

Contextual Meaning and Theory Dependence*

Erich H. Rast
Universidade Nova de Lisboa

Abstract

An overview of linguistic context dependence is given and it is argued that an approach is needed that considers interpretation an inference from truth-conditionally incomplete to more specific contents. However, the theory dependence of idiolects and concepts is another context dependence that poses philosophical problems. It is argued that these are worst for global holism of idiolects and concepts. Since local holism is more tenable, and rational agents also need to have the ability to compartmentalize mutually incompatible theories, the theory dependence of meaning is less problematic than it might appear at first glance.

1 Introduction

Natural language elicits many forms of context dependence. Many of them are overt. For example, indexicals depend on the deictic center I-here-now. However, there are more subtle forms of context dependence in natural language that are less regulated by meaning rules and more pragmatic. Roughly speaking, hearers arrive at an interpretation of what the speaker said based on what they believe the speaker assumes in the context of a conversation. I argue in this article that this interpretation process requires speakers to be able to track other speakers' theories, and that epistemic agents generally must have the ability to consider and compartmentalize theories without necessarily endorsing them.

In Section 2, a brief overview of select phenomena of linguistic context dependence is provided and it is argued that these are overall tractable by understanding interpretation in a context as an inference from often truth-conditionally incomplete to more specific semantic representations. However, there is a more profound and philosophically more challenging context dependence that can be described as a dependence of concepts and lexical meaning on theories. This is laid out in Section 3, in which several problems are discussed that result from the interdependence between lexical meaning and theories. I argue in Section 4 that these problems can be solved by rejecting global meaning holism in favor of local meaning holism and by acknowledging rational epistemic agents' ability to compartmentalize and keep track of theories.

*Email: erich@snafu.de. Postal address: E. Rast, Universidade Nova de Lisboa, Nova Institute of Philosophy, Avenida de Berna 26C 1069-061 Lisboa. This is a preprint. Please only quote the [published version](#).

2 Semantic Contextualism: A Brief Overview

This section provides a brief overview of linguistic context dependence. Much of the modeling of context dependence in the philosophy of language and epistemology is based on [Kaplan \(1989\)](#)'s *Logic of Demonstratives*, which lead to different versions of 'two-dimensional' semantics ([Chalmers, 2006](#)) and corresponding forms of alethic contextualism and relativism. These approaches are the topic of the next section. Section 2.2 addresses their shortcomings and promotes the alternative view that many forms of linguistic context dependence are better described as an inference from potentially truth-functionally incomplete to more specific semantic contents instead of using parameterized modal logics.

2.1 Parameterized contexts

Indexicals and demonstratives are typical overtly context-sensitive expressions. The reference of indexicals like *here*, *I*, and *now* depends on features of the utterance context. Their linguistic meaning partly mandates the resolution of this context dependence. For example, under normal circumstances, *I* refers to the speaker of the utterance, *now* to the time of the utterance, and *here* to the place of utterance. Such a rule 'picks out' the respective referent in a given context of utterance, thereby resolving the context dependence semantically. The result of this enrichment process is a proposition that is true or false in the given circumstances of evaluation. Absolute tenses are also often used indexically. To fully understand a use of the present tense a hearer may have to know the time of utterance, for instance.¹

Understanding utterances with indexicals comes to a degree because the corresponding contextually-provided referents may be determined more or less precisely. In a sense, a hearer understands an utterance of *Yesterday, Bob had an accident* without knowing what day of the month or week it is; the accident happened the day before whatever day is now. However, this minimal understanding may turn out to be insufficient for a given communicative task. For example, when filling out an insurance form, merely knowing that something happened the day before the day on which the utterance took place might not suffice because a calendrical date is expected. A more precise understanding of the utterance could be paraphrased as *Bob had an accident on Friday, the 13th of November 2020*.²

Every indexical allows for such grades of understanding. Sometimes when interpreting a use of *I* it may suffice to know that someone spoke, whoever that may have been; in other cases, the hearer must spatiotemporally locate the speaker before they can rightfully be said to have understood the utterance as a whole. However, the reference to the deictic center acts as a hard constraint in any of those cases. The deictic center usually consists of the speaker, the time of utterance, the actual world, and the place of utterance. It can be shifted in some languages in indirect speech reports and for certain expressions like *local*, *around*, and medical uses of *right* and *left*.³ The hearer, and sometimes even the speaker, can be wrong about this reference, and in that case, they fail to grasp the semantic content of the utterance. This dependence on a fact about the world that is independent of the speaker's intentions is characteristic

of indexicals. Other context-dependent expressions need not involve a deictic center in this way.

In early approaches to indexicals such as [Reichenbach \(1947\)](#), [Burks \(1949\)](#), and [Bar-Hillel \(1954\)](#) a crucial question was whether these could be eliminated from a language that would serve as a foundation for all science. Bar Hillel argued that even though sentences containing indexicals can be substituted with sentences containing no indexicals, the reference to a conventionally fixed origin of a coordinate system cannot be eliminated. In tense logic, [Prior \(1957, 1967, 2003\)](#) famously argued that the logic of becoming and going expressed by operator-based tense logic could not be replaced without significant loss of expressivity by statements that quantify over points in time or time intervals directly and thereby lead to eternally true or false propositions. His arguments for this view were metaphysical and partly hinged on a specific interpretation of McTaggart's Paradox ([McTaggart, 1908](#)). In the Philosophy of Language the irreducibility of the basic indexicals *I*, *here*, and *now* was brought up by [Castañeda \(1967, 1989b,a\)](#) and [Perry \(1977, 1979, 1998b,a\)](#), and has been discussed in numerous follow-up publications. In this debate, the key question was whether thoughts, corresponding truth-functionally complete propositions, and broadly-conceived epistemic states that would ordinarily be expressed using indexicals like *now* and *I*, could be expressed by expressions only containing third person referential terms such as proper names and definite descriptions. There is a certain consensus in the literature that at least *now* and *I* are irreducible in cognition in terms of their expressive power for explaining behavior, which lead to various theories of *de se* belief attributions that take into account the 'essential indexicality' of these indexicals.⁴

A more recent debate started with [Recanati \(2004b\)](#) versus [Cappelen and Lepore \(2004\)](#). It addresses the more general question about linguistic context dependence's pervasiveness and what this means for literal meaning. Much of this discussion concerns the extent to which double-index modal logics can adequately represent linguistic context sensitivity. As part of the philosopher's toolbox, based on [Kaplan \(1989\)](#) and [Lewis \(1980\)](#), various modal logics and their interpretations are used in which contexts and circumstances of evaluation (CEs) are reified as parameters relative to which truth-in-a-model is determined. In Kaplan's two-layered account, for instance, the linguistic meaning (the character) of an expression is a function that in a context yields an intension (the semantic content), which is, in turn, a function that in given circumstances of evaluation yields an extension. Based on such modal logics with contexts and CEs, various contextualist and relativist positions have been developed and contrasted with Cappelen and Lepore's semantic invariantism on the one hand and Recanati's more radical contextualism on the other hand.

Going into the details of this complex debate would go beyond the scope of this contribution. Only a brief summary can be given. According to invariantism, simple clauses are not context-dependent except for the obvious and overt context dependence of indexicals. [Cappelen and Lepore \(2005\)](#) even go so far as to claim that a giraffe can be tall simpliciter. Others such as [Recanati \(2006\)](#) and [MacFarlane \(2007\)](#) found such an approach unsatisfactory. According to the radical contextualism of [Recanati \(2004a\)](#), the literal meaning on which such invariantist positions are based is an 'idle wheel'; instead, according to Recanati pragmatic modulation functions may change linguistic meaning on the fly during semantopragmatic construction of sentence-level content.

In contrast to this, indexicalists like [Stanley \(2004, 2005\)](#) model context-dependent expressions with open argument places bound either by semantic or by pragmatic processes. This use of open argument places makes their accounts slightly different from two-dimensional moderate contextualists who continue to use double-index modal logics to model a richer set of contextual variances than those elicited by overt indexicals. In the approaches based on modal logics with multiple parameters, these parameters are enriched with whatever additional ingredients are needed to get the semantics of context-sensitive expressions right that do not overtly depend on the deictic center. Usually, they are modeled as n -tuples containing all the needed ingredients.⁵

Broadly speaking, two-dimensional accounts come in three different varieties. According to classical contextualism, in a Kaplan-style two-layered modal logic the semantic content of the expression is fixed by some mechanism that takes into account features of the context parameter. If a context-sensitive expression is modeled in this way, then varying contexts will yield different semantic contents. This is the classical model of indexicals. In contrast to this, according to the nonindexical contextualism of [MacFarlane \(2009\)](#) the extension of an expression may depend on the context although the semantic content remains context-invariant. This means that the semantic content—i.e., the proposition expressed by the sentence in a context—is itself context-sensitive. Finally, according to full-fledged alethic relativism in a two-dimensional framework, the semantic content may yield different extensions not depending on the context but depending on non-traditional features of the circumstances of evaluation. Tense operators and modalities work in that way in traditional double-index modal logics because these operators implicitly quantify over time and possible worlds. In the debate between contextualists such as [de Sa \(2008, 2009\)](#) and relativists like [MacFarlane \(2008, 2012, 2014\)](#), relativists have argued that many more expressions may be truth-relative in this sense.

Within this discussion, some authors suggested that certain predicates of personal taste give rise to faultless disagreement between speakers that only a relativist semantics can adequately represent. In such a theory, the extension associated with a semantic content in given circumstances of evaluation is not just relative to times and possible worlds, but also relative to very nontraditional constituents of circumstances of evaluation parameters such as persons. For example, [Lasnik \(2005\)](#) argues that the predicate *fun* is sensitive to an assessor (or, judge, in his parlance) in given circumstances of evaluation. Regardless of who is the speaker of an utterance, in this assessor-relativism an utterance of *Roller coasters are fun* may be true relative to one and false relative to another assessor. Consequently, two people may disagree about an utterance containing such an expression without one being at fault. They may both be right even when they seemingly contradict each other and one of them negates the other's statement. Relative to one assessor the semantic content of the proposition may be true and relative to another assessor the semantic content of its negation may be true. To do justice to this position, it is worth noting that each of the assessors may still be mistaken in such an approach. For example, an assessor might erroneously believe that roller coasters are fun (relative to her); actually riding a roller coaster would make her realize that she was wrong right from the start.

The differences between parameter-based traditional contextualism, nonindexical contextualism, and relativism primarily hinge on the role given to semantic content in theorizing. The idea behind relativist faultless disagreement is that two assessors who disagree faultlessly disagree *about* the same semantic content of an utterance. The relativist argues against the contextualist that two interlocutors would disagree about two different contents according to the contextualist two-dimensional semantics. If the assessor in one context is John and the assessor in another one is Mary, then under a contextualist semantics the content of Mary's beliefs would be the proposition that roller coasters are fun for Mary, and the content of John's beliefs would be the proposition that roller coasters are fun for John. According to the relativist, this cannot count as disagreement because the contents of their beliefs remain compatible with each other. This relativist standard objection to contextualism will play a role in the second part of this article and should be kept in mind.

However, if the peculiar notion of semantic content is not available because the model is not two-layered, if attitudes are modeled in another way, if incompatible contents are not taken as a necessary condition for disagreement, or if the disagreement is modeled on the basis of other content—such as content expressed by pragmatic presuppositions or any other pragmatically derived, non-literal speech act content—, then the differences between parameter-based contextualism and relativism become less critical. Both theories have in common that they model contextual variance in a truth-conditional setting. If a context dependence is linguistically mandated like in the case of the truth-conditions expressed by the use of an indexical, then to some extent these parameterized approaches to context dependence model the contextual resolution process. For example, for *yesterday* the linguistically mandated reference rule is *the day before the day of the utterance*. It generally picks out the right referent and can be formalized in a double-index modal logic in which the date of the utterance is stored in the context parameter (provided that date calculations are available). Likewise, a relativist semantics for predicates like *being fun* and *tasty* states that utterances containing these expressions are true or false relative to an assessor and the suggested interpretation of the semantic apparatus is that respective assessors may differ from the speaker of the utterance.

2.2 Semantic underdetermination and interpretation as inference

Although it is adequate for indexicals at a high level of abstraction, modeling other context-sensitive expressions *as if they were* indexicals can be misleading and inadequate. Many forms of linguistic context dependence are pragmatic, and sentence-level content is often semantically underdetermined. For example, there is no linguistic rule in the meaning of *ready* that determines what a person is ready for. The hearer must figure out what the speaker means by an utterance of (1) *He's ready*. Dubbed 'contextuals' by Rast (2014), such expressions require some additional interpretation; in the case of *ready*, there is a syntactically optional complement clause that is not optional from the perspective of sentence-level semantics. This is similar to cases such as *to buy* which also has syntactically optional argument places for a seller and a price, but from a semantic perspective requires these ingredients to differentiate it from

other transfer verbs like *to obtain*, *to pay*, and *to borrow*.⁶ Other expressions suggest a default interpretation, sometimes very strongly, but neither require it semantically nor syntactically. For instance, (2) *John had breakfast* can be meant to convey that John has had breakfast for the first time in his life, but by default it is taken to express the proposition that John had breakfast at the day of the utterance of (2). Indexicals are also often contextual in this sense in addition to their dependence on the deictic center. For instance, the place denoted by a use of *here* can only be determined on the basis of what has been said so far and assumptions about what the speaker wants to convey, as the place denoted by a use of *here* only needs to contain the deictic center as a mereological part and may otherwise be almost arbitrarily small or large. Depending on what has been said so far and the speaker's intentions, a use of *here* may be intended to convey locations such as *here in this box* (*where the speaker is crouching*), *here in this room*, *here in this building*, *here in this city*, *here in this country*, *here on this continent*, *here on this planet*, and *here in this part of the Milky Way*.

Essentially three approaches have been proposed to deal with these forms of partly conventionalized, yet ultimately pragmatic context dependence. According to [Bach \(2005\)](#), utterances often express only propositional radicals by virtue of conventionalized meaning provided by a shared lexicon. What the speaker meant needs to be inferred from these truth-conditionally incomplete representations by a Gricean interpretation process. [Rast \(2014\)](#) suggests a variation of this approach that models the missing contextual factors as open argument places over which one may existentially quantify to obtain a minimal form of content. For example, the 'existential completion' of (1) is *John is ready for something*. Based on such representations, abductive inference may yield more specific content such as *John is ready to call a cab*. This inference is derived from what has been said so far, from the topic of the conversation and question under discussion, from the interpreter's assumptions about what the speaker believes, and from general common-sense world knowledge. While the mechanisms laid out by [Rast \(2010, 2014\)](#) are very limited, the approach in general is based on the idea of considering interpretation as an inference to the best explanation (IBE). Relevance theory of [Sperber and Wilson \(1986, 2004\)](#) is a third, psychologically motivated approach to interpretation. It is based on bounded rationality. Hearers draw inferences about what the speaker wants to convey but this process competes with economy constraints. As long as logical and set-theoretic representations of semantic content are used, these three approaches can be linked up with the modeling of pragmatic context, common ground, and linguistic score-keeping at discourse level such as [Stalnaker \(1978\)](#), [Lewis \(1979a\)](#), [Barwise and Perry \(1983\)](#), [Stokhof and Groenendijk \(1991\)](#), [Kamp and Reyle \(1993\)](#), [Asher and Lascarides \(2003\)](#), and [Ginzburg \(2012\)](#).

The key to making any of these approaches fruitful is to represent semantic underdetermination of conventionalized meaning in a way that allows the interpreter to infer what the speaker meant based on existing beliefs about the speaker, what has been said so far, the common ground, general world knowledge, and knowledge about the particular communication situation. The approaches primarily differ in the extent to which they are motivated from empirical psychology. Relevance theory strives for empirical adequacy, whereas the Gricean model describes ideal communication situations and ideal interpretation. The IBE approach's degree of idealization lies in-between. It is based on broadly-conceived logical inference mechanisms from

graded belief representations of common-sense ontologies and situational knowledge. All three approaches can be adopted for varying assumptions about the degree of conventionalization of meaning in a shared lexicon.

However, existential completions and Bach's propositional skeletons have to rely on mechanisms that allow for a finite number of existing argument slots to be 'filled in' by the interpretation process. Radical contextualists like [Recanati \(2004a\)](#) do not believe that such mechanisms suffice in general to adequately describe linguistic context dependence because they might not capture creative and poetic language use. Moderate contextualists in turn consider radical contextualism too general and unconstrained since pragmatic modulation functions can, in theory, turn any meaning into another meaning during semantic composition. The position of moderate contextualism is that the number of conventionalized contextual factors—those that are marked in a shared lexicon by semantic argument structures of words—may be large and require a decent amount of sophisticated semantic analysis, yet their number is ultimately finite. Likewise, it is stipulated that the number of rule-governed, broadly-conceived linguistically regulated pragmatic interpretation patterns is also finite.

Speakers and interpreters may occasionally allow contextual shortcuts whose understanding requires general intelligence instead of fixed, rule-based mechanisms. For example, a polite speaker of Japanese may leave out almost any part of speech. Understanding such an utterance and the meaning of not verbalizing part of the speech requires more than just linguistic skills and knowledge, and it is doubtful whether an inference to the best explanation mechanism could adequately explain such cases in sufficient detail in a rule-based manner. However, the existence of such phenomena does not speak against moderate contextualism in the same sense as not understanding someone's explanation of a mathematical problem does not speak against semantics. Understanding an utterance often requires intelligent reasoning that goes far beyond of a speaker's linguistic competence and what can reasonably be expected to be dissected by linguistic theorizing.

The problem of linguistic context dependence is thus principally solvable from the perspective of moderate contextualism. The challenges are in the detail, such as how to find an adequate semantic representation that allows for fruitful descriptions of the inferences that take place when a hearer interprets an utterance, systematic ways of cataloging a language's context-dependent expressions, and how to describe and model these inferences at a desired level of idealization. However, another potential source of context dependence is neither modeled by parameterized modal logics nor by the above mentioned inferential approaches: the possible dependence of meanings and concepts on background theories, opinions, and world views. This context dependence is the subject of the remainder of this article.

3 The Problems of Theory Dependence

Recently, there has been a renewed interest in philosophical aspects of discussions about word meaning. [Plunkett and Sundell \(2013, 2019\)](#) and [Plunkett \(2015\)](#) have argued that disputes are often implicitly about the meaning of words, the adequacy of using words in context, and the appropriateness of contextual norms. If some such

disputes concern word meaning, then one may ask how speakers can understand each other if they presume different word meanings from the start. If, in turn, two speakers defend different theories about a particular topic and these theories characterize or define the word under dispute in different ways, then this leads to various *problems of theory dependence*. The topic has a long tradition in analytic philosophy. The role that theory dependence plays for lexical meaning is crucial for assessing Moore's thesis of 'good' as a primitive and the Paradox of Analysis (Moore, 1903), as well as for a later debate between Quine and Carnap about the internal/external distinction of theories and the notion of analyticity in studies by Carnap (1950) and Quine (1960, 1951).

3.1 The problems

In what follows, the word *theory* shall be understood in the broadest possible sense as including all kinds of nonscientific beliefs, opinions, and world views in addition to scientific theories, approaches, models, and hypotheses. Given that broad understanding, the problems of theory dependence may be summarized as follows. Every theory either directly defines the meanings of words mentioned in it or indirectly characterizes the meanings of words used to formulate it by law-like statements in which those words are used *or* mentioned. Therefore, a definitional account of the meaning of a word central to a theory is directly or indirectly restricted by that theory. So in the context of two different theories, the meanings of words that are central to those theories are restricted in different ways, and, in the worst case, cannot even mean the same because those theories define or indirectly characterize their meanings in different ways.⁷

For example, if it follows from a physical theory that atoms can be split, then an adequate characterization of the meaning of the word *atom* cannot attribute the property of being indivisible to atoms. As another example, Arianism is the Christian belief that Jesus (God the Son) is not co-eternal with God the Father. Someone who believes this doctrine cannot at the same time believe in the trinity, that God the Father, God the Son, and the Holy Spirit are of the same essence. The Arian doctrine thus affects the possible theological characterizations of both *Jesus* and *God*. Historically, the conflict between Arians and Trinitarians led to persecution and violent clashes during the 4th Century AD, and ultimately the official church position was to declare Arianism a heresy. As a third example, consider competing theories of social institutions. According to Searle (1995, 2005), "... an institution is any system of constitutive rules of the form *X counts as Y in C*" (Searle, 2005, p. 10). In contrast to this, Guala (2016) argues that institutions are systems of regulative rules that lead to game-theoretic equilibria. According to Searle, this thesis is incompatible with his definition because, in his account, constitutive rules cannot be reduced to regulative rules. If Searle is right, then *institution* cannot mean the same in both theories.

If the meaning of a central word differs from theory to theory, or at least possible ways of understanding its meaning are restricted in different and sometimes mutually incompatible manners, then two follow-up problems occur. First, it is no longer clear how two competing theories can be about the same topic. For example, why would a theory according to which atoms are indivisible be about what we nowadays call atoms? Related to this, if two agents endorse two different theories A and B

and talk about a term central to those theories, then it is no longer clear how they disagree. The problem is the same as in the relativist critique of contextualism. If two interlocutor's beliefs are such that a certain word has a different meaning, because they endorse different explicit definitions of it or their beliefs characterize its meaning in substantially different ways, then the semantic contents of their beliefs also differ. So why do they not just talk at cross purposes?

Semantic externalists may reply to these worries that only defending a theory or having opinions cannot directly influence public language meaning. According to the most extreme form of externalism, there is no influence at all. The noun *atom* stands for atoms. Whatever theory of atoms we build and whatever beliefs we hold about atoms does not influence what atoms are. The problem with this view is that it conflates word extension with meaning and consequently does not explain meaning change at all. The meaning of *atom* could only change if atoms change, yet it seems that this meaning has changed over the past centuries. A more realistic form of externalism by Cappelen (2018) acknowledges that word meanings change over time, but not fast and not in a way that is under our control. Instead, meaning change is governed by hard to understand, long-term processes within a large speaker community, based on slowly changing patterns of use. These changes may be triggered by changing world views, theories and opinions of all kind that come to be believed by larger groups of speakers, but not merely by discussions between individual speakers.

This *lack of direct control thesis* is a valid point about public language meaning. However, it does not touch the problem's core. Surely, some sort of meanings are discussed in an explicitly metalinguistic dispute in which words are mentioned. If so, then at least some implicit metalinguistic disputes discussed by Plunkett and Sundell (2013) also have to concern word meanings. After all, any such implicit dispute could be turned into an explicit one at any time simply by mentioning the disputed word instead of (seemingly) using it. Maybe the meanings in such disputes are not meanings of public language expressions, and instead the underlying concepts or the meanings of words of idiolects and sociolects change. For instance, Ludlow (2014) argues with many examples that interlocutors adapt their 'microlanguages' to each other in conversations. So even if one does not buy into the theory dependence problem as a thesis about public language, the problem remains at the level of idiolects and concept systems that differ between speakers, whether or not these coincide with public language.

To illustrate this point, consider two early 19th Century physicists discussing and disagreeing about two wave theories of light that are both derived from Augustin Fresnel's theory of luminiferous aether but differ in various details. Neither the correctness of their theories nor the public language meaning of *aether* should have a substantial bearing on the meanings they associate with the word in the context of this discussion. It remains a problem to explain how they disagree about the same topic and why they are not just talking at cross purposes, if they indirectly characterize aether differently or even use different explicit definitions of *aether*. Likewise, consider two ancient fishers discussing whales. Both agree that whales are fish. One of them argues that they are the largest fish of the sea and being the largest fish of the sea is the whale's defining feature. The other one disagrees and claims to have seen larger fish; he thinks that being a fish with a blowhole is the defining feature of a whale. They have false beliefs about whales, some of which enter their putative definitions, and so

their concept systems cannot represent public language meaning from the perspective of semantic externalism. Nevertheless, one might ask how their disagreement can be spelled out in terms of these flawed concepts, given that their conceptual systems differ with respect to the concepts they erroneously associate with the word *whale*.

3.2 Definitional meaning does not imply an epistemic priority of analyticity

A popular reply to the problems of theory dependence is to reject any definitional account of word meaning. In further support of this position, one might first argue based on externalist arguments by Kripke (1972), Putnam (1975), and Burge (1979) that both public language meaning and thought contents are individuated externally by facts of the shared environment. As a classical example, *water* denotes H₂O because it is a natural kind term whose meaning is fixed indexically by virtue of the fact that water is mostly composed of H₂O. Correspondingly, if someone thinks about water, then the contents of that person's thoughts are also individuated externally. As Putnam's Twin Earth example is supposed to show, thinking about water is not the same as thinking about a colorless, odorless, transparent liquid, for instance.

As a bonus, it appears as if such a form of externalism also fared well with Quine's arguments against analyticity. I will argue below that this is not the case but let us consider the argument first. It seems to be very popular. In a definitional theory of concepts or lexical meaning according to which concepts or word meanings are characterized by the theories (in a broad sense) to which these are central, it seems that law-like statements that are taken to be definitory for a word or concept (whether in individual cognition or as a thesis about public language meaning) would make certain statements analytically true that are not. For example, if the property of being the smallest indivisible building blocks of nature with the characteristic properties of chemical elements takes part of a definition of *atom*, then it seems that *Atoms are indivisible* is analytically true. According to Quine (1951) such a notion of analyticity is ill-conceived and hinges on a notion of meaning, which, in turn, circularly presumes analyticity.

Although this sort of externalism may be appropriate for specific words of a public language in a truth-conditional setting, it comes with too many problems as a general theory of lexical meaning. First, in practice, word meanings are not always indirectly characterized but also sometimes defined explicitly. In that case, the meaning of the word under consideration clearly depends on an underlying theory, namely the one that simultaneously lends credibility and adequacy to the definition and uses it. Moreover, there is a gradual scale between the indirect characterization of word meaning and explicit definitions. Often a word is used in ways that amount to defining without making the definition explicit. One may define what *triangle* means more or less precisely, or one may understand it more intuitively based on examples. It is hard to say where the supposed externalist individuation starts and where it ends. Semantic externalists have mostly only provided convincing stories for everyday nouns for empirical objects such as *water*, *tiger*, and *pencil*, and their accounts remain mysterious for words like *democracy*, *dark matter*, *triangle*, and *institution*.

Second, as mentioned above, the rationality of metalinguistic disputes becomes questionable without a definitional approach to word meaning in idiolects (or, theories) and concepts. If these are externally individuated, then why and how could they be disputed? Notice that even though [Plunkett and Sundell \(2013\)](#) argue that metalinguistic disputes can be substantive, some of them are also not substantive. Suppose John and Mary argue about what counts as a chair, and after a while, they agree that stools with only one leg ought not be called *chairs*. If the meaning of *chair* was externally individuated, then this whole discussion would be irrational and pointless. However, although it may be pointless and not substantive, it is clearly rational and concerns the question of which minimal number of legs has an adequately definitory quality for *chair*.

We frequently dispute word meanings and propose various characterizing properties, which are derived from, and relative to, a supporting background theory. It is hard to see how this practice could be based on a systematic error. This does not mean that we should not embrace externalism, it means that we should embrace externalism *and* the theory dependence of word meaning and concepts. The underlying theories are hopefully about reality and not just about figments of our mind. Nevertheless, within each theory words may get their meaning relative to that theory by indirect characterization or explicit definition.⁸ Indexicalist externalism makes sense for a limited number of natural kind terms because the underlying theories are particularly well-confirmed. It does not scale to theories about more contentious topics.

What about the analyticity objection, then? None of Quine's points against analyticity show in my opinion that there is something wrong with a definitional theory of word meaning and concepts. The lesson to learn from [Quine \(1951\)](#) is rather to be careful not to give epistemic priority to any allegedly analytic inference. For even if we appear to arrive at certain conclusions solely by word meaning, this is never the case. From the present point of view, there is no such thing as 'the' meaning of a word. Words get their meanings relative to the theories in which they are used. If such a theory is based on empirical evidence, then whatever we believe in having derived solely on the basis of word meaning hinges on the adequacy and merits of the theory and its supporting evidence. As a pragmatist naturalist, Quine believed that any theory is revisable and needs to be judged on its scientific merits (in proper scientific contexts). Even mathematics is revisable in that sense. From that perspective, seemingly analytic judgments are theory-relative and revisable like any other judgment. If, contrary to this, there was a non-theory dependent word meaning, then analytic judgment could have some epistemic priority. However, according to Quine, any such meaning would be a dubious stipulation and presume an equally dubious notion of analyticity. We cannot attribute any epistemic priority to inferences seemingly derived only on the basis of word meaning because according to the Quine/Duhem Thesis the underlying theories are confirmed or falsified holistically.

From all of this it follows that it is possible to consistently deny the usefulness of analyticity as an epistemic notion without giving up definitional word meaning and the thesis that theories characterize the meaning of words that play a central role in them. We may even continue to speak of analytic judgments (although Quine would not endorse this), as long as no special epistemic priority is given to them. For example, it is perfectly fine to contemplate whether *bachelor* means *unmarried man* or

whether additional conditions need to be met, and it may even be true that under the first definition every bachelor is an unmarried man and vice versa. Talk like this is fine, as long as one keeps in mind that such considerations tell us nothing about the adequacy of that definition and its underlying theory, about the existence of bachelors and unmarried men, and about what other properties bachelors might have. The truth of the analytic statement hinges on the confirmation or falsification of the supporting theory.

If a complete rejection of explicit definitions and implicit characterizations of meanings is implausible for idiolects and concept systems, and if the theory dependence of word meaning and concepts remains compatible with Quine's arguments against analyticity, then the problems mentioned above cannot be ruled out that easily. When are two theories about the same topic or concern the same central words? How can two speakers advocating competing, mutually incompatible theories or world views be said to disagree and talk about the same things?

4 Tackling the Problems

There are several ways to tackle the problem of topic equality of theories. First, there are good reasons to assume we associate some minimal meanings with expressions that are not necessarily truth-functionally complete and represent 'everyday', common-sense word meanings. Rast (2017b,a) suggests the term *core meaning* for these and contrasts them with *noumenal meaning*, which represents what a word *really* means according to our current best understanding and theorizing. For instance, even speakers in the past who were not in a position to know that water consists of H₂O associated with it the core meaning of being a transparent, drinkable liquid. Likewise, we can recognize animals by the way they look under normal circumstances. The core meaning of *whale* is to look like a whale. So if two people disagree about the noumenal meaning of a word, for example, whether whales should be classified as fish or mammals, they may continue to talk about the same topic as long as they sufficiently agree about the associated truth-conditionally incomplete core meaning. Second, Rast (2020) lays out that words can also be associated with measurement operations. Competing theories are about the same topic if associated measurement operations (which may differ across agents and theories) roughly pick out the same extension.⁹ Third, unless a noun is further qualified and distinguished from other uses, the same noun in two different theories A and B is supposed to stand for the same kind of entities in both A and B. Certain words, usually nouns plus qualifying adjectives, act as fixed points around which varying theories are constructed. Choosing the same words for such alleged fixed points tells speakers that two theories are supposed to be about the same kind of entities. This nominal topic equality is a fallible stipulation, but supporting theories are fallible, too, and in a sense also mere stipulations.

The most important mechanism is measurement because measuring roughly the same kind of entities warrants topic continuity; the others do not warrant but rather stipulate it. Taken together, these three mechanisms suffice to explain putative and real topic continuity. However, having an account of topic equality does not solve the problem of a potential drift of word meaning and concepts across speakers. How

do we understand each other, if our background theories, opinions, and world views differ from each other and influence lexical meaning? I believe that the best answer to this question is twofold. First, as argued in the next section, it is only pressing when global meaning and concept holism is assumed. Instead, we should embrace local meaning and concept holism. Second, at least up to a certain degree we are able to, and *have* to be able to, track and entertain different opinions, world views, and theories without endorsing them. Hence, theory dependence is less of a problem for mutual understanding than one might think at first glance. This topic is addressed in Section 4.2.

4.1 The case for local holism

Holism is best understood in opposition to atomism and the arguments that speak against it. A central thesis of semantic atomism is that the meanings of simple words are not generally composed of other words' meanings. A reasonable semantic atomism may acknowledge that there are more complex, morphologically derived words whose meanings are composed out of their parts' meanings. For instance, *consequential* may have a primitive meaning whereas *inconsequential* may have a meaning derived from the former. However, this must be limited to complex words. If the meanings of all words are decomposable into logical combinations of the meanings of other words, then the meaning of every word hinges on the meaning of those other words, which is a form of holism. So the semantic atomist has to assume primitive, non-decomposable meanings, or that simple words have no meaning at all and only serve as syntactic anchoring points in a computational theory of cognition, or—as the more common, externalist response—allow talk about 'meaning' only in a derived sense, for example by assuming that the extensions of simple words individuate their meanings. One form of semantic atomism can be found in works by Fodor (1975, 1987), while Fodor and Lepore (1992) thoroughly discuss arguments against holism without presuming Fodor's contested theory of cognition.

Atomism would provide an elegant solution to the problems of theory dependence if there were not such good counter-arguments against it, whether it concerns public language, idiolects, or concepts. First of all, if the meaning of a word is primitive, then how does it change? This is a generalized form of the earlier argument against indexicalist externalism. Word meanings change within discussions when they are defined explicitly. Consequently, they should also sometimes change when they are characterized implicitly. Likewise, concepts such as the concept of holy trinity can change over time even if they do not match the established current public language word meaning. If they can change over time because our conception of reality changes, then it seems equally reasonable to assume that they can also differ synchronously when different theories of reality are considered, defended, and supposed. There *are* metalinguistic discussions.

Atomists have a hard time explaining such negotiated concept and meaning changes because they do not allow for the logical decompositions of lexical meanings under dispute in metalinguistic discussions. Atomism is also questionable from the perspective of the inferences that can be drawn from word use. Suppose a fixed number of words has a primitive meaning that cannot be further dissected. Suppose α is such a

word. This word α will have one set of consequences relative to theory A and may have another set of consequences relative to theory B . Shouldn't at least some such consequences count as an aspect of the word's meaning? It is hard to see how these different consequences could never be the result of different meanings. Another point against atomism is that some seemingly substantive theses can be turned into explicitly metalinguistic theses and vice versa, and the difference between them is only whether speakers quote linguistic material or not; some ways of talking are even in-between the two. Consider the following examples:¹⁰

- (3) a. Every atom is indivisible.
- b. Atoms are indivisible.
- c. Being an atom entails being indivisible.
- d. Being indivisible is a defining feature of atoms.
- e. An essential aspect of the meaning of *atom* is that they are indivisible.
- f. *atom* means *being a smallest indivisible building block of nature with the characteristic properties of a chemical element*.

Implicit to semantic atomism is the claim that examples like (3a)–(3d) do not concern the meaning of *atom*. Is this really plausible? Although only (3e) and (3f) explicitly mention words, in ordinary conversations the dependence on the natural language is often irrelevant, and all of the above statements characterize atoms in similar ways. In practice, we often define words without mentioning them at all. Even a simple use of a generic like in (3b) can have a 'metalinguistic flavor' in a context where a characterization, explanation, or definition of a word is expected. Neither is an explicitly metalinguistic definition like in (3f) arbitrary, nor does the use/mention distinction clearly indicate whether a word is defined or characterized, or whether a world-level claim is made. This does not mean that the choice between explicit definition and indirect characterization is unimportant or that every law-like statement in which a word is used has a definitory quality for that word. An explicit definition may indicate particular methodological preferences, that it is only conventional or operational and later to be revised, or that a term is theoretical. Nevertheless, semantic atomism presumes a too large divide between the explicit definitions of complex words and the meaning of supposedly primitive words. There is no such gap in practice.

So if we reject semantic atomism, how can holism deal with the problem of theory dependence? To answer this question, holism has to be characterized in more detail. First, holism can apply to words in public language. In this view, the meaning of a public language word depends on and is partially constituted by the meanings of other public language words; if the meaning of a word α changes, then the meanings of words change that are partly constituted by the meaning of α . This is semantic holism as the counterpart to semantic atomism. An analogous thesis may be formulated for idiolects and sociolects, which we may call meaning holism in general. Finally, concept holism concerns individual agents' concept systems, where a concept is a meaning-like representation that is not necessarily associated with a word. For instance, a sculptor

may have a concept for a particular shape, may be able to recognize it and use it while sculpting, without naming it and without there being a name for it in public language. Concept holism is the thesis that a concept c changes whenever other concepts change that partially constitute c .

Since some externalists deny that concepts exist, and it is also controversial whether public language meaning can change in the way relevant for the theory-dependence problem, I will focus in the following discussion on idiolectal meaning holism and for simplicity sometimes abbreviate it as *holism*. What can be said about this type of holism can also be said about the others. The focus shall also be on meaning change. Most of what can be said about meaning change can be transferred to the case when two agents disagree. The main difference between the two cases is that two agents may also differ in *other* beliefs that are peripheral or unimportant to the theories under consideration. This complicates matters, but not in a way relevant for what follows.

Consider the theory change scenario. An agent endorses a theory A but then for some reason starts to suspect that A is not the right theory and endorses theory B instead. Provided that A and B are not compatible with each other (they cannot simultaneously be true), the agent first has to retract A from his total belief base K and then integrate B into K . Although there are well-established formal theories for modeling these kind of processes such as AGM belief revision (Alchourrón et al., 1985) and KM update (Katsuno and Mendelzon, 1992), realistically speaking only some aspects of theory change can be modeled formally. The process is inherently creative and involves theory discovery of B . The retraction of A might not be minimal, it may be based on a shift in perspective and a massive re-evaluation of more beliefs in K than merely those required for A . Therefore, we cannot assume that those beliefs in K which are *prima facie* independent from A remain constant during such a revision. In any case, however, some statements involving words used in both K plus A and K revised by B will likely have different consequences before and after revision. If the effect can be isolated to only one word, then this word's inferential meaning has changed. Whether we are willing to say that its purported idiolectal meaning has also changed depends on whether the law-like statements responsible for the inferential meaning change count as attributing a definitory quality or are of a more accidental character.

For the current purposes, two versions of holism have to be distinguished. According to global holism, whenever the meaning of one word changes relative to a belief base, then the meaning of *all other* words changes, too. Analogously, in the two-agent case, every agent associates a slightly different meaning with each word in their idiolect, or they have slightly different concept systems. Why would this be the case? Generally, the ideas behind this position is that words are only meaningful in larger units like sentences and discourse fragments, and that their meaning has been learned and is indirectly constituted by the network of law-like semantic relations and constraints between words. As Lepore and Fodor (1993) put it, "...meaning holism says that what the word 'dog' means in your mouth depends on the *totality* of your beliefs about dogs, including, therefore, your beliefs about whether Lincoln owned one. It seems to follow that you and I mean different things when we say 'dog'; hence that if you say 'dogs can fly' and I say 'dogs can't fly' we aren't disagreeing." (Lepore and Fodor, 1993, p. 638) Correspondingly, each concept in a concept system depends

on other concepts in this view, and no two agents can learn and internalize exactly the same concept.

Arguments by [Davidson \(1967, 1973\)](#) are sometimes advanced in support of global holism. As a twist on [Quine \(1960\)](#), Davidson suggests to define truth-conditions for a language by Tarski-sentences of the form '*S* is true in language *L* iff. *T*', where *T* specifies the truth-conditions for the sentence mentioned on the condition's left hand side.¹¹ In a radical interpretation situation, when a speaker of *L* would utter a sentence *S* and we have to figure out what this utterance means, we have to apply the Principle of Charity and assume that this speaker's beliefs are mostly true.¹² Based on this assumption, we can make sense of another person's rationality in a radical interpretation situation by attributing beliefs and desires to that person and associating them with our assumptions about what their utterances mean. However, the Principle of Charity can only get one so far. Since the beliefs of the interpreter and the interpreted person only roughly converge, understanding of the other person's language will only ever be a rough approximation in a radical interpretation situation. Moreover, since beliefs depend on each other just like the statements of a theory, the recovery of the other person's language in a radical interpretation situation seems to imply global holism even when the Principle of Charity is applied.

In contrast to global holism, local holism is the position that a meaning change of a word may trigger some finitely many meaning changes but that this does not imply that the whole idiolect changes. For example, suppose John calls any apple or pear an *apple*. He has a persistent misconception that pears were once similar to peaches but have long gone extinct. John later learns about pears and how to distinguish apples and pears by taste and shape like most speakers of English. The change affects *pear* and *apple* in John's idiolect, as well as the concepts of being a pear and being an apple. The incorrect pear concept is eliminated, and a more adequate one is internalized. In terms of theories, we may say that John learns better pear and apple theories. According to local holism, this change might affect related concepts and word meanings such as the meaning of *apple pie* (it's not the same as a pear pie), *juice* (pear juice exists), and *fruit* (pears are fruits, they are not extinct, and taste such-and-such). It will not affect every other word, though. For example, John's idiolectal meanings of *and*, *relation*, *democracy*, *dog*, and *greater than* are not affected. They are not just affected in a barely noticeable and neglectable way. They are not affected at all.

Theory dependence is a huge problem for the global holist. Since people have different beliefs about all kinds of topics, and every difference of beliefs leads to differences in idiolects and corresponding concept systems, even with a generous application of the Principle of Charity two interlocutors will likely talk at cross purposes and fail to fully understand each other. The farther the theories they endorse are apart from each other, the less they understand each other when discussing a topic common to those theories. So it seems at first glance. On a closer look, however, it turns out that the arguments for global holism are relatively weak. There are good reasons for rejecting global holism and accepting local holism instead.

[Fodor and Lepore \(1992\)](#) lay out in detail why many of the arguments for global holism based on [Quine \(1960\)](#) and [Davidson \(1967, 1973\)](#) are not conclusive. One of their points is that language learners and field linguists are never in a radical interpretation situation.¹³ The environment is shared, the agents' cognition works in similar ways,

and inadequate interpretations of utterances can be corrected over time. Speakers also share common features of their perceptions. For instance, a child learning the word *rabbit* from watching a living rabbit sees a rabbit and not rabbit slices like in the famous Gavagai example of Quine (1960). Radical interpretation scenarios are radically skeptic from an epistemic point of view, but successful language learners are not and cannot be radical skeptics. Other arguments by Fodor and Lepore (1992) also undermine the support that radical interpretation and the Quine/Duhem thesis seem to lend to global holism, but addressing them here would go beyond the scope of this contribution.

There is one positive logical argument against global holism that Fodor and Lepore do not endorse. When speakers adapt idiolectal meaning and related concepts to one another, only word meanings and concepts *central* to a given topic need to be revised. For example, children who learn what pears are and how they differ from apples only need to revise fruit- and nutrition-related concepts. There is no need or reason in such a case to revise unrelated concepts like being a tire or being a tiger. There are essentially two reasons for this locality of revisions and why centrality is not an arbitrary stipulation in this context.

On the one hand, the common-sense ontologies encoded by concept systems are hierarchical. An upper ontology represents very abstract concepts such as relations, counting, mereological notions, physical versus abstract objects, physical movement, processes, information transfer, and so forth. In contrast, a lower ontology represents specific knowledge about the world. A change of the beliefs that constitute the lower ontology is unlikely to require a change of beliefs that constitute the upper ontology in a reasonable account of belief revision and theory discovery. On the other hand, theories about specific topics, identified by associated measurement operations, are discernible from other theories and the more general ontology. A lower ontology is divided horizontally into parts that are mostly or entirely independent of each other from a logical perspective. For instance, there are many (onto-)logical relations between tires and pears and these objects can interact in many ways, but beliefs about these relations are regulated by the upper ontology. They might be based on the fact that both are types of manipulable physical objects that can be carried and moved, for example. A revision of the pear concept by integrating new pear and fruit theories does not have to trigger a revision of the tire concept, and likewise for the idiolectal meanings of *pear* and *tire*. So even though there are logical relations between pears and tires, neither is *pear* central to the tire theory nor, vice versa, *tire* central as a term in pear theories.

Although developing a full-fledged account of centrality as a measure of the nearness of terms to the measurable topics of a theory would be a major undertaking, there can be no doubt that *pear* is not just psychologically but also logically nearer to *apple* than *tire* is. Words whose meanings are directly related to each other by law-like statements at the same level of ontological specificity and within the same theory with measurable topics are close to each other, for instance, whereas words whose meanings are characterized in a theory about other measurable topics and whose meanings are only related to each other via law-like statements of the upper ontology (less specific, more abstract) are more distant from each other.¹⁴

Anyone who accepts these kind of examples and the reasoning behind them ought to be wary about Quine's dictum that "[t]he unit of empirical significance is the whole of science." (Quine, 1951, p. 39) Individual theories can be confirmed and rejected without revising other theories, let alone all of science, and changing individual theories need not trigger revisions of the upper ontology that supports them. Confirmation holism is only local. As a consequence of this position, under the indirect meaning characterization thesis and the assumption that word meaning is (at least sometimes) definitional, it follows that a change of idiolectal meanings and concepts only affects words and concepts closely related to the one that changes. Further changes may be triggered, but these are usually local, too, since the underlying common-sense ontology is divided vertically and horizontally.

4.2 Tracking theories

So far, we have talked about beliefs and endorsing theories, and the concept systems and ontologies related to these beliefs. But how are these notions related to each other? The way I understand beliefs in this article, these are types of attitudes that we attribute to agents *de re*, using belief ascriptions of public language. For instance, John who calls both apples and pears *apples* does not believe *de re* that pears do not exist. Maybe he believes *de dicto* that pears don't exist because he is disposed to utter sentences like *Pears don't exist any longer*, but no corresponding *de re* belief can be attributed.

In contrast to *de re* belief, concepts can be described using public language but do not necessarily correspond to words of public language or an agent's idiolect. For instance, when John considers every pear an apple he possesses a primitive apple-pear concept. If he uses the word *apple* to refer to apple-pears (i.e., apples *or* pears), then the idiolectal meaning of *apple* is for him: being an apple-pear.

Other concepts regulate the relations between concepts and, taken together with the concepts they regulate, constitute a concept system. The ontology that corresponds to such a concept system can be described by the embedded sentences we would use when ascribing corresponding *de re* beliefs. Hence, in this way of talking, endorsing a theory can be described as the revision of existing beliefs by a theory. The point of the previous section was that even though this process may affect more beliefs than just those constituting the theory that is replaced, from a logical point of view the ontology constituted by corresponding background beliefs is vertically and horizontally divided into parts, and theory revisions will not generally affect all of an agent's beliefs. An indirect consequence of this view is that an agent's concept system is also usually only affected locally. Thus, we should opt for local holism and the problem of theory dependence becomes less pressing.

However, this picture is not complete. Talking about beliefs can only be understood as a first approximation. We not only endorse theories, we also consider them, suppose them, and deal with them in many other ways that do not imply that an agent fully believes them. This is another important point for explaining meaning disputes.¹⁵

Consider two agents having a dispute that indirectly concerns word meaning. Speaker *g* endorses theory *A* and *h* defends theory *B*, which are both about the same topic with associated measurement operations. Regardless of what has been said in the

previous section, under the local holism thesis the two speakers will misunderstand each other if the idiolectal meaning of a term α differs relative to g 's belief base plus A from the idiolectal meaning of α according to h 's belief base plus B , provided that α is central in one of the theories and some of the inferential meanings in which α differ between h and g have a definitory quality for at least one of the agents, i.e., the speaker considers them constitutive for what it means to be (rightly) called an α . How can the speakers then understand each other? Are they not still only talking at cross purposes insofar as α is concerned?

As hinted above, the answer to this question is that a theory does not have to be endorsed to create mutual understanding. Instead, speakers can consider a theory, and this ability suffices to rule out talk at cross purposes under ideal circumstances. For g to understand theory B , she only needs to consider B 's merits on the basis of a hypothetical revision with B , but need not integrate theory B and thereby give up A . Instead of endorsing other persons' theories, we *track* them. However, it does not stop there. We may also track a theory by hypothetically revising by this theory what we assume that the person(s) who defend the theory believe, i.e., based on our assumptions about the proponents' concept systems.

Even this description is incomplete. As even a cursory look at our practices reveals, humans have the astonishing ability to compartmentalize theories altogether, independently of whether these are endorsed or not. Even if a revision is not hypothetical and a new theory is endorsed, this does not necessarily induce a change of the remaining common-sense ontology. For example, physical theorizing could have triggered radical changes in the everyday concept systems of physicists. After all, time and space are no longer constant in modern physics, and quantum mechanics also has radical implications about the macrophysical world. Nevertheless, the radically different ways modern physics looks at nature have probably not changed physicists' common-sense ontologies in any substantial way. Instead, they can designate an area of 'theoretical physics' in which physical theories revise the background ontology, but this area is compartmentalized from the original common-sense ontology that stays in place.

This compartmentalization is necessary and inevitable. First, sometimes two theories are worth endorsing even though there are good reasons for believing that they are incompatible. This point is particularly important since two theories can be incompatible with each other even when they are not about the same topic. As a typical example, many theoretical physicists believe that Einstein's Theory of General Relativity and Quantum Mechanics are not compatible with each other and that some more general theory will replace them in the future. Nevertheless, it is perfectly rational to endorse both theories at the same time. They are well-confirmed even though they cannot be combined easily. In this case, physicists endorse both of them until a better, more unifying framework has been found. Similarly, it would be incorrect to claim that physicists do not endorse Newtonian Mechanics; they do, they are merely aware that it does not provide accurate descriptions of objects moving at near light speed and does also not describe the behavior of extremely small 'objects.' Physicists endorse Newtonian Mechanics although Relativistic Mechanics can replace it entirely. It is not necessary to use the more complicated relativistic formulas for macrophysical objects at very low speeds.

Second, it is not irrational to retain information, even when it does not meet the requirements for being fully integrated into one's belief system. Whether it is worth and rational to retain a new theory (opinion, world view) may be a complicated matter, but the decisive criterion cannot be that it meets the requirements for being endorsed. Otherwise, learning inductively by corroborating evidence from different information sources would be impossible, for instance, in a scenario where each information source individually does not meet those requirements. A rational agent needs to keep track of theories and supporting evidence that do not meet the standards for being endorsed.

Third, the standards for endorsing theories are also context-sensitive. For instance, it is rational for an agent to endorse a scientifically well-confirmed theory if that agent is not very knowledgeable about the theory's subject matter and domain. Identifying experts and relying on them is an important skill for any rational epistemic agent, since learning is largely a social process. However, it is equally rational for another agent, who is knowledgeable about the theory's subject matter and domain, not to endorse the same theory and merely to consider it. A theory worth endorsing on one occasion may only be worth being aware of in a more skeptical context. It can even be rational to consider or track a theory in one context and completely ignore it in another. For instance, a certain amount of knowledge about religious texts and opinions is needed to understand the world views and motives of religious fanatics. This does not mean that the same knowledge needs to have any influence on one's own world views or needs to play a role in the evaluation of scientific evidence.

Tracking theories means keeping their origins and sources in mind and knowing them well enough for understanding others; it does not imply endorsing them in any way. Hence, the contextualist objection of talking at cross purposes is ill-conceived for theory-based disagreement. Theories neither need to be compatible with each other nor do they need to be co-tenable, believed, or endorsed by speakers in order for them to disagree about them. It is entirely possible to rationally disagree about an aspect of a theory, opinion, or world view that neither of the interlocutors endorses.

When we take into account this ability to compartmentalize and track theories, it is reasonable to also assume that we can deal with the semantic effects of theory dependence under the assumption of local holism. Take the much-discussed *Secretariat is an athlete* example from Ludlow (2008, 2014), for instance. Secretariat was a famous racing horse. Suppose John believes that *athlete* can only be used to denote humans. In his view, part of the definitory properties of athletes is being human. Mary disagrees with him and believes that horses can be athletes, too. Even though a prototypical athlete might be human, only physical prowess and success in competitions are defining characteristics for athletes. Their disagreement is discussed by Plunkett and Sundell (2013) as a typical case of an (implicit) metalinguistic dispute.

Nevertheless, Mary and John can understand each other if they manage to keep track of each other's opinions about athletes. If each of them presupposes a different meaning of *athlete* in their idiolect, this does not automatically lead to misunderstanding. It only leads to a linguistic misunderstanding when one of them does not know the other's opinions about athletes well enough, and does not keep track of the other's athlete theory. Normally, however, speakers are able to keep track of other theories at least to some extent, which includes an ability to recognize the effects of local holism on possible candidates for word meaning. To what extent? From an idealized modeling

perspective, precisely to the extent to which their model of the other's theory about a given topic in the conversation and words central to it matches the other's actual theory.

5 Summary

This article started with an overview of linguistic context dependence. I argued that parameterized contexts do not suffice to represent linguistic context dependence adequately. However, combined with the parameterized context dependence of indexicals and tenses, regarding interpretation as an inference to (usually) more specific semantic contents while presuming the semantic underdetermination thesis leads to a fairly complete account of linguistic context dependence. This is possible only if moderate contextualism is the right position. I have suggested that this is the case because the number of linguistically-regulated, rule-based context-sensitive phenomena in natural languages is finite.

The problem left open by such an approach is the theory dependence of lexical meaning. This theory dependence does not need to occur at the level of public language meaning to become a problem; it also creates difficulties for explaining mutual understanding at the level of idiolectal meaning and concept systems. Although the easiest way to address the problem is by rejecting a definitional approach to meaning and concepts, I have rejected this solution because it creates numerous problems. It does not match the reality of overt and implicit metalinguistic discussions and is forced to draw an inadequately sharp divide between definitions and externally individuated meanings. Semantic atomism can evade this problem, but may not be plausible for other reasons. Especially the meanings of words for abstract objects, complex verb phrases, and compound nouns are hard to explain from the perspective of a stringent semantic atomism. However, the problem of theory dependence remains pressing for semantic holism.

In the final part of the article, it was argued that solving the problem of theory dependence requires two theoretical commitments. First, global holism needs to be given up in favor of the overall more plausible local holism. Since common-sense ontologies are horizontally and vertically divided into parts, the effects of theories on idiolectal meaning are often isolated to these parts. Endorsing or rejecting a theory does not influence all concepts or the meaning of all words in an idiolect but only a select few central to the theory. Second, theory representations of rational agents need to be compartmentalized, as rational agents need to track theories incompatible with their beliefs without endorsing them. If this is true, then it is also reasonable to assume that the requirements of rational theory compartmentalization allow speakers to compartmentalize the effects of theory dependence on concept systems and meaning. An ideal rational speaker would be able to keep track of all theories and information sources in a way that takes into account shifts in lexical meaning due to the different law-like statements with definitory qualities for concepts and expressions that these theories support. Humans are not ideally rational in this sense, of course, yet it is reasonable to assume that they can keep track of someone else's definitions and characterizations in the same way they can keep track of their own theories. Sometimes

they succeed, and then there is no misunderstanding, and sometimes they fail, and there will be talking at cross purposes.

Assembling the parts of the article leads to the following picture. Hearers interpret semantically incomplete content by drawing inferences from it based on existing epistemic representations. A model for this process may be Gricean, an inference to the best explanation, or a more psychological account like Relevance Theory. The belief base relative to which utterances are interpreted need not solely consist of the interpreter's beliefs and endorsed theories. An interpretation may also be based on assumptions about the respective speaker's beliefs and theories (opinions, world views). Endorsing what the speaker said, as well as the underlying theories that need to be presumed in order for the utterance to be believable to be true, is then a second step. This step may require a revision of the interpreter's theories about the topic and a corresponding change of the interpreter's idiolect, adapting to the speaker's idiolect in that respect. However, both world-level and metalinguistic disagreement is possible without this second step.

Notes

¹There are also non-indexical uses of absolute tenses, such as the use of the English present tense in a generic like *Cats are mammals*.

²Cf. Perry (2001) on incremental truth-conditions.

³See Schlenker (2000, 2003) on shifting first person pronouns, and, more generally, Fillmore (1997), Lyons (1977, p. 579), Levinson (1983, p. 64).

⁴See, for instance, Kaplan (1989), Lewis (1979b), Cresswell and von Stechow (1982), von Stechow (1984).

⁵Cresswell (1996) argues that modal logics with finitely many parameters are not expressive enough to deal with the indexical context dependence of arbitrarily long sentences. Instead, full quantification over reified contexts is needed. This argument has largely been ignored by the philosophy of language community.

⁶See Jackendoff (1987, p. 381/2), Jackendoff (1990, pp. 189–194).

⁷Since the discussion in what follows mostly concerns lexical meaning, *word* is used for the linguistic entities under consideration. These are usually nouns (general terms) in examples, but for brevity *word* is also used in a looser sense as a shortcut for linguistic expressions in general. This may include compound nouns, nouns with participial phrases, noninflected verb phrases, and phraseologisms, for instance.

⁸By mentioning indirect characterizations and explicit definitions in this way, I do not want to presume that these are unique phenomena. There is not only a gradual transition between them, they are also umbrella terms for many different, yet related practices such as stipulating meaning postulates, operational definitions, definition as abbreviation, definition by example, definition by systems of axioms, providing prototypical information, specification, abstraction, and so forth.

⁹Without emphasizing measurement operations, Cappelen (2018) also advocates such an extensional notion of topic equality.

¹⁰Clearly, there are two uses of *atom* in the text. The examples are about physical theories, whereas *atomism* suggests a mereological use. Which one is meant is clear from the context. This is not another example of theory dependence but merely a case of ambiguity.

¹¹See Davidson (1973, p. 318).

¹²See Davidson (1967, pp.312–313), cf. Davidson (1973, pp.323–324).

¹³See (Fodor and Lepore, 1992, pp. 73–80).

¹⁴One approach would be to base the account of centrality on a good account of theory revision, which, in turn, would have to take into account theories and their associated measurement operations as units when modeling epistemic entrenchment. Since there is no non-psychological 'logic of theory discovery', however, even such an elaborate approach would remain limited.

¹⁵Endorsing a theory in this context is understood roughly as believing what the theory states. There could be attitudes other than belief at play, for example, a true-holding attitude with less epistemic entrenchment

than belief. As long as it can be attributed de re, this does not impact my position. Resorting to belief should be taken as a simplification.

References

- Alchourrón, C. E., Gaerdenfors, P., and Makinson, D. (1985). On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic*, (50):510–530.
- Asher, N. and Lascarides, A. (2003). *Logics of Conversation*. Cambridge University Press.
- Bach, K. (2005). Context ex machina. In Szabó, Z. G., editor, *Semantics versus Pragmatics*, page 16–44. Oxford UP, Oxford.
- Bar-Hillel, Y. (1954). Indexical Expressions. *MIND*, 63:359–76.
- Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. MIT Press, Cambridge, Mass.
- Burge, T. (1979). Individualism and the Mental. *Midwest Studies in Philosophy IV*, page 73–121.
- Burks, A. (1949). Icon, Index and Symbol. *Philosophical and Phenomenological Research*, 9(4):673–689.
- Cappelen, H. (2018). *Fixing Language*. Oxford University Press.
- Cappelen, H. and Lepore, E. (2004). *Insensitive Semantics*. Blackwell, Oxford.
- Cappelen, H. and Lepore, E. (2005). A tall tale: In defence of semantic minimalism and speech act pluralism. In Preyer, G. and Peter, G., editors, *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford University Press.
- Carnap, R. (1950). Empiricism, semantics, and ontology. *Revue Internationale de Philosophie 4 (1950)*, 4:20–40.
- Castañeda, H.-N. (1967). Indicators and quasi-indicators. *American Philosophical Quarterly*, 4:85–100.
- Castañeda, H.-N. (1989a). Direct reference, the semantics of thinking, and guise theory (constructive reflections on david kaplan’s theory of indexical reference). In Almog, J., Perry, J., and Wettstein, H., editors, *Themes from Kaplan*, page 105–144. Oxford University Press, Oxford, New York.
- Castañeda, H.-N. (1989b). *Thinking, Language, Experience*. University of Minnesota Press, Minneapolis, Minn.
- Chalmers, D. (2006). Two-dimensional semantics. In Lepore, E. and Smith, C. B., editors, *Oxford Handbook of the Philosophy of Language*, chapter 23, page 574–606. Oxford University Press, Oxford.

- Cresswell, M. J. (1996). *Semantic Indexicality*. Kluwer, Dordrecht.
- Cresswell, M. J. and von Stechow, A. (1982). De re belief generalized. *Linguistics and Philosophy*, 5(4):503–535.
- Davidson, D. (1967). Truth and meaning. *Synthese*, 17(3):304–323.
- Davidson, D. (1973). Radical interpretation. *Dialectica*, 27 (1973):314–328.
- de Sa, D. L. (2008). Presuppositions of commonality: An indexical relativist account of disagreement. In Garcia-Carpintero, M. and Kölbel, M., editors, *Relative Truth*. Oxford University Press, New York.
- de Sa, D. L. (2009). Relativizing utterance-truth? *Synthese*, 170(1):1–5.
- Fillmore, C. J. (1997). *Lectures on Deixis*. CSLI Publ., Stanford. Revised version; first publ. 1972 based on a lecture 1971 in Santa Cruz.
- Fodor, J. (1975). *The Language of Thought*. Harvard University Press, Cambridge MA.
- Fodor, J. (1987). *Psychosemantics*. MIT Press.
- Fodor, J. and Lepore, E. (1992). *Holism: A Shopper's Guide*. Blackwell Publishers, Oxford / Cambridge.
- Ginzburg, J. (2012). *The Interactive Stance*. Oxford University Press, Oxford.
- Guala, F. (2016). *Understanding Institutions: The Science and Philosophy of Living Together*. Princeton University Press, Princeton and Oxford.
- Jackendoff, R. (1987). The status of thematic relations in linguistic theory. *Linguistic Inquiry*, 18(3):369–411.
- Jackendoff, R. (1990). *Semantic Structures*. MIT Press.
- Kamp, H. and Reyle, U. (1993). *From Discourse to Logic*. Kluwer.
- Kaplan, D. (1989). Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals. In Almog, J., Perry, J., and Wettstein, H., editors, *Themes from Kaplan*, page 481–564. Oxford University Press, Oxford, New York.
- Katsuno, H. and Mendelzon, A. O. (1992). On the difference between updating a knowledge base and revising it. In Gärdenfors, P., editor, *Belief Revision*, page 183–203. Cambridge University Press, Cambridge, UK.
- Kripke, S. A. (1972). Naming and Necessity. In Harman, G. and Davidson, D., editors, *Semantics of Natural Language*, page 253–355. D. Reidel Publishing Co., Dordrecht, Boston.
- Lasersohn, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28(6):643–686.

- Lepore, E. and Fodor, J. (1993). Précis of holism: A shopper's guide. *Philosophy and Phenomenological Research*, LIII(3):637–640.
- Levinson, S. C. (1983). *Pragmatics*. Cambridge University Press, Cambridge, UK.
- Lewis, D. (1979a). Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8:339–359.
- Lewis, D. K. (1979b). Attitudes De Dicto and De Se. *Philosophical Review* 88, (4):513–543.
- Lewis, D. K. (1980). Index, Context and Content. In Kanger, S. and Öhman, S., editors, *Philosophy and Grammar*, page 79–100. D. Reidel Publishing Co., Dordrecht, Boston.
- Ludlow, P. (2008). Cheap contextualism. *Philosophical Issues*, 18:104–129.
- Ludlow, P. (2014). *Living Words*. Oxford University Press.
- Lyons, J. (1977). *Semantics*, volume 1. Cambridge University Press, Cambridge, UK.
- MacFarlane, J. (2007). Relativism and disagreement. *Philosophical Studies*, 132(1):17–31.
- MacFarlane, J. (2008). Truth in the garden of forking paths. In Carcía-Carpintero, M. and Kölbel, M., editors, *Relative Truth*, page 81–102. Oxford UP, Oxford.
- MacFarlane, J. (2009). Nonindexical contextualism. *Synthese*, 166:231–250.
- MacFarlane, J. (2012). Relativism. In Fara, D. G. and Russell, G., editors, *The Routledge Companion to the Philosophy of Language*, page 132–142. Routledge, New York.
- MacFarlane, J. (2014). *Assessment Relativity: Relative Truth and its Applications*. Clarendon Press, Oxford.
- McTaggart, J. E. M. (1908). The unreality of time. *Mind*, 18:457–74.
- Moore, G. E. (1903). *Principia Ethica*. Cambridge University Press, Cambridge.
- Perry, J. (1977). Frege on Demonstratives. *Philosophical Review*, 86:474–497.
- Perry, J. (1979). The Problem of the Essential Indexical. *Noûs*, 13:3–21.
- Perry, J. (1998a). Myself and I. In Stamm, M., editor, *Philosophie in Synthetischer Absicht*, page 83–103. Klett-Cotta, Stuttgart. (Festschrift für Dieter Henrich).
- Perry, J. (1998b). Rip Van Winkle and Other Characters. *The European Review of Analytical Philosophy, Volume 2: Cognitive Dynamics*, page 13–39.
- Perry, J. (2001). *Reference and Reflexivity*. CSLI Publications, Stanford.
- Plunkett, D. (2015). Which concepts should we use?: Metalinguistic negotiations and the methodology of philosophy. *Inquiry*, 58(7-8):828–874.
- Plunkett, D. and Sundell, T. (2013). Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint*, 13(23):1–37.

- Plunkett, D. and Sundell, T. (2019). Metalinguistic negotiation and speaker error. *Inquiry*, page 1–27.
- Prior, A. (1967). *Past, Present, and Future*. Oxford University Press, Oxford.
- Prior, A. N. (1957). *Time and Modality*. Oxford UP, Oxford.
- Prior, A. N. (2003). Changes in events and changes in things. In Hasle, P., Øhrstrøm, P., Braüner, T., and Copeland, J., editors, *Papers on Time and Tense*, chapter 1. Oxford UP, Oxford.
- Putnam, H. (1975). The Meaning of ‘Meaning’. In Putnam, H., editor, *Mind, Language and Reality*, number 2 in Philosophical Papers, page 215–271. Cambridge University Press, Cambridge, UK.
- Quine, W. V. O. (1951). Two dogmas of empiricism. *The Philosophical Review*, 60(1):20–43.
- Quine, W. V. O. (1960). *Word and Object*. John Wiley and Sons, New York.
- Rast, E. (2017a). Metalinguistic value disagreement. *Studia Semiotyczne*, XXXI(2):139–159.
- Rast, E. (2017b). Value disagreement and two aspects of meaning. *Croatian Journal of Philosophy*, 17(51(3)):399–430.
- Rast, E. (2020). The theory theory of metalinguistic disputes. *Mind and Language*, forthcoming (publ. online first):1–19.
- Rast, E. H. (2010). Plausibility revision in higher-order logic with an application in two-dimensional semantics. In Arrazola, X. and Ponte, M., editors, *Proceedings of the logKCA-10 – Proceedings of the Second ILCLI International Workshop on Logic and Philosophy of Knowledge, Communication and Action*, Donostia/San Sebastian. ILCLI, University of the Basque Country Press.
- Rast, E. H. (2014). Context as assumptions. In Lihoreau, F. and Rebuschi, M., editors, *Knowledge, Context, Formalism*, page 9–39. Springer.
- Recanati, F. (2004a). Deixis and Anaphora. In Szabó, Z. G., editor, *Semantics vs. Pragmatics*. Oxford University Press, Oxford.
- Recanati, F. (2004b). *Literal Meaning*. Cambridge University Press, Cambridge, UK.
- Recanati, F. (2006). Crazy Minimalism. *Mind and Language*, 21(1):21–30.
- Reichenbach, H. (1947). *Elements of Symbolic Logic*. Macmillan, New York.
- Schlenker, P. (2000). *Propositional Attitudes and Indexicality: A Cross-Categorical Approach*. PhD thesis, MIT, Massachusetts. Dissertation 1999, MIT; minor changes in 2000.
- Schlenker, P. (2003). A Plea for Monsters. *Linguistics and Philosophy*, (26):29–120.

- Searle, J. (1995). *The Construction of Social Reality*. Penguin, London.
- Searle, J. (2005). What is an institution? *Journal of Institutional Economics*, 1(1):1–22.
- Sperber, D. and Wilson, D. (1986). *Relevance: Communication and Cognition*. Blackwell, Oxford.
- Sperber, D. and Wilson, D. (2004). Relevance theory. In *The Handbook of Pragmatics*, page 607–632. Blackwell, Oxford: Blackwell.
- Stalnaker, R. (1978). Assertion. In Cole, P., editor, *Pragmatics*, page 315–332. Academic Press, New York. (= Syntax and Semantics Vol. 9).
- Stanley, J. (2004). On the linguistic basis for contextualism. *Philosophical Studies*, 119(1-2):119–146.
- Stanley, J. (2005). Semantics and context. In Preyer, G. and Peter, G., editors, *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, page 221–253. Oxford UP, Oxford.
- Stokhof, M. and Groenendijk, J. (1991). Dynamic Predicate Logic. *Linguistics and Philosophy*, 14(1):39–100.
- von Stechow, A. (1984). Structured Propositions and Essential Indexicals. In Landman, F. and Veltman, F., editors, *Varieties of Formal Semantics*, page 385–403. Kluwer, Dordrecht.