt, and so shared, meanings. In other words, the linguistic production of speakers exhibits variability. This variability is caused by the fact that what a speaker means by

[25] Other philosophers, such as Dummett (1986, 1994) or Weiss (2010), use the notion of shared language as an important part of their argumentation.

uttering a sentence depends partially on her general knowledge and beliefs as well as on the broader circumstances of a conversation in which the sentence is uttered.

Many communitarists admit that the notion of shared language alone does not suffice for the explanation of what makes communication successful. Because of that Lewis (1969, 1979) and Brandom (1994, 2000) stick to the notion of scorekeeping in language games and Borg (2004) sticks to the distinction between semantic and pragmatic features of natural language processing.

However, if communication includes processes beyond a simple application of a shared language, the question arises what criteria of individuation should be used for demarcation of shared languages from within the variable linguistic production of speakers. If communitarists believe that speakers within a community share a language and maintain it in time, then it should be possible to track and demarcate the shared language (the shared set of meaningful expressions and syntactic rules) from their linguistic production.

In this chapter, I discuss one of the most influential strategies which communitarists can adopt to demarcate shared languages – the strategy for maintaining coherently the intuitive idea of natural languages as shared and the idea that the linguistic production of speakers exhibits variability. In short, a communitarist can claim that *not all the aspects of the linguistic production of speakers* are relevant for demarcating what the shared language of a community is.[26] One way how to distinguish between relevant and irrelevant aspects of the

[26] In some sense, we can say that the strategy is an allusion to Chomsky (1965) and his distinction between competence and performance because only some performances (actual acts of uttering a sentence) are acceptable as relevant data for demarcation of shared languages. Notice that the competence-performance distinction has been heavily criticized (Labov 1971, 468; Noonan 1999,

linguistic production is to stick to some version of semantic-pragmatic distinction. In this chapter, I discuss the currently predominant version of this distinction – minimal semantics as advocated by Borg (2004, 2012).

The aim of this chapter is to show that the strategy faces serious problems. The problems lie in the fact that meanings of expressions in natural languages are indeterminate, as well as in its two-step model of communication and understanding currently criticized on empirical grounds (Cosentino et al. 2017).

In the first part of this chapter, I will discuss why a coherent view on natural languages should take into consideration the linguistic production of speakers as well as how the variability of the linguistic production challenges the intuitive idea of natural languages as shared. In the second part of this chapter, I will present the strategy mentioned above in detail.

I believe that the intuitive idea of natural languages as shared is so pervasive that it is often accepted without an explicit reflection. Nevertheless, the aim of the chapter is not to argue that we should abandon the notion of shared language, but to point out weaknesses of the strategy for demarcation of shared languages and to open a discussion for future improvement. Any alternative to communitarism currently under discussion brings more problems than solutions and their prospects in providing a coherent view on natural languages are, at least for now, poor.[27]

---

21) for arbitrary preference of some data over other. So if communitarists rely on the restrictions on "performance", then similar objections should apply to them as well.

[27] Semantic holism as a version of individualistic approach based on the notion of idiolects is an alternative (Rapaport 2000, 2003; Pollock 2014). Another alternative could be Ludlow (2014) and his idea of microlanguages. See Drobňák (2018) for a discussion of why Quine (1960) is a proponent of an approach which relies on the notion of idiolects as well.

3.1 A Common Ground

I start from a simple assumption: any philosophical account of natural languages should take into consideration how speakers understand expressions and sentences and, subsequently, how they use them in communication. The linguistic production of speakers, as an outcome of their competence, should be of *some* relevance for any philosophical explanation of natural languages simply because natural languages are those languages which naturally developed within communities of speakers and they are used by those speakers in communication.

On the other hand, it also sounds intuitively acceptable that linguists, semanticists or philosophers of language should abstract from actual linguistic practices of speakers when providing explanations of how natural languages work. There are many reasons for abstracting – including defectiveness and sloppiness of the actual linguistic production of ordinary speakers. Despite that, abstractions cannot be completely arbitrary. There must be some connection between the results of abstractions and the linguistic production of speakers, otherwise there would be no justification that those results fit a particular natural language.[28] In other words, there must be some criteria of individuation for shared languages which tell us how to abstract from the linguistic production of speakers.

Because of that, the idea that linguists, semanticists or philosophers of language abstract from *some aspects* of the linguistic production of speakers

[28] Such a result of an abstraction can be, for example, a semantic model of a language.

sounds more reasonable. The question which aspects are relevant and which should be overlooked is then decided by criteria of individuation. In such case, the actual linguistic production of speakers is considered to be a reliable source of data about natural languages, but we are finical in delimiting which aspects of the linguistic production count as a reliable source.

To sum up, any coherent view on natural languages should somehow reflect upon the actual linguistic practice, otherwise it is not clear what makes it to be about a natural language. In particular, since the actual linguistic production is carried out by speakers in communication, any coherent view on natural languages should be able to give compatible answers to three questions:

a) what natural languages are;
b) what role particular speakers have in maintaining meaningful expressions of natural languages;
c) how communication relates to the previous questions.

3.1.1 The Challenge

The biggest challenge in providing the answers for a), b), and c) is that the traditional and very intuitive idea of language does not fit very well the way how communication works.

The traditional view in semantics or philosophy of language is that a language is a set of meaningful expressions and syntactic rules. Such a view on

language is implied if the principle of compositionality is accepted. Standard approaches that aim at delivering semantic models of languages assume that meanings of words and syntactic rules are sufficient for composing meanings of sentences (understood usually as truth conditions). The lexicon of a language, i.e. the set of meaningful expressions, serves as a stock of building blocks for compositionality. If I learn the meaning of an expression, I can use it (together with other expressions and syntactic rules) to compose basically an infinite number of sentences. The role of expressions as building blocks for compositionality is facilitated by the fact that the majority of expressions have determinate and context-invariant meanings.

Such a view on languages has a very high explanatory potential. First of all, a language that mostly consists of expressions with one determinate and context-invariant meaning is easy to learn and share. If we assume that different speakers can acquire and share the meanings of expressions (or concepts as their mental representations), we can easily answer all three questions above. We can explain what a natural language is by saying that it is the set of meaningful expressions and syntactic rules which is shared by some speakers. By learning the shared language, speakers become competent and maintain the language for next generations. Then, a group of people counts as a linguistic community if and only if almost all its members share the same language. In the same way, we can say that different speakers understand each other because they share the pre-learnt language and their communication is successful because they ascribe the same meanings to the same expressions.

The biggest problem of this approach is that such an explanation of communication does not have much support if we look at how communication works. Davidson (1986, 1994) argues that the way how understanding is reached shows to be untenable that all speakers simply assign the same pre-learnt meanings to the same expressions in communication. Such a view on understanding is untenable, because how we understand a sentence may be influenced by contextual cues present in a conversation. Cues can be intentionally incorporated into conversation by a speaker, or they might be a result of accidental circumstances in which a conversation takes place. As Davidson (1986: 439) demonstrates on Donnellan's (1968) use of the sentence 'There's glory for you', even this sentence can be understood as 'There's a nice knockdown argument for you' if the conditions are right. Because of these conversational shifts in meaning, the linguistic production of speakers exhibits variability.

However, the point I want to emphasize is not about the explanation of communication. As I mentioned above, communitarists often admit that the explanation of successful communication requires more than just a simple application of a shared language. Rather, the point I want to emphasize is about demarcating the boundaries of shared languages. If our linguistic production (the way we use expressions) is a source of data for demarcating shared languages and, at the same time, there is a realm of the linguistic production that exhibits variability, then the variability in data indicates that expressions do not have determinate and context-invariant meanings and so it problematizes the possibility of demarcating shared languages.

The variability of the linguistic production leads Davidson to the conclusion that there is nothing that would correspond to the standard view on languages as sets of meaningful expressions and syntactic rules and we should get rid of it. As he claims "there is no such thing as a language, not if a language is anything like what many philosophers and linguists have supposed" (Davidson 1986: 446) and he believes that "we must give up the idea of a clearly defined shared structure which language-users acquire and then apply to cases" (Davidson 1986: 446). I agree that the argument is valid, but only under the assumption that all the linguistic production is taken to be a relevant source of data about shared languages. If the circumstances of particular conversations influence our linguistic production and all the linguistic production is taken to be a relevant source of data about natural languages, then communitarists are losing the demarcation criterion for what counts as a shared language.

Rejecting the assumption that all the linguistic production is a reliable source of data about natural languages might help communitarists to avoid Davidson's conclusion, but the variability of the linguistic production (as a fact about natural languages and their use) still poses a challenge for them. The challenge for communitarists is to give clear criteria for which aspects of the linguistic production of speakers count as a reliable source of data for demarcating shared languages and which aspects should be considered to be irrelevant. If communitarists want to preserve the notion of shared language, then they have to explain the existence of the variability of the linguistic production along the existence of shared languages.

3.2 Communitarism

One strategy how to avoid the challenge is to buy the bullet and to accept that the variability of the linguistic production shows that meanings of expressions in natural languages are, in some sense, variable. Buying the bullet does not necessarily mean the loss of the notion of shared language. The idea might be that, even though there are several ways how an expression is used in communication, the ways are well-recognized and shared by speakers within a community. The context-invariant meaning can be explicated as a compound of several contextual values and the shared language as a set of expressions with complex and variable meanings and syntactic rules. Such a view on meaning is sometimes labelled 'rich meaning approach' but, as far as I know, this approach to meaning and shared languages has not been spelled out in detail so far.

The biggest problem of this strategy is that it is not clear whether meaning understood in this way can be compositional and so whether creating a semantic model of a language would be possible. Another problem of this approach is that if meanings are complex compounds, acquiring such languages would be much more demanding (and probably almost impossible). Even if this strategy allows communitarists to save the notion of shared language (in the new sense), such a notion of shared language would not be able to play the same explanatory role as the standard notion of shared language was supposed to play – causing new complications and problems that must be solved.

3.2.1 Which Aspects of Linguistic Production?

Another strategy for preserving the intuitive notion of shared language relies on setting a clear boundary between those aspects of linguistic production which are shared by all speakers and those aspects which can vary from speaker to speaker, from conversation to conversation. As I said earlier, the notion of shared language is often taken for granted without an explicit reflection so it is hard to find an explicit proponent of this strategy. However, I believe that the strategy can be naturally linked to Grice's (1957, 1961) distinction between the semantic and pragmatic features of content. I believe that, if asked, many philosophers would stick to an explanation in line with the Gricean distinction between semantics and pragmatics – we share a language with regards to semantics and the variability of the linguistic production is caused by pragmatics.

In particular, a proponent of communitarism can claim that only those aspects of our linguistic production which are relevant for semantic features of content serve as a reliable source of data about shared languages. If communitarists can succeed in demarcating which aspects of our linguistic production corresponds to the semantic features of content, then they basically succeed in avoiding the challenge. Even though this might not be its primary purpose, minimal semantics as advocated by Borg (2004, 2012) can serve as a very good background for accomplishing this task.

According to Borg (2004, 2012), formal semantics should deal with the literal meaning of sentence-types and expressions. More specifically, formal semantics should provide a model of a language which is able to state what each sentence of a language means solely on the basis of syntactic features of sentences and semantic properties of its constituents (particular expressions of a language). Stating this standardly amounts to stating truth conditions of sentences.[29]

> What matters for semantic operations on a formal account just are the (local) syntactic properties of representations. So, on this kind of picture, grasp of meaning would seem to be in principle amenable to a (Turing-style) computational explanation. If, say, we treat grasp of literal linguistic meaning as the canonical derivation of truth conditions for sentences, for example, along the lines of Larson and Segal 1995, then semantic understanding can form part of a genuine language module, for this is clearly a function which is encapsulated and computational. Knowledge of meaning, on this kind of account, consists of knowledge of a proprietary body of information (the lexicon for the language) and knowledge of a set of rules operating only on that information, rules which consist of formal transformations of the data. (Borg 2004: 81)

Knowing the "proprietary body of information" requires knowing how to categorize objects under particular expressions of a language, i.e. knowing which

[29] According to Borg, this allows to incorporate also syntactically triggered context-sensitive expressions. Overt context-sensitive expressions (e.g. demonstratives, indexicals, tensed expressions) may count as such.

expressions are related to which concepts. Notice that, according to Borg, T-sentences maps natural language sentences to "Mentalese" so it makes sense to say that concepts are mental representations of meanings and so categorization of objects under expressions is relevant for semantic processing of sentences. If this is so, then referential aspects of our linguistic production can provide relevant data for demarcating literal meanings, as referential aspects of our linguistic production indicate how a speaker categorizes objects.[30]

Aspects of our linguistic production which are not syntactically encoded are a matter of what can be implied by uttering a sentence and belong to pragmatics. They are irrelevant for semantic meaningfulness of expressions and syntactic processing of sentences. But most importantly, if minimal semantics is adopted as a background theory for communitarism, then we can say that the aspects of our linguistic production which are not syntactically encoded are irrelevant for the demarcation of shared languages.

More generally, there are two aspects of minimal semantics which make this theory appealing for communitarists:

a) Since meaning of a sentence is syntactically encoded, it is possible to determine the literal meaning of a sentence (and so to understand its literal meaning) without any information about the circumstances of a

[30] "It would also fall within the purview of the language faculty to calculate the mental representation of the truth-condition for the natural language sentence 'The cat is on the mat', where what is constructed is a language of thought sentence which exhibits connections to the external world just to the extent that the language of thought expressions out of which it is constructed exhibit such relations (to put it crudely, *since CAT hooks up to cats, and MAT hooks up to mats, the truth conditions for the natural language sentence 'the cat is on the mat' turns on how things stand with some cat and some mat*)." (Borg 2004: 24, emphasis added)

conversation. It can be done solely on the basis of information about the literal meaning of lexical units and the syntactic structure of a sentence and this information is accessible by all speakers in all circumstances simply by hearing a sentence.[31]

b) Since minimal semantics is closely linked to the modular theory of mind,[32] models provided by formal semantics are supposed to be models of a specific linguistic module which is responsible for semantic processing. In general, this module is considered to be innate and this gives us a reason to assume that different speakers process literal meaning of sentences in the same way. In other words, different speakers ascribe the same meanings to the same expressions/sentences.

If a communitarist adopts minimal semantics, she can claim that the aspects of our linguistic production which are related to syntactically encoded truth conditions of sentences serve to demarcate what shared languages are. Since congruence on concepts (categorization of objects) matters for stating truth conditions of sentences, only the referential aspects of our linguistic production are relevant for meaningfulness of particular expressions. A language is then a set of meaningful expressions and syntactic rules with regards to syntactically encoded truth conditions of sentences and referential aspects of the linguistic production of speakers. If we add the assumption that such a language is an outcome of our

[31] "What minimalism specifies is the content a competent language user is guaranteed to be able to recover, given adequate lexical resources." (Borg 2012: 63)

[32] Borg (2004) overtly discusses the modular theory advocated by Fodor (1983, 1998, 2000). Other modular approach can be found in Chomsky (1971, 1975, 1986, 2000).

innate semantic processing module, we can expect that all speakers within a community share a language. This allows communitarists to save the notion of shared language and use it in further explanations. For example, it can be used to state a demarcation criterion for linguistic communities: What makes a group of speakers a linguistic community is the fact that they all share a language in the aforementioned sense – that they all share semantic processing with regards to truth conditions of sentences and they agree on categorization of objects falling under particular expressions.

Modular theory of mind also answers what role particular speakers play in establishing and maintaining a natural language. As syntactic processing is innate, it does not require any special effort. We are all disposed to process sentences syntactically in the same way simply by virtue of being normal human beings. All we need to do is to show our successors which expressions refer to which objects in the world to the extent that they are able to grasp corresponding concepts.[33]

As long as we agree on which objects fall under 'blood', 'hands', 'the room', etc. in the sentence 'The man over there left the room with blood on his hands', we can all (semantically) process and understand the sentence in the same way. Surely, there can be much more implied by uttering the sentence (e.g. that the man is a killer), but minimal semantics allows communitarists to discriminate minimal standards which must be shared by all speakers and it allows

[33] Allowing that the process of "grasping concepts" can be, at least partially, innately driven: "Finally, then, it seems that we might recognize a third way in which to understand what a module is, for we might view a module as a combination of our two previous accounts, so that a cognitive module comprises a proprietary body of information together with a proprietary set of rules or processes operating over that information. Again, both the rules and the representations they operate on are usually thought to be given innately; thus we have a model of a module as an innate and dedicated cognitive processor" (Borg 2004: 76).

communitarists to demarcate natural languages in terms of these minimal shared standards.[34]

3.2.2 Communitarism and Communication

If communitarists adopt minimal semantics, the most natural view on communication may be a two-step model: semantic processing first, pragmatic processing second.[35] When a hearer hears a sentence, she first unconsciously processes it via a semantic module. The result is that she understands what a sentence means (semantic understanding). In the next step, all pragmatic information about a speaker and other circumstances intervenes and a hearer comes up with an interpretation what a speaker might want to imply by uttering this sentence (pragmatic understanding).

The reason why this approach to communication might be appealing for communitarists is that the notion of shared language plays an important role in the explanation of how communication works. According to this proposal, pragmatic processing is possible only with a background semantic processing. To reach a

[34] This is not to say that Borg by itself is a proponent of this view. My only assumption in this paper is that her view can be used to demarcate shared languages in such a way and that such a view might be intuitively appealing for many communitarists.

[35] This is the view held by Borg as well. In general, Borg does not think that formal semantics should be able to explain how communication works, and she does not aim at giving such an explanation. But by setting minimal semantics into a modular theory of mind, she sets the idea of minimal semantics into a broader view on how semantic and pragmatic aspects of understanding relates to each other. And this relation indicates the two step model: "On the one hand, then, semantic knowledge is important and special – without it we would be robbed of the ability to interpret the meanings of words and sentences and thus linguistic communication would be impossible. Yet, from another perspective, semantic knowledge is quite unimportant and peripheral – without all the other kinds of knowledge we have, semantic understanding would be pretty much worthless." (Borg 2004, 263)

pragmatic understanding, which is usually what we care about in communication, a hearer must be "on the same page" with a speaker with regards to literal meanings of sentences. This requires that they both share a language with regards to truth conditions of sentences and categorization of objects under particular expressions. If this is not the case, then the initial data required for pragmatic processing might lead a hearer astray.

To sum up, minimal semantics a) is able to preserve the notion of shared language by delimiting truth-conditional and referential aspects of linguistic production as relevant for demarcation of shared languages and b) relies on the notion of shared language in the explanation of communication works by postulating a congruent semantic understanding as a precondition for pragmatic understanding.

3.3 Problems of the Strategy

There are two problems for communitarists adopting this strategy. The first problem is related to the minimal standards which must be globally shared by all speakers. The second problem is related to the two-step model of communication.

3.3.1 Global Sharing

The requirement of a shared language with regards to the sharing of meanings of particular expressions seems to be too strong to expect. The problem is that meanings of many expressions in natural languages are not fully determinate and context-invariant.

The point about context-invariance can be demonstrated on examples of free pragmatic enrichment. On the basis of what a hearer might know or find out during conversations, her understanding of the verb 'stop' in the sentence 'The policeman stopped the car' may vary depending on whether the policeman was standing at the road, sitting in the car or chasing the car.[36] As Recanati argues, the circumstances of a conversation in such cases influence not only pragmatic aspects of content (i.e. what is implicated by the sentence) but also the meaning of the sentence because each way of stopping the car (by issuing a proper signal, by pushing a brake pedal, or by firing a warning shot) is related to different truth conditions.

In this subchapter, I will put the topic of context-invariance aside and I will focus on the indeterminate nature of meaning in natural languages in detail. The indeterminate nature of meaning in natural languages is often flouted because it is usually taken to be a problem of a small number of expressions only; i.e. vague expressions. Vague expressions share one characteristic feature – objects categorized under them can be ranked on a scale ranging from those which certainly belong to a category to those which certainly do not. Even though there are some uncertain cases, we all have a clear idea of a scale on which we move,

[36] The example is a modification of the example from Recanati (2004, 2010). The same point could be demonstrated by the example of painted leaves as discussed in Travis (1997).

i.e. we have well established and shared criteria of categorization. For example, most of the people would agree that the percentage of head surface without hair or the density of hair is among relevant criteria for 'bald'. Different speakers may diverge on how they actually set thresholds, but it does not necessarily mean that their concepts diverge as well, i.e. that they use different criteria for categorization of objects.

I believe that the indeterminate nature of meaning in natural languages is much more widespread. First of all, it concerns all expressions in natural languages. For any expression of a natural language we can find circumstances in which an application of a criterion of categorization is unclear or undecided. If this is so, then it opens a possibility that different speakers use different criteria of categorization, i.e. that they assign different (though probably similar) concepts with one expression. Notice, that this is not vagueness as it is standardly understood. The problem I discuss here is that we do not know whether some criteria of categorization are relevant, while in the case of vagueness we know what relevant criteria are (we know the scale) but we do not know the exact thresholds.

Let us demonstrate this on the case of the expression 'actor'. Most people would agree that an actor is a person whose profession is acting in films or television and the number of appearances in movies or whether acting is the main source of income are among relevant criteria of categorization.[37] In 2011, Susan Orlean (2011) published an article in The New Yorker about Rin Tin Tin, a movie

[37] At least, this is a definition of the term provided by Oxford dictionary. See: https://en.oxforddictionaries.com/definition/actor

star from twenties. He was a real star of that times – starring in more than 20 movies of Warner Bros., received the Abraham Lincoln humanitarian award, the medal for distinguished service and New York Mayor gave him a key to the city. He was at the peak of his career in 1929 when he received the most votes for the best actor for the Academy Award. The only trouble was that Rin Tin Tin was a German Shepard and the members of the Academy decided that a dog cannot win the prize for the best actor.

Notice that the question whether a dog can be an actor is not a case of vagueness. It is not a matter of how we decide to set the thresholds on a standard scale – Rin Tin Tin was the main character of many movies and he was paid for his acting. He was even famous for real acting, as opposed to merely appearing on the stage (he was able to build the atmosphere of a scene by his facial expressions and so on). The question was whether being a human being is a relevant criterion for categorization of objects under 'actor' and there is no vagueness in that, there is no blur area of problematic cases. And yet, there was no definite answer to this question.

The case can be interpreted in two ways and both of them undermine the idea of minimal shared semantic standards and semantic processing. Firstly, we can say that the meaning of 'actor', or a corresponding concept, has been indeterminate before the voting and only after realizing this indeterminacy different people made it a bit more precise.[38] If this was the case, then it is hard to

[38] Notice that different people made different decisions so if this was the case, then the term became ambiguous. For voters, 'actor' can include dogs; for members of the Academy it can include human beings only.

say what sharing indeterminate meanings/concepts amount to. How can we say that two speakers shared the same meaning of 'actor', if it was not clear what the meaning was? How can we decide whether a concept possessed by one speaker is the same as a concept possessed by another speaker if it is indeterminate which criteria of categorization are constitutive for the concept? A natural response to this worry would be to say that those speakers possessed similar concepts or that their understanding of expressions partially overlapped. However, as far as I know, there is no viable theory of concept/meaning similarity currently under discussion.[39]

Another interpretation of the Rin Tin Tin case is that even before the case it was determinate whether dogs fall under the concept of actor but the question did not arise.[40] In such case, the unusual circumstances forced people to compare their understanding of 'actor' with respect to the categorization of dogs. People who voted for Rin Tin Tin believed that a dog can count as an actor; the members of the Academy were a bit more conservative in their criteria of categorization. The understanding of 'actor' within a linguistic community has been challenged and it uncovered discrepancies between concepts possessed by different speakers and so meanings assigned to the same expression. If this was the case, then clearly the idea that all speakers within a community share semantic understanding on the level of particular expressions, i.e. ascribe the same meanings to the same expressions, does not have much support.

[39] See Fodor and Lepore (1999) for the evaluation of Churchland's (1986, 1993) notion of meaning similarity.
[40] This is certainly an oversimplification. At least, it had never received so much attention.

The idea that expressions of natural languages are indeterminate is not new. Waismann's (1945) idea of open-texture goes in the same direction and Gauker (2017), as a current proponent of the idea of open-texture, overtly argues that it is problematic to simply assume that we all share the same fully determinate concepts (even though, for most of the cases our criteria of categorization deliver overlapping results).[41] Wilson (1982) proposes a thought experiment which aim to demonstrate that our criteria of categorization are often influenced by accidental features of situations in which decisions are made. Ludlow (2014) argues that our criteria of categorization are dynamic, i.e. they can change from conversation to conversation.

The lesson to be learned from the Rin Tin Tin case is that there is never a guarantee that there is a special realm of semantic processing which is shared by all speakers. Even in the case like 'actor', in which we usually assume our congruence in concepts without any doubts, we can find differences among members of a linguistic community.[42]

The most important point, however, is that we can never rule out a possibility that we stumble upon such indeterminacy or differences in concepts/meanings for any expression of a natural language because we can never assess all possible circumstances in which an expression can be used. For any expression, there is a possibility that there are some circumstances which may reveal differences in categorization which have not been noticed before. If this is

[41] However, see Shapiro (2006) for a critical discussion. Shapiro argues that open-texture should count as a kind of vagueness. As far as I can see, the discussion does not have a winner so far.

[42] A similar point leading to a conclusion that decisions about categorization of particular objects depend on accidental features of particular situations, and so it is hard to expect congruence among all speakers, has been raised in Wilson (1982, 2006).

so, then the assumption that there is a special realm of semantics on the level of particular expressions, which is shared by all speakers, seems at least problematic and deserves more attention by any communitarist adopting this strategy.

3.3. 2 Two-step Model of Communication

The second problem of this strategy is the two-step model of communication. The idea that we first semantically process what we hear and only after that we start pragmatic processing has been undermined by recent empirical research. Werning and Cosentino (2017) and Cosentino et al. (2017) show that free pragmatic enrichment intervenes even at the early stages of semantic processing of sentences. More specifically, free pragmatic enrichment helps us to modulate word meanings before we semantically process a sentence. The research was focused on neurological activity of subjects during processing of congruent vs. incongruent noun-verb combinations.[43] More specifically, research teams tested how neurological activity depends on a context of a sentence. A context was presented to subjects as a short story and it served as information necessary for the free pragmatic enrichment. Beside other combinations, researchers tested how neurological activity changes when we combine a congruent context[44] with

[43] Cosentino uses funnel-pour as an example of congruent combination and funnel-hang (a coat) as an example of incongruent combination.

[44] A context inducing congruent noun-verb combination. In the case of funnel-pour, that would be a context of standard procedures in chemical laboratory or in wine cellar.

congruent and incongruent noun-verb combinations and when we combine an incongruent context[45] with congruent and incongruent noun-verb combinations.

According to the two-step strategy, subjects should first semantically process literal meaning of the whole sentence and only then pragmatic processing should take place. If this is so, then we can predict that the context of a sentence should not influence neurological activity related to semantic processing of particular words (or noun-verb combinations). More specifically, a neurological activity at the time of 400 ms after hearing a verb (N400 component)[46] should be the same regardless of a context. However, the research shows that a neurological activity after hearing a verb is significantly affected in those cases in which congruent context is followed by an incongruent noun-verb combination and vice versa (incongruent context and congruent noun-verb combination). In other words, the same noun-verb combination elicits different neurological activity in different contexts. If we assume that the neurological activity (N400 component) corresponds to the contribution of a particular word to a sentence meaning processing, then the difference shows that context influence a word's semantic contribution before the semantic processing of a sentence is finished.

More research has to be done to find out how exactly we should interpret the N400 component,[47] but all interpretations currently under discussion hold that

[45] A context inducing incongruent noun-verb combination. In the case of funnel-hang (a coat), that would be a context of creative work at an art class.

[46] The phenomenon of peaking neurological activity at the time of 400 ms after semantically oriented stimuli has been reported for the first time by Kutas and Hillyard (1980) and confirmed several times after that in different experimental settings (Baggio et al. 2008; Kutas and Hillyard 1984; Kutas et al. 1984; Kutas and Federmeier 2011)

[47] See Baggio and Hagoort (2011); Brouwer and Hoeks (2013); Hagoort et al. (2009) for discussion.

it represents the contribution of particular expressions to the semantic processing of a sentence. Differences in neurological activity varying in accordance with different contexts indicate that pragmatic processing takes place even before semantic processing of a sentence is finished and so it undermines the two-step model of understanding and communication. But if the two-step model of communication is undermined, then the idea of a shared language based on minimal shared standards is problematized as well. There seems to be no special realm of semantic processing on the level of particular expressions which would be shared by all speakers.

3.4 Conclusions

In this chapter, I discussed one strategy which communitarists can adopt for maintaining coherently the idea of a natural language as shared and the idea that the linguistic production of speakers exhibits variability in communication. According to the strategy, only some aspects of the linguistic production of speakers are relevant for demarcating a shared language. This strategy can be naturally supported by some version of semantic-pragmatic distinction, if semantic features of content are considered to be shared by all speakers. The first problem of this strategy is that meanings of expressions in natural languages are indeterminate and so it is hard to say what sharing meanings of expressions, and so sharing a language, might amount to. The second problem is its two step model

of communication and understanding which has been currently criticized on empirical grounds.

4. Meaning-constitutive Inferences

Despite the progress in making inferentialism more rigorous, accomplished thanks to the work of Robert Brandom (1994, 2000) and others,[48] inferentialism still faces many objections. One of the traditional objections focuses on the analytic-synthetic distinction. According to inferentialism, meaning depends on inferences held valid by speakers. Clearly, so the objection goes, not all the inferences we make can be meaning-constitutive and therefore inferentialism has to include a satisfactory version of the analytic-synthetic distinction. Since this is a Sisyphean task, the reputation of inferentialism seems to be corrupted.

A promising attempt to answer the objection can be found in Peregrin (2014b).[49] Even though answering the objection is not his main objective, Peregrin argues that meaning is a matter of inferential rules and only the subset of all the valid inferences for which there is a *widely shared corrective behaviour* among members of some community corresponds to rules (and is therefore meaning-constitutive). Unfortunately, Peregrin does not discuss what counts as "widely shared". This opens a way to a possible objection if we make use of Peregrin's proposal in the metasemantic debate on meaning constitutiveness. Someone can claim that the criterion of "widely shared corrective behaviour" may

[48] For more references on progress in inferentialism see e.g.: Boghossian (2003, 2012), Peregrin (2006, 2010, 2012), Shapiro (2004).

[49] Brandom (2007) offers a response to this objection, but Fodor and Lepore find it unsatisfactory. He builds on Sellars's (1949: 296) idea that "conceptual connections are just the lawful ones" (Brandom 2007: 661). However, this is a weak response as it leads to a consequence that if speakers are wrong about laws, then the words they use mean something else as what the speakers intend them to mean. The view that our words can mean something else as what we intend them to mean is highly controversial. For criticism of such a view, although in a different context, see Schwarz (2013).

be an interesting theoretical proposal, but it is excessively vague and therefore it is of no use.

In this chapter, I argue for the empirical plausibility of Peregrin's proposal. The aim of the chapter is to show that we can find examples of meaning-constitutive linguistic action which take place in specific communication situations. The idea is supported by examples of meaning modulation from Ludlow (2014). Meaning modulation is first and foremost a phenomenon which can be observed in communication. Speakers often change, adjust and discuss meanings of the words they use for various purposes. In general, we can understand meaning modulation as a tool which facilitates successful communication by deciding open questions about a meaning of a word, by making some of the features of a meaning more explicit or by changing some of the features of a meaning. A paradigmatic example of a meaning modulation is the discussion of whether Pluto should be a planet. It can be seen – from the semantic point of view – as a discussion and clarification of the meaning of the word 'planet'.

What I find interesting about meaning modulations is the way how they can be (and often are) used within communities. Some of the modulations not only facilitate successful communication, but also serve to settle precedents which are subsequently adopted and followed by other members of a linguistic community. They are part of more general social mechanisms which operate on the level of whole communities and which constitute new meanings.

If Peregrin is right and meaning is determined by inferential rules, then situations of meaning modulation should support his criterion of widely shared corrective behaviour: we should be able to find a specific type of modulations in which an outcome of the modulations is publicly available and modulated in such a way that it has a potential to establish a widely shared corrective behaviour. I believe that binding modulations – a specific type of modulations – satisfy this condition.

4.1 Preliminaries

The objection mentioned above was explicitly formulated by Fodor and Lepore (2001, 2007), who attempt to criticize several aspects of inferentialism. Among others, they argue that inferentialism has a problem stating which inferences are meaning-constitutive.

According to them, there are many inferences which are actually made but which are not/should not be semantically relevant. If inferentialism is a doctrine that meaning is an inferential role, i.e. a set of inferences in which an expression plays a role, an inferentialist needs to delineate clear boundaries of meaning-constitutive inferences. Fodor and Lepore believe that this means that an inferentialist has to revive the well-known analytic-synthetic distinction to distinguish between meaning-constitutive and "utterly contingent" inferences. As Quine (1951) persuasively showed, this seems to be a task doomed to failure.

> We're also not clear what Brandom thinks about the status of utterly contingent inferences like "If it's a plant in my backyard and it's taller than 6 feet, then it's a tree". He does apparently endorse the idea that "[the concept-constitutive inferences] must include … those that are materially [sic] correct" (MIE, p. 657). But what he gives as examples are two he borrows from Sellars: "A is to the East of B" ⊢ "B is to the West of A" and "Lightning is seen" ⊢ "Thunder will be heard soon". We find this puzzling since the first of these strikes us as arguably conceptually necessary (whatever that means) and the second strikes as arguably nomologically necessary (whatever that means). So even if we granted that both are concept-constitutive, we would still want to know whether clear cases of purely contingent hypotheticals are too; and, if they aren't, how Brandom proposes to do without an analytic/synthetic distinction. (Fodor and Lepore 2007: 680–681).

4.1.1 Inferentialism and Inferential Role Semantics

It is important to understand the difference between Inferential Role Semantics (IRS) and inferentialism as advocated by Peregrin before responding to the objection.[50] According to IRS as proposed by Boghossian (1993, 2012) the

[50] Despite the fact that Fodor and Lepore address their criticism to Brandom, I will mention his work only to a very limited extent in this paper. The reason is exegetical. Fodor and Lepore present Brandom's views in a way that more or less fits IRS (Boghossian's approach). Peregrin argues that this is a misinterpretation of Brandom and builds a response to the objections on what he sees as a more "Brandomian view". Instead of entering an exegetical discussion, I will talk about Boghossian's inferentialism (IRS) as a target of criticism and Peregrin's inferentialism as a response to the objection.

meaning of an expression can be understood as its inferential role. The inferential role is then explained as a set of all the valid inferences in which the expression takes part. Therefore, to understand a meaning of a sentence is to know which other sentences are inferentially connected to the sentence. If the inferential role is understood as a set of *all* the valid inferences related to a sentence and meaning is an inferential role, then all the inferences should be meaning-constitutive – and so such a view is problematized by the objection mentioned above.

Such a view is problematic for more reasons. The approach of IRS is individualistic in nature – inferences which are part of the inferential role of an expression are determined by the dispositions of a *particular speaker* and her ability to distinguish valid and invalid inferences. As Boghossian puts it, it is determined by an ability to "*infer* from S1 to S2, but not to S3" (Boghossian 1993: 73). Such an approach opens the way once more for the objection mentioned above. If the inferential role depends on the practices of particular speakers, it is not clear how to delineate the boundary between meaning-constitutive and non-constitutive inferences. In particular contexts, some inferences which we would normally be inclined to call meaning-constitutive may be less important (e.g. for successful communication) than some contingent inferences. More importantly, if an inferential role of a sentence is the set of *all* the valid inferences in which a sentence appears, then different speakers ascribe (slightly) different meanings to the same sentence.[51] Which inferences a speaker includes in the inferential role of

[51] It is an open question if small differences in inferential roles are acceptable. Fodor and Lepore build their objection on the assumption of fully shared meanings because they do not see any viable similarity-based alternative. See Fodor and Lepore (1999) for their discussion of meaning similarity. I believe that the assumption of fully shared meanings is problematic because it does

a sentence depends on his personal experience, and this is a highly subjective factor.

On the other hand, inferentialism as advocated by Peregrin is in some sense independent of the abilities of particular speakers. As Peregrin puts it: "Language is essentially public, and as such it cannot rest on private associations" (Peregrin 2014b: 45). Meaning is established in the social interactions of many speakers. Additionally, meaning persists within a community only through the existing normative attitudes of speakers – speakers hold some inferences to be correct and by their corrective behaviour force others to conform. If someone ascribes a set of inferences to a sentence which is not in accordance with the publicly established meaning, then she is just wrong and misunderstands the expression (and so she is a legitimate target of criticism).

It must be emphasized that Peregrin's version of inferentialism is not completely independent of the abilities and practices of speakers. I agree that the meaning of a word is independent of the inferential practices of each particular speaker – I cannot change what a word means within a linguistic community solely by changing my own inferential practices. But the meaning of an expression still depends on what the majority of speakers (a minority with high semantic authority, maybe)[52] holds and protects as correct – i.e. it depends on the actual practices of *many* speakers.[53]

---

not correspond to the actual linguistic practice of speakers; regardless of the fact whether there is any alternative. However, this is not the topic of this paper and I postpone the discussion for another occasion.

[52] For the sake of simplicity, I will talk about a majority of speakers from this point onward. But in many contexts (e.g. in the case of scientific terminology), we cannot expect that a majority of speakers really knows all the correct inferences. Semantic authority plays a significant role in

Of course, even in this "communal" setting, the sets of all the valid inferences related to particular sentences by individual speakers can vary. So how can Peregrin avoid the objection and distinguish meaning-constitutive and non-constitutive inferences? According to Peregrin, inferential roles should be understood as sets of *inferential rules* or, in some sense, as sets of inferences which correspond to inferential rules. An example of inferential rules can be 'X is a dog ⊢ X is an animal' or 'X is a dog ⊢ X is not a cat'. Peregrin also accepts inferential rules linking a sentence to some extralinguistic factors, which can have the form 'X is a dog ⊢ …' in which the three dots indicate some action that is inferable from the sentence, e.g. not irritate X. The inferential role of the sentence 'Laika is a dog' includes the inferences 'Laika is a dog ⊢ Laika is an animal'; 'Laika is a dog ⊢ Laika is not a cat'; 'Laika is a dog ⊢ …' (not irritate Laika). But the sentence also appears in "utterly contingent" inferences such as 'Laika is a dog ⊢ Laika cannot enter John's apartment' or 'Laika is a dog ⊢ Laika can be off-leash in many areas of Central Park'.

Peregrin's key to deciding which inferences correspond to inferential rules (i.e. are meaning-constitutive) and which are "utterly contingent" lies in the widely shared corrective behaviour of speakers. Corrective behaviour is any kind of behaviour by which speakers respond to the language use of other speakers. This includes positive as well as negative reactions – rewards in the case of

---

language distribution and preservation and it has to be taken into account. In fact, we can understand 'majority of speakers' as a group of speakers with higher semantic authority or we can simply talk about a majority of speakers whose opinion is semantically relevant.

[53] A similar point was emphasized by Koreň (2017).

correct inferences and warnings and punishments in the case of incorrect inferences.

> There is an inferential rule in force for a given language if the speakers of the language tend to see some inferences that violate the rule as incorrect. (Peregrin 2014b: 58).

and elsewhere:

> And what I call a propriety, or an (implicit) rule, grows out of such attitudes resonating throughout the surrounding society. (Peregrin 2014b: 10).

In the second quote Peregrin talks about attitudes, but the attitudes of speakers matter only because they can be expressed behaviourally via corrective behaviour. If I tend to see some inferences as valid, then I tend to correct speakers who violate them. What is even more important is the phrase 'resonating throughout the surrounding society'. While inferences such as 'Laika is a dog ⊢ Laika is an animal' are *publicly well known* and widespread and there is an established *widely shared corrective behaviour* of speakers related to such inferences, inferences such as 'Laika is a dog ⊢ Laika cannot enter John's apartment' depend on the knowledge of particular speakers and so the (relevant) majority of speakers is not able to evaluate their validity. If the speakers are not able to evaluate the validity of such inferences, then they are not able to use corrective practices either and such inferences cannot be meaning-constitutive.

To sum up, we can decide which inferences are meaning-constitutive (i.e. correspond to inferential rules) by evaluating for which inferences there is a widely shared corrective behaviour among the members of a community. In some sense, Peregrin's proposal serves as a criterion of meaning-constitutiveness.

4.2 Meaning Modulations

The question is if we can find something that corresponds to the inferential roles as proposed by Peregrin at the level of natural languages and linguistic communities: If there are some inferences which are "widely shared" or if Peregrin's proposal is an unreasonable abstraction. I believe that we can find examples of meaning modulations which show that there are social mechanisms important for establishing of new meanings at the level of whole communities. But before I focus on the empirical plausibility of the criterion in more detail, I will present the topic of meaning modulations in general and briefly sketch its relation to inferentialism. Let us have a look at the conversation from the TV series The Apprentice in which Donald Trump (and his aide Caroline) discuss an incident that involved a contestant (Ivana) which happened while she was dealing with a given task – to sell a candy bar:

01	Trump:	Ivana. You flashed a group of people.
02	Ivana:	Look (…) This…
03	Trump:	No, no, no. Did that happen?

04	Ivana:	It happened? But it happened for a reason.
05	Trump:	Why?
06	Ivana:	Because I knew (…) Okay we had gone through a lot of product (…) We only had…
07	Trump:	What does flash mean? You ripped down your pants? What does that mean?
08	Ivana:	I was wearing (…) I was wearing a bikini (…) and (…) and let's not blow this out of proportion. I was wearing bikini shorts.
09	Caroline:	We haven't said anything yet so relax.
10	Ivana:	More: I know. I know. I'm just really defensive about this because…
11	Trump:	Go ahead I'd like to hear that.
12	Ivana:	Um.
13	Trump:	But you did flash.
14	Ivana:	I did. But it was a gimmick. It was a gimmick, just like (…)[54]

Ludlow (2014) uses the example to show how we – more or less implicitly – modulate/litigate meanings within conversations. In this particular case, the word 'flash' has been questioned. In lines 02 and 04 Ivana accepts Trump's accusation

[54] The example first appeared in Sidnell (2010), who used it to show how we use the communicational tool of repair – how we go back in conversation to deal with troubles in understanding. The original transcript conventions used by Sidnell are not important in this context and I decided to use a much simpler transcript: '(…)' indicates a pause made by a speaker and '…' indicates interruption of the speaker by another speaker.

of flashing with slight hesitation. Probably, she hesitates for more reasons but as the conversation shows later on, she does not agree that what she did is an evident case of flashing. In line 07 Trump indicates that he is not sure about the meaning of the word (despite the fact that he introduced it into the conversation) and Ivana tries to cash in on Trump's doubts: in line 08 she indicates that flashing should not apply in cases in which someone is wearing a bikini (and so she discusses the boundaries of the meaning of 'flash'). However, Trump does not accept her modulation and forces Ivana to admit that her behaviour was clearly a case of flashing (line 13). Ivana finally defers to Trump and admits that she flashed (14).

Situations like this are interesting for metasemantics in several ways. Most importantly, such situations are quite common and, as Ludlow argues, they should show that meanings are in general underdetermined and meaning modulations serve to specify the meanings for particular conversations and speakers. In other words, the shared language of community is a myth. There are only microlanguages which are created and modulated on the fly and very often include only the speakers who are present, without any impact on other speakers.[55] In the Trump example, it does not matter if there is a correct meaning of 'flash'. It may even happen that an act counts as flashing only if the exposed body is naked and so Ivana did not flash. But Ivana deferred to Trump's

[55] Even though Ludlow focuses on different phenomena, he basically follows Davidson (1986, 1994) in his conclusion about shared language as a myth.

understanding of the word and her acceptance settles what the word means within their conversation.[56]

Examples like this can easily be "translated" in the inferentialist's terms. We can say that speakers discuss or disagree on the validity of some inferences. In this particular case, it can be an intralinguistic inference 'You flash ⊢ You are completely naked'. Obviously, Trump does not accept the inference, but Ivana would be happy to accept it. An advantage of inferentialism is that it can find meaning litigations in even less obvious circumstances. Let us have a look at the conversation from the TV series Gilmore Girls where Lorelai and her mother Emily dispute whether the offer of a lunch is still on if someone changes their previous plans:

01	Emily:	Stop being so dramatic. I just showed up for lunch…
02	Lorelai:	What do you mean you showed up for lunch?
03	Emily:	Our lunch, at 1:00. You, me, Rory – the three of us. We're having lunch, aren't we?
04	Lorelai:	I didn't think so.
05	Emily:	You didn't?
06	Lorelai:	Well, no, but (…)
07	Emily:	When you invited your father and me for the weekend, you said it included a lunch with you and Rory.

[56] It would be interesting to look at how rational the game of giving and asking for reasons is if understood in terms of meaning modulations. As far as I can see, Ivana did not defer to Trump because he was right or because he offered rational reasons for why her behaviour counts as flashing. She deferred because he was an authority in general – it was his show; he was a judge and her prospective employer. However, this is not the aim of this paper.

08	Lorelai:	Well, yes, I know, but that was before you left.
09	Emily:	What does my leaving have to do with anything?
10	Lorelai:	Well, when you left, you weren't here anymore. You were gone, so we just assumed lunch was...
11	Emily:	Where's Rory?
12	Lorelai:	Okay, see, you left, so (…)
13	Emily:	She's not here, is she?
14	Lorelai:	No.
15	Emily:	Didn't she know about the lunch?
16	Lorelai:	Yes, mom, she knew about the lunch, but you (…) so we (…) and she (…) I'll call her.
17	Emily:	I'll wait.

From the inferentialist's perspective, the conversation can be reconstructed as a dispute over the validity of the inference 'You cancel your previous plans ⊢ You cancel the rest of the plans as well'. The validity of the inference is proclaimed in line 08 by Lorelai and challenged by Emily in the next line 09. Emily ignores Lorelai's repeated appeal to accept it and Lorelai finally defers to Emily in line 16. The example is clearly a case of meaning modulation/litigation from the perspective of IRS. According to IRS, meaning depends on all the inferences held valid by particular speakers and the validity of an inference is in question here, therefore we can conclude that the meaning is in question. What is more, we can conclude that Lorelai has changed/adjusted her understanding of the sentences

'You cancel your previous plans' and 'You cancel the rest of the plans as well' during the conversation.

4.2.1 Meaning-constitutive Modulations

However, the situation is less obvious from the perspective of Peregrin's inferentialism. Not all inferences are meaning-constitutive, i.e. not all inferences are maintained and reinforced by the widely shared corrective behaviour of a community of speakers. In the same manner, not all modulations can be meaning-constitutive as well. If we want to show that Peregrin's criterion presented earlier is empirically plausible, then we should be able to find litigations/modulations which establish meaning-constitutive inferences. In short, we should be able to discern meaning-constitutive modulations.[57]

The modulations presented in the previous examples are made within small groups of people (the conversation between Trump and Ivana is followed by a small group of contestants and judges, the conversation between Emily and

[57] It is generally accepted in the philosophy of language that meaning change is a long-term, unconscious process. If there are any changes in meanings, they are usually implicitly adopted by speakers in the same way as most of the expressions of a language are learnt. Such a view is typical of Wittgenstein (1953), but also of Peregrin (2014b) and discussed in more detail in Peregrin (2014a). A similar view on meaning change and acquisition, discussed in the context of deciding signalling systems, can also be found in Lewis (1969: 129). In the following subchapters, I focus on examples of explicit meaning modulations and intentional acceptance of their results. By doing so, I do not intend to claim that this is the only way in which meaning can be adopted by speakers and become widely shared. I focus on explicit examples because the social mechanisms which are applied in the distribution and adoption of a new meaning in such cases are much more evident and so easier to document and analyse. I even think that both views are partially compatible. I can imagine a situation in which a meaning is settled in an explicit modulation, but after some time the modulation is forgotten and the meaning is adopted implicitly by new speakers.

Lorelai is private) and there is no indication that these modulations should be applied globally as a precedent for other speakers.[58] On the contrary, it is likely that even Ivana and Lorelai will not follow the results of those modulations in future and their deference is only pretended. Since the modulations that were presented did not establish widely shared corrective behaviour, these modulations cannot be meaning-constitutive. Of course, in some cases, similar modulations which take place in personal communication may play an important role in the concept formation of particular speakers, but they are not important from the perspective of entire linguistic communities.

Now let us have a look at a different example. In 2000, Hayden Planetarium demoted Pluto from the status of a planet in their newly opened exhibition.[59] The decision was unusual at that time and it triggered a wave of criticism. One year later, the New York Times published a front page article called "Pluto Not a Planet? Only in New York", in which the author calls the decision "unilateral" and cited several astronomers who criticized the head of the planetarium, Dr. Neil de Grasse Tyson. The article started a "witch hunt" – Dr. de Grasse Tyson received many letters and emails from ordinary people demanding an explanation and renouncement of his view. But the article triggered an academic debate about the definition of 'planet' as well, and as the debate very soon showed, there were no clear criteria for calling an astronomical object a 'planet'. The International Astronomical Union therefore decided to redefine the

[58] Both conversations are from TV series and they both have been seen by millions of people. However, this does not change the main aim of those examples – to show that there is an everyday phenomenon which *is* usually private.

[59] For a longer overview of the case see Weintraub (2007).

term and the new definition did not apply to Pluto anymore: Pluto was officially relegated and pronounced to be a "dwarf planet" in 2006. Despite the fact that this decision raised a new wave of discussions, after 10 years we can say that it is generally accepted by the vast majority of astronomers, as well as non-experts.

When Ludlow presents the examples of Trump and Pluto, he admits that there is a difference – namely in the explicitness of the modulation. In the case of Pluto, astronomers explicitly discussed the meaning of the term 'planet', while in the case of Trump's conversation with Ivana, the litigation over the meaning of 'flash' was to a large extent implicit. As Ludlow puts it, in cases like Pluto "we are consciously aware of disputes about word meaning" (Ludlow 2014: 39). As far as I can see, we can identify more differences and all of them are surprisingly well suited to a delineation of the class of meaning-constitutive modulations. The differences lie in:

A) the intentions of the speakers
B) the number of participants in a modulation
C) information flow and its general accessibility

A) *The intentions of the speakers*

Even though Lorelai and Emily were engaged in modulation, they do not have any reason to look for the most acceptable modulation. Emily wants to have a lunch with her daughter and granddaughter and her position in the litigation follows from this aim. Lorelai defers to Emily's modulation because she knows she has no

chance of convincing her. Neither the intentions of Emily nor the intentions of Lorelai are directed towards the most plausible solution. In fact, it does not matter if there is any plausible modulation; even if there is, it would most probably be ignored. On the other hand, in the case of Pluto, the members of the International Astronomical Union try to find an acceptable modulation – acceptable with regard to the future use of the term within the whole community and with regard to possible future discoveries. In fact, in 2006 there was at least one known object of a size similar to Pluto and potentially there are more such objects in our solar system. The decision that Pluto is a planet would therefore lead to ad hoc decisions about the status of objects in our solar system or to a possible extreme increase in the number of planets. The declassification of Pluto is therefore a result of a reasonable debate looking for plausible solutions[60] for the whole astronomical community, and this was part of the intentions of the committee which was responsible for a redefinition.

In general, we can distinguish two types of modulations on the basis of the intentions of the speakers. On the one hand we have modulations which are intended to serve personal aims, with no intention of attaining a plausible consensus with other speakers. On the other hand, we have modulations which aim at plausible solutions with regard to generally acceptable objectives. We can

[60] I admit that the talk about intentions and the most plausible solutions may be somewhat loose. A reformulation of Dennett's idea of an ideal agent could be used to make the talk rather more rigorous. Dennett (1971) suggests that we can predict someone's behaviour by treating her as an ideally rational agent who uses the best means to attain her aims. Similarly, we can define the most plausible modulation as the one which best suits the aims of the community, where aims are a result of general consensus. In the case of Pluto the decision to declassify it was the most plausible solution with regard to more aims – it avoids a possible extreme enlargement of the number of planets and it allows a more rigorous definition of a 'planet' and a more accurate classification of objects in our solar system in general.

call the first kind ad hoc modulations and the second kind binding modulations. From the perspective of Peregrin's inferentialism, we can say that the ad hoc modulations are not meaning-constitutive, while the binding modulations are meaning-constitutive – only binding modulations are full-fledged *meaning* modulations.

But the intentions of speakers are important for the distinction between modulations only because they lead to a difference in the expected consequences in the behaviour of the speakers. Since the litigants in ad hoc modulations follow particular personal aims, we can expect that even a speaker who enforces a modulation will not be consistent in the use of the expression when compared to her past and future conversations. As a result of achieving her aim, a speaker has no reason to follow the modulation any more. Moreover, a speaker may not follow the modulation in the context of her different aims. Since other speakers do not expect that a speaker will follow the modulation, they do not have any reason to adjust and apply their own corrective behaviour so as to be in accordance with the modulation in future conversations as well.

On the other hand, the reasons which led to the decision about Pluto are a result of a debate with regard to generally acceptable objectives. Binding modulations are intended from the outset to settle a *widely shared* consensus followed by a majority of a linguistic community in the future and this means that some individuals have to adapt from time to time. But since the outcome of binding modulations is supposed to be generally acceptable, we have good reason

to assume that most speakers will systematically follow the modulation in future conversations, regardless of their initial position in meaning litigation.

B) *The number of participants in a modulation*

Another notable difference between ad hoc and binding modulations lies in the number of speakers who participate in the modulations. Ad hoc modulations are usually incidental and appear in small groups of people, even in one-to-one conversations very often. On the other hand, binding modulations are usually open to all the speakers of a relevant linguistic community. In the case of Pluto, a part of the astronomical community decided which modulation would be in use, but the discussion was open to non-experts as well. Even small children sent letters to Dr. de Grasse Tyson. It does not matter whether their opinion was taken into consideration or not. What is most important is that they took part in the litigation and by doing so they designated themselves as members of a relevant community to which the litigation – and its result – applies. This is an important point when compared to ad hoc modulations. If Lorelai refuses to follow the modulation proposed by Emily, her status as a member of any linguistic community will not be harmed in any sense regardless of the fact that she took part in the litigation. But someone's refusal to follow the decision about Pluto can be seen as a reason for the enforcement of corrective practices and, in an extreme case, a reason for her detachment from a linguistic community.

Moreover, the example of Pluto is a rather specific binding modulation. The term 'planet' belongs almost exclusively to astronomy and so astronomers

have some semantic authority in litigations. This is why the opinion of non-experts was not taken seriously. But there are many examples where the authority is not so clear and the role of "vox populi" is much bigger. This is the case of words such as 'marriage' or 'person'. These words became the centre of attention as they appeared in the press, at academic conferences, and in courtrooms. But the question whether an unborn child is a person is not only a legal, medical, or religious matter. It is, first and foremost, a social matter. The results of these modulations will directly influence the everyday lives of many people and therefore a wide public discussion plays an important role in the final decision.

C) *Information flow and its general accessibility*

Since 2001, the New York Times has published more than twenty articles and commentaries on the topic of Pluto's demotion, most of them written by the journalists Kenneth Chang (KC) and Dennis Overbye (DO).

This is only a short list:

Jan. 29, 2002 "Planet or No, It's On to Pluto" (KC)

Jul. 30, 2005 "Planet or Not, Pluto Now Has Far-Out Rival" (KC/DO)

Oct. 4, 2005 "9 Planet? 12? What's a Planet, Anyway? (DO)

Feb. 2, 2006 "Icy Ball Larger Than Pluto. So, Is It a Planet?" (KC)

Aug. 16, 2006 "For Now, Pluto Holds Its Place in Solar System" (DO)

Aug. 22, 2006 "Pluto Seems Poised to Lose Its Planet Status" (DO)

Aug. 24, 2006 "Pluto Is Demoted to Being a Dwarf Planet" (DO)

Aug. 25, 2006 "Vote Makes It Official: Pluto Isn't What Is Used to Be" (DO)

Aug. 25, 2006 “And Now There Are Eight” (Editorial)

Sep. 1, 2006 “Debate Lingers Over Definition for a Planet” (KC)

Dec. 24, 2006 “Dwarf Planet” (DO)

Jun. 12, 2008 “Not a Planet, but a Plutoid” (KC)

Jan. 12, 2009 “How Many Planets Do You Want in the Solar System?” (KC).[61]

The focus here is on the New York Times because it is one of the most influential newspapers in the world, but it is basically arbitrary. We can find a similar list of articles about Pluto in practically any newspaper. The interest of journalists in the topic caused an extensive information flow which ensured that the information about the current status of Pluto (and so about the current state of the meaning of ‘planet’) was distributed among the members of the relevant linguistic community. This is hardly an accidental feature. Of course, even an ad hoc modulation *can* exceptionally become the centre of attention. But this fact does not change the main point: an extensive information flow is an important component of meaning-constitutive modulations because it creates favourable conditions for a distribution of the new meaning within a relevant linguistic community. A modulation cannot become widely shared if it is not generally accessible by a majority of speakers.

When combined, the three points related to binding modulations (the intentions of speakers to follow a modulation, a large number of participants in a

[61] It is worth noticing that the opinions of the New York Times journalists changed radically as the discussion proceeded. The journalists adopted the view of the International Astronomical Union without much hesitation, despite their initial criticism.

modulation, and an extensive information flow) constitute ideal conditions for their results to become widely shared and so to establish a widely shared corrective behaviour with regard to a particular set of inferences. Of course, the fulfilment of these conditions does not necessarily guarantee that the new meaning will be adopted and we can easily find borderline situations. This is, for example, the case of the word 'polyarchy', promoted by Robert Dahl within the field of political science. Dahl (1956, 1971, 1984) argued that the contemporary political system in the USA is not democracy, but polyarchy. Democracy is a system in which all the citizens are considered to be equal in political decisions, while in polyarchy control over governmental decisions is constitutionally vested in elected officials. While his distinction was well known, globally discussed, and later on generally accepted in the field, the word 'polyarchy' has never replaced the word 'democracy' within the "linguistic community of political scientists". It is hard to say why this was the case. A possible explanation might be that there was no need to start using the new word because political scientists in 1956 knew very well that 'democracy' did not mean anymore what it used to mean in Ancient Greece. The meaning of the word 'democracy' has changed with emerging modern republics and so there was no need to adopt 'polyarchy'.

However, even if 'polyarchy' is an example of an unsuccessful modulation, I do not think that the existence of borderline cases causes any problems in our current context. What is sufficient for the purpose of supporting Peregrin's criterion is that there are at least some examples in which binding modulations were successfully adopted by a community. The existence of such

examples shows that there are general mechanisms for establishing new meanings on the level of whole communities – even though they might fail from time to time. Mechanisms which are related to binding modulations are exactly those mechanisms which guarantee that there is a widely shared corrective behaviour for particular inferences and so the inferences correspond to inferential rules.

4.3 Conclusions

The conclusion of this chapter may seem rather subtle: an outcome of a meaning modulation can hardly become widely shared if a majority of speakers does not know about the modulation, if the speakers do not take part in it, and if the speakers do not intend to follow it. But the mere fact that there are modulations the outcomes of which are widely shared has interesting consequences for the discussion of meaning-constitutive inferences. It shows that Peregrin's criterion of meaning-constitutiveness can be empirically supported – that there are social mechanisms thanks to which a set of inferences corresponding to inferential rules can become widely shared.

I agree that this is not exactly what Fodor and Lepore had in mind when they discussed the analytic-synthetic distinction. Peregrin is not able to give fixed and finite lists of analytic and synthetic inferences. But, actually, this can be seen as an advantage. Peregrin can get rid of the analytic-synthetic distinction in its traditional (problematic) form. His criterion does not depend on any "intrinsic" semantic properties of sentences/words and so it is not circular in defining

semantic properties and analyticity.[62] On the contrary, the criterion for meaning-constitutiveness based on the notion of widely shared corrective behaviour follows the dynamics of natural languages and this is a feature worth keeping.

[62] In other words, it does not define semantic properties as those which correspond to analytic inferences and analytic inferences as those which are part of meaning.

5. Inferentialism and Context

The idea that the content of an utterance depends on context is widely recognized in the philosophy of language. However, exactly how context influences content is the subject matter of a long-lasting discussion. A fairly standard view holds that there is a distinction between the meaning of a sentence and the content of an utterance. The literalists' approach (Grice 1957, 1961; Borg 2004, 2012), which differentiates between semantic and pragmatic aspects of content, is especially influential. Since the seventies, however, we have been able to find many examples which aim at undermining this standard distinction between semantics and pragmatics. The examples are often provided by proponents of contextualism, including Travis (1997) and his painted leaves example and many other examples discussed by Recanati (2004, 2010, 2012a, 2012b).[63] Despite their differences, both literalism and contextualism share one thing – the discussion is done within the framework of truth-conditional semantics.

Beside the literalism-contextualism debate, we have been able to see a growth of inferentialism, as an alternative to any truth-conditional account of semantics, in the last few decades.[64] The biggest difference between truth-conditional and inferential semantics is in how meaning is understood in the first

[63] If not stated otherwise, I will refer to contextualism as advocated by Recanati in the rest of the paper. Beside Travis and Recanati, Searle (1992) also belongs among the proponents of contextualism. Other examples of the influence of context on content can be found in Rumelhart (1993).

[64] Despite the initial skepticism about inferentialism, its recent variations seem to withstand many traditional objections. See, for example, Peregrin (2014b, 2017), Brandom (2007), or the recent developments of Adjukiewicz (1934) and his directival theory of meaning as advocated in Grabarczyk (2017).

place. Meaning is a matter of truth conditions within the framework of truth-conditional semantics and it is a matter of inferential relations between sentences within the framework of inferentialism. The difference in their understanding of the nature of meaning is considered to be so huge that many semanticists find the doctrines to be incommensurable.

At the same time, inferentialism is a kind of a use-theory of meaning which partially follows the tradition of late Wittgenstein (1953) and, as such, it belongs among the approaches which favour pragmatics over semantics. Taking this assumption into account, there should be some connection between inferentialism and contextualism. Surprisingly, there has been no attempt so far to show whether inferentialism can accommodate contextualists' examples of the influence of context on content, or to explore their connections otherwise.

In this chapter, I would like to focus on the framework of normative inferentialism, as advocated by Brandom (1994, 2000) and Peregrin (2014b), and to clarify its position with regard to the literalism-contextualism debate. This general aim includes several partial aims. First of all, I would like to present the inferential view on meaning and to discuss how the distinction between *the inferential potential* and *the inferential significance* of sentences can fit the standard distinction between the meaning of a sentence and the content of an utterance. On the basis of this, I would like to present a specific inferential view on understanding and semantic deviations. In the third part of the chapter, I would like to discuss consequences that the inferential specification of meaning has for

its view on such phenomena as ambiguity, free pragmatic enrichment, or conversational implicature.

Throughout the chapter, I will restrict myself to discussing the inferential view on the meaning of sentences as I believe that this is the perspective from which inferentialism can bring the most interesting results. However, the fact that I will not discuss meaning on the subsentential level of words and predicates does not mean that inferentialism does not endorse compositionality.[65] The question of how compositionality works is at the heart of the debate between literalism and contextualism, but explaining the inferential position with regard to compositionality requires links to be made to the debate on a more general level first. This chapter can be seen as an attempt to make such general links – links on which a future discussion of compositionality and many other topics (such as indexicals, conventional implicature, polysemy, disagreements, or meaning similarity) could be provided from the inferential point of view.

5.1 Inferentialism on Meaning

As it was previously mentioned, inferentialism understands meaning via inferential relations of sentences. For example, the meaning of the sentence 'Fido is a dog' depends somehow on its relations to other sentences such as 'Fido is an animal', 'Fido is not a cat', or 'Fido eats bones'. Simply put, the main idea behind

[65] For a discussion of the inferential view on compositionality from a formal point of view see Peregrin (2009); for a discussion of the meaning of subsentential expressions as inferential roles see Peregrin (2006, 2014b: Ch.3).

inferentialism is that if we understand language as a huge web of inferentially interconnected sentences, then we could specify the meaning of a sentence via its position within this web.[66]

The fact that sentences are inferentially interconnected also plays a crucial role in the explanation of how languages are employed in communication. From this perspective, communication is understood first and foremost as an activity of language games, and the game of giving and asking for reasons is especially important. As Brandom (1994) argues, asserting a sentence, and more generally any move in a language game, bears with it commitments and entitlements. If a speaker asserts 'Fido is a dog' (in the most usual circumstances), she is committed to agreeing with the assertion that 'Fido is an animal' and the hearer is entitled to demand this commitment. Importantly, this game should be carried out in accordance with the inferential structure of a language.

The most straightforward way to specify the position of a sentence within the inferential structure of a language is to make a list of sentences from which a sentence can be inferred and sentences which can be inferred from a sentence (plus some collateral premises). However, in the game of giving and asking for reasons, a sentence can typically be backed up by several different reasons, i.e., any sentence can be inferred from several different sets of premises. Sentences tend to be intersections of several separate inferential paths within the inferential structure of a language and the sets of premises corresponding to those paths may

[66] This does not necessarily mean that there is no connection between language and the world. For example, we can follow Sellars (1974) and beside intra-linguistic transitions we can accept language entry (a speaker responds linguistically to a perceptual situation) and language exit (a speaker responds behaviorally to a linguistic activity) transitions as well.

be incompatible. This indicates that a simple list of sentences might not be sufficient for the specification of the position of a sentence, because lists may turn out to include mutually contradictory sentences. To avoid this problem, inferentialists must perform more detailed discrimination between sentences within lists.

5.1.1 The Inferential Potential as a Specification of Meaning

A more rigorous way to specify the position of a sentence can be found in Peregrin (2014b). The position of a sentence is specified as

> a pair of sets, one representing what A is inferable from (in particular containing the sequences of sentences from which A is inferable), the other representing what is inferable from A (together with possible collateral premises – in the form of triples the first two constituents of which represent the collateral premises and the third constituent a conclusion). (Peregrin 2014b: 50)

Note that Peregrin shows that "assuming the reflexivity and transitivity of the inference relation" (Peregrin 2014b: 50), each of the two halves of the pair of sets is capable of representing the whole pair.

For the sake of simplicity, we will focus only on the second part of the pair – the set which consists of ordered pairs of sets of collateral premises and

sentences that can be inferred from the sentence being analysed and those collateral premises.[67] The sets of ordered pairs can then be further divided into subsets which consist of one or more ordered pairs and which represent separate inferential paths in which a sentence can occur. With regard to this, meaning can be explicated as the inferential potential of a sentence (IP) and, subsequently, specified as a set of ordered pairs.

> If we now try to represent IP(A) by A→ (rather than A←), i.e., as {<{A1, ..., An}, An+1> | {A1, ..., An, A} ⊢ An+1}, IP(A) comes to relate {A1, ..., An} to An+1 if one's belief A, on the background of her beliefs A1, ..., An, warrants her belief that An+1; or that given the collateral commitments to A1, ..., An, the commitment to A brings about the commitment to An+1. Generalizing IP so that it applies also to sets of sentences (instead of merely single sentences), so that IP(M) = {<{A1, ..., An}, An+1> | {A1, ..., An} ∪ M ⊢ An+1}... (Peregrin 2014b: 51)

To summarize this proposal, the meaning (the inferential potential) of a sentence M can be specified as a set of ordered pairs:

$IP(M) = \{\langle X_1, D_1\rangle, \langle X_2, D_2\rangle ..., \langle X_n, D_n\rangle\}$

where $X_1 = \{A_1, ..., A_n\}$ (represents collateral premises) and D_1 is a sentence which can be inferred from $\{A_1, ..., A_n, M\}$; $X_2 = \{B_1, ..., B_n\}$ and D_2 is a sentence

[67] The abbreviated form 'set of ordered pairs' will be used henceforth.

which can be inferred from {B_1, ..., B_n, M}; X_n = {C_1, ..., C_n} and D_n is a sentence which can be inferred from {C_1, ..., C_n, M}.

For example, a simplified version of the inferential potential of the sentence 'Fido is a dog', as it may be used in everyday communication, can be

IP('Fido is a dog') = {<{'Fido barks', 'Fido eats bones', 'Fido is a living being', ...}, 'Fido is an animal'>, <{'Fido barks', 'Fido does not eat bones', 'Fido is not a living being', 'Fido is made of metal and plastic', ...}, 'Fido is a robot'> ..., <{'Fido does not bark', 'Fido does not eat bones', 'Fido is not a living being', 'Fido is made of plush', ...}, 'Fido is a toy'>}.

5.1.2 Context and the Inferential Specification of Meaning

At this point, one feature of the inferential specification of meaning deserves special attention with regard to the literalism-contextualism debate. According to contextualists, the meaning (truth conditions) of a sentence depends on the context in which the sentence is uttered. Note that contextualists understand context very broadly in comparison to how it has been understood since Kaplan (1978) or Lewis (1970, 1980) and how it is often understood by literalists.

According to a fairly standard view, e.g. Stanley and Szabo (2000), indexical expressions, such as 'I', 'here', or 'now', are specific because their extensions change from context to context. More specifically, four parameters of indexical expressions are usually considered to be context-dependent – agent,

time, place, and possible world. Context-dependent parameters of indexical expressions are represented by an n-tuple which specifies the values of relevant parameters for a context and those n-tuples are part of the formal representation of the meaning of those expressions. For example, the meaning of the sentence 'I saw Darth Vader there' can be represented in such a way that it includes a specification of person and place <CA, CP>. If pronounced by Luke Skywalker at the end of "The Empire Strikes Back", the particular specification of relevant contextual parameters for the sentence is <Luke Skywalker, Cloud City>.[68]

An explicit representation of contextual parameters within the literalists' specification of meaning is possible because the number of parameters is small and indexical expressions by themselves are easily identifiable. However, according to contextualists, basically any information about the world, the circumstances of a conversation, or a speaker may be contextually relevant for the meaning of a sentence. What is more, the examples provided by contextualists show that, in natural languages, not only indexical expressions are context-sensitive. For many other expressions we can find some relevant contextual parameters which modulate their content. For example, when Recanati (2004) argues that the meaning of the sentence 'The policeman stopped the car' can shift between contexts, depending on whether the policeman was standing by the road or was driving the car is presented as a contextually relevant parameter.

If this is so, then literalism, with its rather narrow view on context-dependence, would require serious amendments to be able to provide an

[68] Omitting the fact that the world is fictional.

appropriate specification of meaning for natural languages. Contextually relevant parameters could be explicitly represented by establishing new indices but, eventually, literalism would end up with thousands of parameters – turning an attempt to specify meaning in natural languages into chaos. With regard to this threat, contextualists usually give up their effort to represent context explicitly and use the notion of context very loosely. Contextually relevant information is often stated in the form of a short text before or after a sentence that is being analysed is mentioned. And a specification (or a formal representation) of meaning starts only after the contents of particular expressions are contextually modulated.[69] With a little simplification, we can say that the influence of context is mentioned but not represented within the framework of contextualism.

One of the biggest advantages of the inferential specification of meaning is that it includes context, as understood by contextualists, from the outset. The inferential potential of a sentence, as a specification of its meaning, includes collateral premises and sentences which can be inferred from the sentence and those premises. Depending on what we count as collateral premises, i.e., how we understand the inference relation, sentences included in inferential potentials could be understood as an explicit representation of context in the contextualists' sense.[70]

---

[69] That is not done very often by contextualists as they consider the contextual modulation of the meanings of particular words/predicates to be more important (or interesting) than a subsequent specification of meaning on the level of sentences. The idea is that if we contextually modulate the meanings of particular expressions, then the specification of meaning can continue standardly (whatever we take to be a standard specification of meaning).

[70] Which inferences count as constituents of the inferential potential, according to normative inferentialists, will be discussed in subchapter 5.1.4.

For example, the fact that the policeman was standing by the road can be explicitly represented by the sentence ‘The policeman was standing by the road’ within the inferential potential of the sentence ‘The policeman stopped the car’. Clearly, the sentence ‘The policeman was driving the car’ is also part of the inferential potential of ‘The policeman stopped the car’ but it belongs to a different set of collateral premises (as it corresponds to a different inferential path). A simplified version of the inferential potential of ‘The policeman stopped the car’ can be specified as

IP(‘The policeman stopped the car’) = {<{‘The policeman was driving the car’, ...}, ‘The policeman depressed the brake pedal’>, <{‘The policeman was standing by the road’, ...}, ‘The policeman raised his arm forward’> ...}.

5.1.3 The Inferential Significance

Inferentialists are aware that not all the inferential paths related to a particular sentence are relevant in particular conversations. When someone utters a sentence, she does so in the context of a conversation. For our purposes, let us say that if a sentence is implicitly or explicitly taken for granted by the participants in a conversation, we can say that it is part of the context of a conversation. Brandom talks about context in this sense as about the particular perspective which participants in conversation have with regard to meaning and because of that he talks about meaning as being “perspectival”.

> And since propositional and so conceptual contents of all sorts are constituted by the broadly inferential proprieties of practice in which they are caught up, such contents are essentially social and perspectival in nature. The propositional content of a claim or commitment can be specified only from some point of view; that it would be differently specified in definite ways from other particular possible social perspectives (that is, scorekeepers occupying such perspectives) is an essential part of its being the content it is. (Brandom 1994: 197)

Since meaning is always understood with regard to a certain perspective on the part of participants, inferentialism counts as a version of contextualism. Simply put, normative inferentialists hold to the principle that the context-dependent content of an utterance has primacy over the context-invariant meaning of a sentence.

Even though the basic idea of perspectival meaning has been presented by Brandom, his writings give us few clues as to how it should be understood with regard to the task of a specification of meaning. This aspect of normative inferentialism has been developed further by Peregrin, who explicates the perspectival meaning via the notion of the inferential significance.

> However, having explicated the potential as the above kind of function, their relationship turns out to be quite straightforward: the inferential significance of A within the context C is the value of the inferential potential of A for C.

> But this should not be read as claiming that potentials are prior to significances; a sentence has an inferential potential to the extent to which the employment of A becomes invariant across contexts, i.e., to which there emerge context-independent rules (which we explicate in terms of the function). (Peregrin 2014b: 51)

One way to spell out what Peregrin means by context in this quote is to say that he has in mind sentences which are taken for granted within a conversation and which are represented within the inferential potential as constituents of sets X_1, X_2, …, X_n. The inferential significance is then understood by Peregrin as "a matter of the inferential consequences of a sentence in a particular context" (Peregrin 2014b: 66). So, if IP(M) = {<X_1, D_1>, <X_2, D_2>..., <X_n, D_n>} and the premises that are taken for granted in the given contexts include all of X_1, ..., X_i, but none of X_{i+1}, ..., X_n, then the inferential significance is IS(M) = { D_1, ..., D_i}.

Roughly speaking, if understood in this way, the inferential significance represents commitments to which a speaker is committed by uttering a sentence and which a hearer is entitled to demand with regard to the context of a conversation. If understood in this way, the inferential significance does not represent the meaning of a sentence, or the content of an utterance. It is a theoretical tool specifically designed for the framework of inferentialism and it is hard to find a counterpart for it within the literalism-contextualism debate.

Here, I suggest a slightly different way to spell out the idea of perspectival meaning. My suggestion is to specify the inferential significance of a sentence as a subset of the inferential potential, or, more specifically, to specify the inferential

significance by that subset of the inferential potential which includes sentences implicitly or explicitly taken for granted within a particular conversation. For example, if IP(M) = {<X_1, D_1>, <X_2, D_2>..., <X_n, D_n>} and the premises that are taken for granted in the given contexts include all of X_1, ..., X_i, but none of X_{i+1}, ..., X_n, then the inferential significance is IS(M) = {<X_1, D_1>, ..., <X_i, D_i>}.

So if the inferential potential of the sentence 'Fido is a dog' is

IP('Fido is a dog') = {<{'Fido barks', 'Fido eats bones', 'Fido is a living being', ...}, 'Fido is an animal'>, <{'Fido barks', 'Fido does not eat bones', 'Fido is not a living being', 'Fido is made of metal and plastic', ...}, 'Fido is a robot'> ..., <{'Fido does not bark', 'Fido does not eat bones', 'Fido is not a living being', 'Fido is made of plush', ...}, 'Fido is a toy'>}

and the sentence 'Fido is a living being' is taken for granted within a conversation, then the inferential significance of the sentence can be specified as the subset of the inferential potential which includes 'Fido is a living being'

IS('Fido is a dog') = {<{'Fido barks', 'Fido eats bones', 'Fido is a living being', ...}, 'Fido is an animal'>}.

If the inferential potential represents the meaning of a sentence and the inferential significance is specified as a subset of the inferential potential for a particular context, then the inferential significance basically represents the meaning of a

sentence for a particular context.[71] With a fair bit of simplification we can say that the distinction between the inferential potential and the inferential significance represents the inferential version of the distinction between the meaning of a sentence and the content of an utterance.

However, the inferential distinction between the meaning of a sentence and the content of an utterance, as understood here, does not fit the view of either literalists or contextualists precisely. As opposed to contextualism, inferentialism holds that it is possible to specify context-invariant meaning. But the inferential notion of context-invariant meaning does not play the same explanatory role as it does within the framework of literalism.

If we understand the inferential significance as the value of the inferential potential for a particular context, then the meaning of a sentence (its inferential potential) is context-invariant in the sense that it includes all possible contextual values, i.e., all possible subsets which may be specified as its inferential significances. This kind of view is sometimes labeled a "rich meaning" approach.

From this perspective, the content of an utterance (its inferential significance) does not need to be explained as a deviation from, or modulation of, literal context-invariant meaning. In the same way, the process of contextual modulation does not need to be understood as a process in which the meaning of a sentence is changed into the content of an utterance. The process of contextual modulation is a process in which we delimit a specific part of meaning, or in

[71] Alternatively, we could coin a new term and call the inferential significance thus specified the actual inferential potential or the inferential potential for a particular context. Not much hinges on the terminology as long as we distinguish the current understanding of the inferential significance from Peregrin's previous proposal.

which we specify a particular aspect of meaning. The content of an utterance is not an outcome of a change of meaning; it is an outcome of a specification of meaning for a particular context.

5.1.4 Semantic Deviations and Meaning-constitutive Inferences

The inferential view on the distinction between the meaning of a sentence and the content of an utterance presupposes a specific view on semantic deviations. The framework of normative inferentialism allows us to distinguish between contextual modulations which are in accordance with the context-invariant meaning and contextual modulations which deviate from the context-invariant meaning. The answer to the question which contextual modulations are specifications of meaning and which are deviations from meaning depends on what belongs to the inferential potential of a sentence. Answering this question requires a little detour.

There are two major approaches to what determines the ordered pairs within inferential potentials that are currently under discussion. The first approach is in line with inferential role semantics as advocated by Boghossian (1993, 2005, 2012); the second approach is in line with normative inferentialism. According to the first approach, ordered pairs within inferential potentials are determined by the actual dispositions of particular speakers and their ability to make inferences. From the perspective of inferential role semantics, the meaning of a sentence

depends on inferential relations in which the sentence occurs in the language of thought of a particular speaker.

> Let's suppose that we think in a language of thought and that there are causal facts of the following form: the appearance in *O*'s belief box of a sentence S_1 has a tendency to cause the appearance therein of a sentence S_2 but not S_3. Ignoring many complications, we may describe this sort of fact as consisting in *O*'s disposition to infer from S_1 to S_2, but not to S_3. Let's call the totality of the inferences to which a sentence is capable of contributing, its *total inferential role*. (Boghossian 1993: 73)

If we follow this approach and we assume that the set of all the inferences in which a sentence actually appears (the "total inferential role") determines ordered pairs within the inferential potential of a sentence, then we may easily end up with the view that the meaning of a sentence varies from person to person. The ability of speakers to infer from sentence to sentence partially depends on their previous experience and general knowledge of the world and this is a rather subjective factor. For example, if a speaker S_1 believes that kynophobia is the fear of cats, then the inferential potential of the sentence 'Fido is a dog' in her language of thought does not include the ordered pair <{'Michael Jackson had kynophobia'...}, 'Michael Jackson was afraid of Fido'>. If a speaker S_2 believes that kynophobia is the fear of dogs, then the inferential potential of the sentence 'Fido is a dog' in her language of thought includes the ordered pair – and so the sentence differs in meaning between S_1 and S_2. Such a view is a legitimate target

of criticism by Fodor and Lepore (2001, 2007) for not being able to distinguish between meaning-constitutive and non-constitutive inferences.[72]

The second approach is advocated by Peregrin (2014b). According to this approach, the inferential potential does not include every inference that any particular speaker is disposed to make. Rather, the ordered pairs within inferential potentials are determined only by those inferences which almost all the speakers within a community accept as correct.

This is a consequence of a specific view of normative inferentialists on how natural languages evolved. According to normative inferentialism, natural languages are understood as outcomes of interlocking normative attitudes of speakers. As a result of a constant usage of corrective behavior, the acceptance of some inferences becomes widespread among almost all the members of a community and their acceptance is demanded as a sign of basic linguistic competence. These inferences are considered to be semantically relevant and these inferences form the inferential structure of language.[73]

In Peregrin's terminology, only those inferences which correspond to inferential rules determine ordered pairs within inferential potentials and only

[72] Fodor and Lepore explicitly address the objection to Brandom. I have always found Brandom's views hard to pinpoint and I do not intend to enter exegetical discussions. I believe that a) the objection fits the inferential role semantics better than normative inferentialism and b) at least Peregrin's version of normative inferentialism avoids the objection. See Drobňák (2017) for a detailed discussion of how Peregrin's view can be used to avoid the objection of Fodor and Lepore.

[73] The criterion of widely accepted inferential relations can be used as a response to the objection of Fodor and Lepore (2001, 2007). I agree that the criterion is far from the standard analytic/synthetic distinction that Fodor and Lepore expected – first and foremost, the criterion is vague. However, I see the vagueness of the criterion as its virtue. The criterion of widely accepted inferences is naturalistic in its nature. Its message is clear – if we want to know something about the semantics of natural languages, we have to go out and see how ordinary speakers use them in communication. This is the only source of data about natural languages that we have. If some cases turn out to be vague, then so be it. That is how natural languages work.

those inferences correspond to inferential rules which are widely accepted throughout the whole community of speakers.

> There is an inferential rule in force for a given language if the speakers of the language tend to see some inferences that violate the rule as incorrect. (Peregrin 2014b: 58)

and elsewhere

> And what I call a propriety, or an (implicit) rule, grows out of such attitudes resonating throughout the surrounding society. (Peregrin 2014b: 10)

For example, since the inferential relation between the sentence ‘Fido is a dog’ (together with the premise ‘Michael Jackson had kynophobia’) and the sentence ‘Michael Jackson is afraid of Fido’ is widely accepted, the ordered pair <{‘Michael Jackson had kynophobia’...}, ‘Michael Jackson was afraid of Fido’> is part of the inferential potential of the sentence ‘Fido is a dog’. On the contrary, since the inferential relation between the sentence ‘Fido is a cat’ (together with the premise ‘Michael Jackson had kynophobia’) and the sentence ‘Michael Jackson is afraid of Fido’ is not widely accepted, the ordered pair <{‘Michael Jackson had kynophobia’...}, ‘Michael Jackson was afraid of Fido’> is not part of the inferential potential of the sentence ‘Fido is a cat’.

Now, let us come back to discussing the question which contextual modulations are specifications of meaning and which are deviations from

meaning. If the inferential potential is understood in line with normative inferentialism, then we can distinguish between contextual modulations which are in accordance with the inferential potential and modulations which contravene the inferential potential. If the inferences drawn within a particular context correspond to a proper subset of the inferential potential of a sentence,[74] then the contextual modulation is understood as a specification of the meaning of a sentence for a particular context. If the inferences drawn within a particular context do not correspond to a proper subset of the inferential potential, then it is a case of semantic deviation. For example, if

IP('Fido is a dog') = {<{'Michael Jackson had kynophobia', 'Fido is a living being' ...}, 'Michael Jackson was afraid of Fido'>, <{'Fido barks', 'Fido eats bones', 'Fido is a living being', ...}, 'Fido is an animal'>, <{'Fido barks', 'Fido does not eat bones', 'Fido is not a living being', 'Fido is made of metal and plastic', ...}, 'Fido is a robot'> ..., <{'Fido does not bark', 'Fido does not eat bones', 'Fido is not a living being', 'Fido is made of plush', ...}, 'Fido is a toy'>}

then

<{'Michael Jackson had kynophobia', 'Fido is a living being' ...}, 'Michael Jackson was afraid of Fido'>

[74] The subset can consist of one or more ordered pairs.

counts as a specification of meaning and

<{‘Michael Jackson had kynophobia’, ‘Fido is a living being’ ...}, ‘Michael Jackson was not afraid of Fido’>

counts as a semantic deviation.[75]

This is not to say that semantic deviations cannot play a positive role in communication. On the contrary, we often rely on modulations which deviate from inferential potentials. Some examples of figurative language or conceptual engineering might serve as a nice demonstration. However, it is important to emphasize that understanding such modulations requires careful outlining of which inferences a speaker considers to be correct and which inferences she considers to be incorrect. In general, understanding such modulations is much more demanding than understanding standard cases. What is more, such modulations do not have to be accepted by a hearer, or they might be accepted only conditionally and the speaker might be a legitimate target of criticism.[76]

5.2 Understanding a Sentence/an Utterance

The inferential view on the relation between the meaning of a sentence and the content of an utterance as presented in the previous subchapter opens the way for

[75] Assuming that no other collateral premises are taken for granted.
[76] Very similar conversational situations are discussed by Ludlow (2014) under the name ‘meaning negotiations’.

a specific view on understanding. To fully understand a sentence in general requires tacit knowledge of all its possible contextual values, i.e., its inferential potential. To understand an utterance requires a hearer to be implicitly or explicitly able to single out an appropriate subset from the inferential potential of a sentence, i.e., its inferential significance.[77] This is usually done automatically. An explicit reconsideration takes place only if communication runs into obstacles.

Communication is then understood as a social art of manipulation with the inferential significance. It is a manipulation with the significance on the side of speakers because, by framing the context of a conversation, a speaker can influence what inferential significance her subsequent utterances will have. And it is a manipulation with the significance on the side of hearers because understanding an utterance is an activity of elimination of subsets of inferential potentials on the basis of what is known by a hearer about the context of a particular conversation.

Relations between particular ordered pairs within inferential potentials play the key role in the process of reaching understanding of an utterance. Sentences which are exclusively constituents of a single ordered pair within the inferential potential are especially important. Such sentences can serve as triggers which help participants in a conversation to single out effectively the most appropriate candidate for the content of an utterance (the inferential significance). The exclusivity of triggers explains their use in specific conversational situations,

[77] Such a view opens the interesting question of whether the recent neurological research on semantic and pragmatic understanding, e.g. Cosentino et al. (2017), could be reinterpreted from the inferential perspective. Because of the complexity of the topic, I have opted to postpone its discussion until another time.

e.g., in situations in which a newcomer is introduced into the context of a conversation.

As a demonstration of how triggers are employed in communication, let us take the inferential potential of ‘Fido is a dog’. There is only one ordered pair within the inferential potential of ‘Fido is a dog’ which includes ‘Fido is a robot’ and so ‘Fido is a robot’ can play the role of a trigger. Let us say that the speaker S_2 enters an ongoing conversation and the first sentence of this conversation is the only thing S_2 knows about the context of the conversation.

S_1: “...and all of a sudden Fido barked.”

S_2: “Oh, I didn’t know you have a dog.”

S_1: “It is a robot which my dad brought from his business trip to Japan.”

S_2 was right in inferring ‘Fido is a dog’ from what she has heard as ‘Fido barks’ is a constituent of several ordered pairs within the inferential potential of ‘Fido is a dog’. But since S_2 did not know anything else about the context at this point, there were still two ordered pairs that could be specified as separate subsets of the inferential potential, i.e., two possible candidates for the inferential significance of ‘Fido is a dog’. It might have been the ordered pair which includes ‘Fido is a robot’ or it might have been the ordered pair which includes ‘Fido is an animal’.

The first part of the last sentence, ‘It is a robot’, serves as a trigger. It can be seen as an explicit proclamation that ‘Fido is a robot’ is part of the context.[78] S_1 used the sentence ‘It is a robot’ simply because on the basis of this proclamation, S_2 was able to single out effectively the appropriate subset of the inferential potential, i.e., to determine the content of ‘Fido is a dog’ for the context of the conversation.

5.2.1 Partial Understanding

From the theoretical point of view, I believe that there is one phenomenon related to linguistic understanding which is in favor of inferentialism. Inferentialism, as presented here, can naturally embrace partial understanding of an utterance.[79] First of all, we sometimes have a problem to understand an utterance if it is uttered “out of the blue”, or if we enter an ongoing conversation. From the perspective of inferentialism, the reason why it is problematic to understand utterances in such situations is that if the context is not specified or unknown, the meaning of a sentence leaves too many candidates for the content of an utterance open. A typical feature of such situations is that a hearer may ask for clarification by

[78] The sentence carries much more information, of course. A clarification of the context was not its only purpose.

[79] Note that partial understanding of an utterance is not the same phenomenon as partial understanding of an expression as discussed, e.g., by Verdejo and de Donato Rodríguez (2015). We can understand fully a sentence and, at the same time, we do not have to understand an utterance at all. I believe that the framework of inferentialism can explain what partial understanding of a sentence (or its constituents) amounts to, but it is not the primary focus of this paper.

asking “What *exactly* do you mean by that?” (or something similar) or by reporting that she *“sort of”* understands an utterance.

Since the process in which we single out an appropriate inferential significance can be gradual (because we can discover the context sentence by sentence), understanding an utterance should come by degrees as well. Depending on what we know about the context, we can understand an utterance better or worse. By increasing the number of sentences known to be part of the context of a conversation, we can enhance our understanding of an utterance. For example, if we enter a conversation about Fido and we hear ‘Fido is a dog’ without hearing anything else, we may only guess which subset of its inferential potential is the content of the utterance. But if we hear ‘Well, it is not exactly a living pet’ later, we know at this point that we can eliminate those subsets of the inferential potential which include ‘Fido is a living being’, with the resulting inferential significance at this point to be

IS(‘Fido is a dog’) = {<{‘Fido barks’, ‘Fido does not eat bones’, ‘Fido is not a living being’, ‘Fido is made of metal and plastic’, ...}, ‘Fido is a robot’> ..., <{‘Fido does not bark’, ‘Fido does not eat bones’, ‘Fido is not a living being’, ‘Fido is made of plush’, ...}, ‘Fido is a toy’>}.

At this point, we were able to narrow the inferential potential on the basis of information about the pertinent context, but the information helped us to narrow the inferential potential only to some degree. There are still at least two ordered

pairs which can be specified as separate subsets of the inferential potential and so the content of the utterance can be specified more precisely.[80] For this reason, our understanding of the utterance at this point was only partial. If the conversation goes on and we hear 'At least, it barks', we are able to eliminate another ordered pair of the inferential potential and to understand the utterance fully. In other words, we are able to determine the content of the utterance fully

IS('Fido is a dog') = {<{'Fido barks', 'Fido does not eat bones', 'Fido is not a living being', 'Fido is made of metal and plastic', ...}, 'Fido is a robot'>}.

Situations in which we are not able to understand the content of an utterance fully because of a lack of information about context are quite common. In the following real-life example, a user asks other users for help in the SeatGuru forum because she/he is not able to determine the content of an utterance fully on the basis of the information available to her/him, i.e., she/he is not able to single out the appropriate inferential significance. This is the initial post from the user Katystava

> I am new to SeatGuru and am having trouble determining what exactly is meant by their terminology, "limited recline". Does that mean the seat reclines somewhat, or not at all, or it depends on the aircraft? My question

[80] This is not to say that to specify the content of an utterance fully always requires the inferential significance to be reduced to a single ordered pair. Such a reduction may lead to an inappropriate over-specification in many situations.

> pertains specifically to United 747, Row 61, which is the very last row in the plane just before the galley.[81]

The question posted in the forum can naturally be reinterpreted from the inferential point of view. It seems that Katystava knows what ordered pairs constitute the inferential potential of the sentence 'Seats may have a limited recline'[82] in general, i.e., she/he understands the meaning of the sentence. The second sentence in her/his post clarifies which ordered pairs from the inferential potential she/he was able to determine as possible candidates for the content of the utterance on the basis of information provided by SeatGuru and on the basis of her/his general knowledge of the practices of airlines. The problem is that this information was not sufficient to determine the content of the utterance fully as it still left at least three ordered pairs available. So in our terminology, Katystava understood the utterance partially at the time of posting the question.

The inferential significance of the sentence, as understood by Katystava at the time of posting, included three ordered pairs. The first ordered pair included the sentence 'Seats recline somewhat' as its constituent, the second ordered pair included 'Seats do not recline at all' as its constituent and the third ordered pair included 'Seats recline in some aircraft but not in others' as its constituent.

---

[81] The full conversation is available at https://www.tripadvisor.com/ShowTopic-g1-i12084-k3887710-Limited_recline_what_does_that_mean_exactly_United_747-SeatGuru.html

[82] This sentence is used as a description of particular seats on seatguru.com. In what follows, I will talk about the occurrence of this sentence on the website as about an utterance of the sentence. Even though the sentence was not actually pronounced, the discussion about its meaning in the forum clearly indicates that the sentence has a specific content in the context of its use on the website. So the example is an example of context-dependence of content, even though it is not a standard case of sentence-utterance distinction.

Katystava asked other users of SeatGuru, who might have more information about the relevant context, to help her/him to fully determine the content of the utterance, i.e., to eliminate inappropriate ordered pairs. And the answers provided by other users did exactly this. They indicated which of the three ordered pairs should be eliminated and by doing this they helped Katystava to reach full understanding of the utterance:

> For this type of aircraft they do recline but not as far as the other economy seats. The seat will recline until it hit the wall separating it from the galley.

> As kegler notes, some seats do have a *reduced* recline ability - usually this is the rows/seats just forward of a bulkhead, lavatory unit or cabin divider. But to say that they have *no* recline, is technically incorrect.

By stating ‘But to say that they have *no* recline, is technically incorrect’ the second answer suggests that the ordered pair which includes the sentence ‘Seats do not recline at all’ should be eliminated as the sentence is not part of the relevant context. By stating ‘they do recline but not as far as the other economy seats’ and ‘some seats do have a *reduced* recline ability’, both answers suggest that the appropriate inferential significance is constituted by the subset consisting of the ordered pair which includes ‘Seats recline somewhat’.

5.3 Ambiguity, Free Pragmatic Enrichment, and Conversational Implicature

After presenting a general inferential view on meaning and understanding, I will focus on particular phenomena related to natural languages and communication. I believe that the inferential specification of meaning, especially with regard to the distinction between the inferential potential and the inferential significance, is very well suited to representing and explaining phenomena that occur in natural languages such as ambiguity, free pragmatic enrichment, or conversational implicature.

5.3.1 Inferentialism on Ambiguity

Let us demonstrate how the inferential specification of meaning can be used to explain and accommodate ambiguity through the example of the sentence 'There is a bank around the corner'. Standardly, we would say that the sentence is ambiguous, because it includes the ambiguous expression 'bank'. The expression has two distinguishable meanings and it refers to two distinct kinds of objects – to financial institutions and to land at the edge of a river. If this is so, then the sentence 'There is a bank around the corner' has two meanings, depending on what is meant by 'bank' in a particular context.[83]

[83] See e.g. Gillon (1990) for a discussion of the standard view on ambiguity.

From the perspective of inferentialism, we could represent both "meanings" of the sentence as part of the single inferential potential. Clearly, depending on what is meant by 'bank', the sentence 'There is a bank around the corner' is part of different inferential paths. So if 'bank' refers to a river bank, then corresponding ordered pairs within the inferential potential of 'There is a bank around the corner' include sentences such as 'There is water around the corner' or 'You can cool down in a river'. On the other hand, if 'bank' refers to a financial institution, then corresponding ordered pairs within the inferential potential of 'There is a bank around the corner' include sentences such as 'You need money quickly' or 'You can take out a loan'. A simplified version of the inferential potential with regard to ambiguity can be

IP('There is a bank around the corner') = {<{'You need to get money quickly', 'A bank is a place where you can take out a loan', 'Taking out a loan is a quick way to get money', ...}, 'You can take out a loan around the corner'>, <{'A bank is an edge of a river', 'A river is a place where you can cool down if it is hot', 'It is hot today', ...}, 'You can cool down in a river'>, ...}.

Interestingly, there is nothing special about the specification of the meaning of ambiguous sentences if compared to the specification of the meaning of non-ambiguous sentences. The inferential potential of a sentence is specified standardly and the process of understanding can proceed standardly as well. To understand an utterance of an ambiguous sentence requires a hearer to be able to

single out an appropriate subset of the inferential potential on the basis of sentences that are taken for granted within a conversation.

Since we specify only one inferential potential for ambiguous sentences, strictly speaking, ambiguous sentences have only one meaning. Such a characterization of ambiguity makes the inferential view simple and counter-intuitive at the same time. It is simple because of the reasons mentioned above: a) the specification of meaning is standard, and b) understanding in such cases is explained standardly.

The counter-intuitiveness of this view seems to be obvious as well. Treating ambiguous sentences standardly may seem like pretending that there is nothing special about them but linguists, as well as lay speakers, have a strong intuition that there is a difference between ambiguous and non-ambiguous sentences/expressions. I understand the worry, but I do not think that it should be a reason for abandoning the inferential view on ambiguity.

The difference between (what are usually understood as) different meanings of a sentence can be recognized within a single inferential potential of a sentence. From the inferential point of view, ambiguity can be characterized relationally – as a matter of relations between ordered pairs within an inferential potential. One way to do that is to focus on the differences and overlaps between sentences within ordered pairs of an inferential potential. Our prediction is that in cases of ambiguity we can identify roughly two groups of ordered pairs in which all the ordered pairs within the first group a) share some sentences and b) (in the most radical case) do not share any sentence with the ordered pairs within the

second group. For example, in the case of ‘There is a bank around the corner’, sentences concerning finances might be almost exclusively shared within the first group and sentences concerning rivers might be almost exclusively shared within the second group. Even though we do not have any data regarding inferential potentials gathered so far, I believe that the prediction sounds reasonable.

5.3.2 Inferentialism on Free Pragmatic Enrichment

The inferential explanation of free pragmatic enrichment follows a very similar pattern. The meaning of sentences which are used as examples of free pragmatic enrichment can be specified standardly as their inferential potentials. What is explained by contextualism as different modulations of the meaning of a sentence can be explained by inferentialism as different contents that might be specified in particular contexts, i.e., as different inferential significances. The explanation of the understanding of such sentences/utterances can be standard as well – to understand the content of an utterance requires the ability to single out the appropriate subset of the inferential potential.

Let us use a modified version of the example ‘The policeman stopped the car’ from Recanati (2004) to demonstrate it. The policeman can stop the car a) by raising his arm forward, b) by depressing the brake pedal, or c) by firing a warning shot. According to contextualism, these three ways of stopping the car establish three modulations of the verb ‘stop’ and thus of the sentence ‘The policeman stopped the car’. From the perspective of inferentialism, we can

represent different ways in which a policeman can stop the car by classifying them into three distinct ordered pairs within the inferential potential of ‘The policeman stopped the car’:

IP(‘The policeman stopped the car’) = {<{‘The policeman was driving the car’ ... }, ‘The policeman depressed the brake pedal’>, <{‘The policeman was standing by the road’ ... }, ‘The policeman raised his arm forward’>, <{‘The policeman was driving a car that was chasing the car’ ... }, ‘The policeman fired a warning shot’>}.

So from the perspective of inferentialism, sentences which serve as examples of free pragmatic enrichment have only one meaning. This is because one inferential potential is sufficient to capture different modulations of their meaning. However, as in the case of ambiguity, the fact that the meaning of such sentences is represented by one inferential potential does not mean that we cannot distinguish cases of free pragmatic enrichment from other cases (e.g., from cases of ambiguity).

The characterization of the cases of free pragmatic enrichment can be relational again. We can say that what makes cases of free pragmatic enrichment special is the sharing of some sentences among all the ordered pairs within the inferential potential and, at the same time, the presence of unique sentences within particular ordered pairs. The sentence ‘The car did not move’ can serve as an example of a sentence which may be a constituent of all the ordered pairs within

the inferential potential as it is inferable from any set of collateral premises of 'The policeman stopped the car'. On the other hand, the sentences 'The policeman raised his arm forward', 'The policeman depressed the brake pedal', and 'The policeman fired a warning shot' are examples of sentences which are unique constituents of particular ordered pairs. At the end, their uniqueness is the reason why we used them to represent different modulations when we were specifying the inferential potential.

An interesting consequence of the relational view on free pragmatic enrichment is that it allows us to quantify how close or distant particular modulations are. In other words, the inferential specification of meaning allows us to talk about similarity of content. The ratio between the number of sentences that are shared and the number of unique sentences within ordered pairs of inferential potentials can be used as a measure of the similarity of particular contents which a sentence may have in particular contexts. Testing whether this measure fits our intuitive judgments of the similarity of modulations would require the gathering of data about inferential potentials. As mentioned before, we do not have such data yet. However, I believe that the prediction that the degree of similarity of modulations depends somehow on the number of sentences shared within ordered pairs and on the number of sentences which are unique to particular ordered pairs within inferential potentials sounds reasonable.

5.3.3 Inferentialism on Conversational Implicatures

Conversational implicatures are the most problematic of all the three phenomena discussed in this subchapter. Their explanation requires a slight change of perspective, but the resulting explanation is still in the spirit of the inferential view on meaning and understanding. Conversational implicatures are standardly understood as cases in which by saying one thing a speaker means something else. The inference from what was said to what was meant is standardly explained as taking place in accordance with some conversational principles (e.g., Gricean maxims).[84]

If we want to explain implicatures from the perspective of inferentialism, we need to understand how they differ from the normal cases of understanding. In the normal cases, the sentence being processed is known by a hearer (as it is standardly uttered by a speaker). The task for a hearer is to find out which sentences pertain to a context and, by doing so, to determine the inferential significance, i.e. the content of the utterance. In short, in normal cases, a hearer knows a sentence and needs to find its significance. The situation is reversed in the case of implicatures. The sentence of which the content should be determined is not actually uttered by a speaker. A speaker utters a sentence which is part of the inferential significance of some unspoken sentence. If the context is specific enough, a hearer is able to specify the inferential significance and her task is to find out to which sentence this significance belongs. In short, in the case of conversational implicatures, a hearer knows the significance, but needs to find a

[84] The original formulation of maxims can be found in Grice (1957, 1961). Neo-Gricean approaches which modify the original maxims can be found in Sperber and Wilson (1986), or Carston (2002).

sentence. The task for a hearer in the case of implicatures in many respects resembles the game of charades. A speaker gives clues to a hearer without actually using the sentence she has in mind.

Let us demonstrate it through a typical example of conversational implicatures.

S_1: "I need to get petrol."

S_2: "There's a garage just around the corner."

This example would standardly be interpreted as a case in which S_2 wants to convey a different meaning from that the sentence which is uttered standardly has. The different meaning is the one which would be standardly associated with the sentence 'You can get petrol in the garage around the corner'. If we analyse the case in accordance with the inferential view on meaning and understanding, the result is basically the same but the explanation of how a hearer gets to the result is different. When S_2 uttered 'There's a garage just around the corner', her aim was to indicate what inferential significance the sentence she has in mind has. The task for S_1 was to find out a) which inferential significance was indicated and b) to which sentence the inferential significance belongs.

The first part of the task is facilitated by the first sentence of the conversation 'I need to get petrol'. This sentence is explicitly part of the context. As the information that a garage is a place where people buy petrol is a generally known fact, we can consider the sentence 'A garage is a place where you can get

petrol' to be part of the context implicitly. When S_1 hears 'There's a garage just around the corner', she is in a position in which she is able to specify the subset which plays the role of the inferential significance for a sentence, but she does not know the sentence yet:

IS(?) = {<{'You need to get petrol', 'A garage is a place where you can get petrol', ...}, 'There is a garage around the corner'>}

At this point, she has to "scan" through the inferential potentials of different sentences to find out which sentence includes such a subset. Notice that this is a computationally much more demanding process than to take the inferential potential of a single sentence and to single out a subset of it. Because of that, one of the predictions of the inferential view is that understanding in the case of implicatures takes longer. The question of how exactly we do that, and whether we use some shortcuts in this process, is a matter for empirical research.

Such research would, again, require the gathering of data about inferential potentials, but there are some expectations that we may reasonably have. For example, the more specific a context is, the faster we should find the sentence out because the more specific a context is, the narrower the set of sentences which we consider as possible candidates is. This explains why understanding conversational implicatures out of context can be problematic. Moreover, we may expect that speakers use implicatures only under the right conditions, that is to say, only if the context is specific enough for a relatively easy specification of the

inferential significance and its assignment to a particular sentence. I leave it open as a matter for future discussion what exactly those “right conditions” are or in what respect a context has to be “specific”.

5.4 Conclusions

The aim of this chapter was to show how the inferential specification of meaning and content can be used to explain several phenomena related to the use and processing of natural languages. A specific inferential view on semantic understanding and deviations and on partial understanding of an utterance has been presented. Moreover, the last subchapter of this chapter includes an inferential explanation of such phenomena as ambiguity, free pragmatic enrichment and conversational implicatures.

I admit that the view on understanding is somewhat sketchy, as are the explanations presented in the last subchapter, and there are many details that must be elaborated. However, it is not possible to present a full-fledged view on every aspect of natural languages within the scope of one chapter of dissertation. My hope is that this chapter shows normative inferentialism as a view with a potential to become a general framework for natural languages which can compete with such views as literalism and contextualism. As a part of this project, I tried to present small clues as to how normative inferentialism could enter the discussion between literalism and contextualism, not only on theoretical, but also on empirical grounds. After all, it is a doctrine oriented towards natural languages

and their use in communication from its onset so it should not be very surprising that it may have testable predictions.

However, there is still a long way ahead. A proper and detailed discussion of the inferential view on ambiguity, free pragmatic enrichment, and conversational implicature might be the first step. I hope that this chapter laid some foundations for such detailed analysis. Developing an inferential view on other phenomena related to natural languages, such as compositionality, indexicals, conventional implicature, polysemy, disagreements, meaning similarity, and so on, might be the second step.

There is even more work to do with regard to empirical research based on normative inferentialism. Gathering data about widely accepted inferential relations between sentences, i.e., gathering data about inferential potentials, might be a necessary precondition for the success of such an endeavour. One way to gather such data would be by devising a specific Inferential Acceptability Judgment Task. Gathering such data for all the sentences within a natural language would be a project doomed to failure, but it could deliver reliable results for at least some parts of particular natural languages.

The view presented in this chapter does not deliver absolute and decisive solutions to all those issues. But this has not been the aim of the chapter anyway. My hope is only that this chapter helps to outline a specific inferential approach that may deliver explanations of the aforementioned linguistic phenomena and that it presents normative inferentialism as a doctrine worth exploring. Whether this outline will gain a more precise shape is a matter for future discussions.

References

Adjukiewicz, Kazimierz. 1934. “Sprache und Sinn.” *Erkenntnis* 4(1), 100–138.

Baggio, Giosuè, van Lambalgen, Michiel, Hagoort, Peter. 2008. “Computing and recomputing discourse models: an ERP study.” *Journal of Memory and Language* 59(1), 36–53.

Baggio, Giosuè and Hagoort, Peter. 2011. “The balance between memory and unification in semantics: a dynamic account of the N400.” *Language and Cognitive Processes* 26(9), 1338–1367.

Boghossian, Paul. 1993. “Does an Inferential Role Semantics Rest upon a Mistake?” In: *Philosophical Issues 3*. Ed. Enrique Villanueva. Atascadero: Ridgeview Press, 73–88.

Boghossian, Paul. “Blind Reasoning.” *Proceedings of the Aristotelian Society* 77(1), 225–248.

Boghossian, Paul. 2005. “Is Meaning Normative?” In: *Philosophy - Science - Scientific Philosophy*. Eds. Christian Nimtz and Ansgar Beckermann. Paderborn: Mentis, 205–218.

Boghossian, Paul. 2012. “Inferentialism and the Epistemology of Logic: Reflections on Casalegno and Williamson.” *Dialectica* 66(2), 221–236. https://doi.org/10.1111/j.1746-8361.2012.01303.x

Borg, Emma. 2004. *Minimal Semantics*. Oxford: Oxford University Press.

Borg, Emma. 2012. *Pursuing Meaning*. Oxford: Oxford University Press.

Brandom, Robert. 1994. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.

Brandom, Robert. 2000. *Articulating Reasons: An Introduction to Inferentialism*. Cambridge, MA: Harvard University Press.

Brandom, Robert. 2007. "Inferentialism and Some of Its Challenges." *Philosophy and Phenomenological Research* 74(3), 651–676. https://doi.org/10.1111/j.1933-1592.2007.00044.x

Brouwer, Harm and Hoeks, John. 2013. "A time and place for language comprehension: mapping the N400 and the P600 to a minimal cortical network." *Frontiers in Human Neuroscience* 7(758), 1–12.

Carston, Robyn. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.

Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.

Chomsky, Noam. 1971. "Recent contributions to the theory of innate ideas." In: *Philosophy of Language*. Ed. John Searle. Oxford: Oxford University Press, 121–129.

Chomsky, Noam. 1975. *Reflections on Language*. New York: Pantheon.

Chomsky, Noam. 1986. *Knowledge of Language*. Westport, CT: Praeger.

Chomsky, Noam. 1992. "Explaining language use." *Philosophical Topics* 20(1), 205–231.

Chomsky, Noam. 1995. "Language and nature." *Mind* 104(413), 1–61.

Chomsky, Noam. 2000. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.

Churchland, Paul. 1986. “Some Reductive Strategies in Cognitive Neurobiology.” *Mind* 95(379), 279–309.

Churchland, Paul. 1993. “State-Space Semantics and Meaning Holism.” *Philosophy and Phenomenological Research* 53(3), 667–672.

Cosentino, Erica, Baggio, Giosuè, Kontinen, Jarmo, Werning, Markus. 2017. “The Time-Course of Sentence Meaning Composition. N400 Effects of the Interaction between Context-Induced and Lexically Stored Affordances.” *Frontiers in Human Neuroscience* 8(813), 1–17. https://doi.org/10.3389/fpsyg.2017.00813

Dahl, Robert. 1956. *A preface to democratic theory.* Chicago: The University of Chicago Press.

Dahl, Robert. 1971. *Polyarchy: participation and opposition.* New Haven: Yale University Press.

Dahl, Robert. 1984. “Polyarchy, pluralism, and scale.” *Scandinavian Political Studies* 7(4), 225–240.

Davidson, Donald. 1986. “A Nice Derangement of Epitaphs.” In: *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson.* Ed. Ernie Lepore. Cambridge: Blackwell, 433–446.

Davidson, Donald. 1991a. “Epistemology Externalized.” *Dialectica* 45(2/3), 191–202.

Davidson, Donald. 1991b. “Three Varieties of Knowledge.” In: *A. J. Ayer Memorial Essays: Royal Institute of Philosophy Supplement 30*. Ed. Allen Phillips Griffiths. Cambridge: Cambridge University Press, 153–166.

Davidson, Donald. 1994. “The Social Aspect of Language.” In: *The Philosophy of Michael Dummett.* Eds. Brian McGuinness and Gianluigi Oliveri. Dordrecht: Kluwer Academic Publishers, 1–16.

Dennett, Daniel. 1971. *The Intentional Stance*. Cambridge, MA: MIT Press.

Donnellan, Keith. 1968. “Putting Humpty Dumpty Together Again.” *The Philosophical Review* 77(2), 203–215.

Drobňák, Matej. 2017. “Meaning-constitutive Inferences.” *Organon F* 24(1), 85–104.

Drobňák, Matej. 2018. “Quine on Shared Language and Linguistic Communities.” *Philosophia* 46(1), 83–99.

Dummett, Michael. 1986. “A Nice Derangement of Epitaphs: Some Comments on Davidson and Hacking.” In: *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. Ed. Ernie Lepore. Cambridge: Blackwell, 459–476.

Dummett, Michael. 1994. “Reply to Davidson: The Social Aspect of Language.” In: *The Philosophy of Michael Dummett.* Eds. Brian McGuinness and Gianluigi Oliveri. Dordrecht: Kluwer Academic Publishers, 257–262.

Fodor, Jerry. 1983. *Modularity of Mind.* Cambridge, MA: MIT Press.

Fodor, Jerry. 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.

Fodor, Jerry. 2000. *The Mind Doesn't Work That Way*. Cambridge, MA: MIT Press.

Fodor, Jerry and Lepore, Ernie. 1999. "All at Sea in Semantic Space: Churchland on Meaning Similarity." *Journal of Philosophy* 96(8), 381–403.

Fodor, Jerry and Lepore, Lepore, Ernie. 2001. "Brandom's Burdens: Compositionality and Inferentialism." *Philosophy and Phenomenological Research* 63(2), 465–481. https://doi.org/10.1111/j.1933-1592.2001.tb00119.x

Fodor, Jerry and Lepore, Ernie. 2007. "Brandom Beleaguered." *Philosophy and Phenomenological Research* 74(3), 677–691. https://doi.org/10.1111/j.1933-1592.2007.00045.x

Føllesdal, Dagfinn. 2008. "Introduction." In: *Confessions of a Confirmed Extensionalist and Other Essays*. Eds. Dagfinn Føllesdal and Douglas Quine. Cambridge, MA: Harvard University Press, 1–7.

Gauker, Christopher. 2017. "Open Texture and Schematicity as Arguments for Non-referential Semantics." In: *Meaning, Context, and Methodology*. Eds. Sarah-Jane Conrad and Klaus Petrus. Berlin: DeGruyter, 13–30.

Gillon, Brendan. 1990. "Ambiguity, Generality, and Indeterminacy: Tests and Definitions." *Synthese* 85(3), 391–416. https://doi.org/10.1007/BF00484835

Glüer, Kathrin. 2013. "Convention and Meaning." In: *A Companion to Donald Davidson*. Eds. Ernie Lepore and Kirk Ludwig. Chichester: Wiley-Blackwell, 339–360.

Grabarczyk, Paweł. 2017. "Directival Theory of Meaning Resurrected." *Studia Semiotyczne* 31(1), 23–43. https://doi.org/10.26333/sts.xxxi1.03

Grice, Paul. 1957. “Meaning.” *Philosophical Review* 66(3), 377–388. https://doi.org/10.2307/2182440

Grice, Paul. 1961. “The causal theory of perception.” *Proceedings of the Aristotelian Society*, Supplementary Volume 35, 121–152.

Hagoort, Peter, Baggio, Giosuè, Willems, Roel. 2009. “Semantic unification.” In: *The Cognitive Neurosciences*. Ed. Michael Gazzaniga. Boston, MA: MIT Press, 819–836.

Kaplan, David. 1978. “Dthat.” *Syntax and Semantics* 9(1), 221–243.

Hornsby, Jennifer. 2008. “Davidson and Dummett on the Social Character of Language.” In: *Knowledge, Language, and Interpretation On the Philosophy of Donald Davidson*. Eds. Maria Cristina Amoretti and Nicla Vassalo. Heusenstamm: Ontos, 107–122.

Hylton, Peter. 2007. *Quine*. New York, London: Routledge.

Kemp, Gary. 2017. “'Let the ideas fall where they may': Quine, Publicity, and Pre-established Harmony.” *Protosociology* 34(1), 59–72.

Koreň, Ladislav. 2017. “Peregrinova inferencialistická koncepcia jazyka a logiky (1. část).” *Filosofický časopis* 65(1).

Kutas, Marta and Federmeier, Kara. 2011. “Thirty years and counting: finding meaning in the N400 component of the event-related brain potential (ERP).” *Annual Review of Psychology* 62(1), 621–647.

Kutas, Marta and Hillyard, Steven. 1980. “Reading senseless sentences: brain potentials reflect semantic incongruity.” *Science* 207(4427), 203–205.

Kutas, Marta and Hillyard, Steven. 1984. “Brain potentials during reading reflect word expectancy and semantic association.” *Nature* 307(5947), 161–163.

Kutas, Marta, Lindamood, Timothy, Hillyard, Steven. 1984. “Word expectancy and event-related brain potentials during sentence processing.” In: *Preparatory States and Processes*. Eds. Sylvan Kornblum and Jean Renquin. Englewood Cliffs, NJ: Erlbaum Press, 217–237.

Labov, William. 1971. “The notion of ‘System’ in creole studies.” In: *Pidginization and Creolization of Languages*. Ed. Dell Hymes. Oxford: Cambridge University Press, 447–472.

Lewis, David. 1969. *Convention: A Philosophical Study.* Cambridge, MA: Harvard University Press.

Lewis, David. 1970. “General Semantics.” *Synthese* 22(1/2), 18–67.

Lewis, David. 1975. “Languages and Language.” In: *Minnesota Studies in the Philosophy of Science*, *Volume VII*. Ed. Keith Gunderson. Minneapolis: University of Minnesota Press, 3–35.

Lewis, David. 1979. “Scorekeeping in a Language Game.” *Journal of Philosophical Logic* 8(3), 339–359.

Lewis, David. 1980. “Index, Context, and Content.” In: *Philosophy and Grammar*. Eds. Stig Kanger and Sven Ōhman. Dordrecht: Springer, 79–100.

Ludlow, Peter. 2014. *Living Words: Meaning Underdetermination and the Dynamic Lexicon*. Oxford: Oxford University Press.

Machery, Edouard et al. 2004. “Semantics, cross-cultural style.” *Cognition* 92(3), B1–B12.

Machery, Edouard, Olivola, Christopher, De Blanc, Molly. 2009. “Linguistic and metalinguistic intuitions in the philosophy of language.” *Analysis* 69(4), 689–694.

Noonan, Michael. 1999. “Non-Structuralist Syntax.” In: *Functionalism and Formalism in Linguistics: General papers*. Ed. Michael Darnel. Philadelphia: John Benjamins Publishing Company, 11–33.

Orlean, Susan. 2011. “The Dog Star.” *The New Yorker*. Available at: https://www.newyorker.com/magazine/2011/08/29/the-dog-star

Peregrin, Jaroslav. 2006. “Meaning as an Inferential Role.” *Erkenntnis* 64(1), 1–35. https://doi.org/10.1007/s10670-005-4698-1

Peregrin, Jaroslav. 2009. “Inferentialism and the Compositionality of Meaning.” *International Review of Pragmatics* 1(1), 154–181. https://doi.org/10.1163/187731009X455875

Peregrin, Jaroslav. 2010. “Inferentializing Semantics.” *Journal of Philosophical Logic* 39(3), 255–274.

Peregrin, Jaroslav. 2012. “Inferentialism and the Normativity of Meaning.” *Philosophia* 40(1), 75–97.

Peregrin, Jaroslav. 2014a. “Implicitní pravidla.” *Organon F* 21(3), 381–398.

Peregrin, Jaroslav. 2014b. *Inferentialism: Why Rules Matter*. London: Palgrave Macmillan.

Peregrin, Jaroslav. 2018. “Is inferentialism circular?” *Analysis* 78(3), 450–454. https://doi.org/10.1093/analys/anx130

Pollock, Joanna. 2014. *Mental Content, Holism and Communication*. PhD. thesis. Available at: https://www.era.lib.ed.ac.uk/handle/1842/9853

Quine, Willard. 1951. “Two Dogmas of Empiricism.” *The Philosophical Review* 60(1), 20–43.

Quine, Willard. 1960. *Word and Object*. Cambridge, MA: MIT Press.

Quine, Willard. 1969a. “Epistemology Naturalized.” In: *Ontological Relativity and Other Essays*. Ed. Willard Quine. New York: Columbia University Press, 69–90.

Quine, Willard. 1969b. “Natural Kinds.” In: *Ontological Relativity and Other Essays*. Ed. Willard Quine. New York: Columbia University Press, 114–138.

Quine, Willard. 1969c. “Ontological Relativity.” In: *Ontological Relativity and Other Essays*. Ed. Willard Quine. New York: Columbia University Press, 26–69.

Quine, Willard. 1969d. “Propositional Objects.” In: *Ontological Relativity and Other Essays*. Ed. Willard Quine. New York: Columbia University Press, 139–161.

Quine, Willard. 1969e. “Speaking of Objects.” In: *Ontological Relativity and Other Essays*. Ed. Willard Quine. New York: Columbia University Press, 1–25.

Quine, Willard. 1970a. “Methodological Reflections on Current Linguistic Theory.” In: *Confessions of a Confirmed Extensionalist and Other Essays*. Eds. Dagfinn Føllesdal and Douglas Quine. Cambridge, MA: Harvard University Press, 215–227.

Quine, Willard. 1970b. “Philosophical Progress in Language Theory.” *Metaphilosophy* 1(1), 2–19.

Quine, Willard. 1974. *The Roots of Reference*. LaSalle: Open Court.

Quine, Willard. 1975. “Mind and Verbal Dispositions.” In: *Confessions of a Confirmed Extensionalist and Other Essays*. Eds. Dagfinn Føllesdal and Douglas Quine. Cambridge, MA: Harvard University Press, 244–256.

Quine, Willard. 1979. “Use and Its Place in Meaning.” In: *Meaning and Use*. Ed. Avishai Margalit. Dordrecht: Reidel, 1–8.

Quine, Willard. 1981. *Theories and Things*. Cambridge, MA: Harvard University Press.

Quine, Willard. 1985. “States of Mind.” In: *Confessions of a Confirmed Extensionalist and Other Essays*. Eds. Dagfinn Føllesdal and Douglas Quine. Cambridge, MA: Harvard University Press, 323–326.

Quine, Willard. 1989. “Mind, Brain, and Behaviour.” In: *Confessions of a Confirmed Extensionalist and Other Essays*. Eds. Dagfinn Føllesdal and Douglas Quine. Cambridge, MA: Harvard University Press, 347–351.

Quine, Willard. 1992. *Pursuit of Truth*, second edition. Cambridge, MA: Harvard University Press.

Quine, Willard. 1995. *From Stimulus to Science*. Cambridge, MA: Harvard University Press.

Quine, Willard. 1996. “Progress on Two Fronts.” In: *Confessions of a Confirmed Extensionalist and Other Essays*. Eds. Dagfinn Føllesdal and Douglas Quine. Cambridge, MA: Harvard University Press, 473–477.

Quine, Willard. 1997. “The Growth of Mind and Language.” In: *Confessions of a Confirmed Extensionalist and Other Essays*. Eds. Dagfinn Føllesdal and Douglas Quine. Cambridge, MA: Harvard University Press, 182–191.

Quine, Willard. 1999. “I, You, and It: An Epistemological Triangle.” In: *Confessions of a Confirmed Extensionalist and Other Essays*. Eds. Dagfinn Føllesdal and Douglas Quine. Cambridge, MA: Harvard University Press, 485–492.

Quine, Willard. 2000. “Reply to Stoutland.” In: *Knowledge, language, and logic: questions for Quine*. Eds. Alex Orenstein and Petr Kotatko. Boston and London: Kluwer.

Rapaport, William. 2000. “How to Pass a Turing Test: Syntactic Semantics, Natural Language Understanding, and First-Person Cognition.” *Journal of Logic, Language, and Information* 9(4), 467–490.

Rapaport, William. 2003. “What Did You Mean by That: Misunderstanding, Negotiation, and Syntactic Semantics.” *Minds and Machines* 13(3), 397–427.

Recanati, François. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.

Recanati, François. 2010. *Truth-Conditional Pragmatics*. Oxford: Oxford University Press.

Recanati, François. 2012a. “Contextualism: Some Varieties.” In: *The Cambridge Handbook of Pragmatics*. Eds. Keith Allan and Kasia Jaszczolt. Cambridge: Cambridge University Press, 135–149.

Recanati, François. 2012b. “Pragmatic Enrichment.” In: *Routledge Companion to Philosophy of Language*. Eds. Delia Fara and Gillian Russell. New York: Routledge, 67–78.

Risjord, Mark. 2007. “Who are We: Dissolving the Problem of Cultural Boundaries.” *The Modern Schoolman*, 84(2/3), 205–215.

Rumelhart, David. 1993. “Some Problems with the Notion of Literal Meanings.” In: *Metaphor and Thought*. Ed. Andrew Ortony. Cambridge: Cambridge University Press, 71–82.

Schwarz, Wolfgang. 2013. “Against Magnetism.” *Australasian Journal of Philosophy* 92(1), 17–36.

Searle, John. 1992. *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.

Sellars, Wilfried. 1949. “Language, Rules, and Behavior.” In: *John Dewey: Philosopher of Science and Freedom*. Ed. Sidney Hook. New York: Dial Press, 289–315.

Sellars, Wilfried. 1974. “Meaning as Functional Classification.” *Synthese* 27(3/4), 417–437. https://doi.org/10.1007/BF00484606

Shapiro, Lionel. 2004. “Brandom on the Normativity of Meaning.” *Philosophy and Phenomenological Research* 68(1), 141–160.

Shapiro, Stewart. 2006. *Vagueness in Context*. Oxford: Clarendon Press.

Sidnell, Jack. 2010. *Conversation Analysis: An Introduction.* New York: Wiley-Blackwell.

Sperber, Dan and Wilson, Deirdre. 1986. *Relevance: Communication and Cognition*. Oxford: Blackwell.

Stanley, Jason and Szabo, Zoltán. 2000. “On Quantifier Domain Restriction.” *Mind and Language* 15(2/3), 219–261. https://doi.org/10.1111/1468-0017.00130

Stoutland, Frederick. 2000. “Individual and Social in Quine’s Philosophy of Language.” In: *Knowledge, language, and logic: questions for Quine*. Eds. Alex Orenstein and Petr Kotatko. Boston and London: Kluwer, 181–194.

Travis, Charles. 1997. “Pragmatics.” In: *A Companion to the Philosophy of Language*. Eds. Bob Hale and Crispin Wright. Oxford: Blackwell Publishers, 87–107. https://doi.org/10.1111/b.9780631213260.1999.00006.x

Verdejo, Victor and de Donato Rodríguez, Xavier. 2015. “Partial Understanding and Concept Possession: A Dilemma.” *Ratio* 28(2), 153–162. https://doi.org/10.1111/rati.12065

Waismann, Friedrich. 1945. “Verifiability.” *Proceedings of the Aristotelian Society,* Supplementary Volume 19, 119–150.

Weintraub, David. 2007. *Is Pluto a Planet?* Princeton: Princeton University Press.

Weiss, Bernhard. 2010. *How to Understand Language*. Durham: Acumen.

Werning, Markus and Cosentino, Erica. 2017. “The Interaction of Bayesian Pragmatics and Lexical Semantics in Linguistic Interpretation.” In: *Proceedings of the 39th Annual Conference of the Cognitive Science Society*. Eds. Glen Gunzelmann et al. Austin, TX: Cognitive Science Society, 3504–3509.

Wilson, Mark. 1982. “Predicate Meets Property.” *The Philosophical Review* 91(4), 549–589.

Wilson, Mark. 2006. *Wandering Significance: An Essay on Conceptual Behavior*. Oxford: Oxford University Press.

Wittgenstein, Ludwig. 1953. *Philosophical Investigations*. Oxford: Blackwell.