

Relevance, Resonance, and Meaning

How do participants in ordinary conversation reliably converge on their interlocutor's intended meaning? A commonsense answer is that hearers possess a set of rules which they use to associate each of a speaker's expressions with some conventional meaning. But this answer fails to account for the fact that people often use expressions to mean something which differs from their conventional meaning; for instance, "John is a calculator" can be used to mean *John is good at math*, and a sarcastic utterance of "Economics is so much fun" can mean *Economics is boring*. Thus, as the above examples suggest, hearers must determine what is said (i.e. the 'literal' content of an utterance) before they can infer what a speaker intends to communicate, which goes beyond the conventional meaning of their utterance. But there is still one problem; like what is communicated, what is said is not fully determined by conventional meanings and varies in relation to the context of utterance; for instance, a 'red pen' might have red ink, but be black on the outside, whereas a 'red apple' is red on the outside, but not the inside. Given that what is said is not fully determined by conventional meanings, how do hearers determine what is said?

Though the determination of what is said has traditionally been accounted for by a process of semantic composition that is insulated from pragmatic influence (see Grice 1975, described below), recent work in the philosophy of language suggests that pragmatic factors play a crucial role in semantic composition (François Recanati, 2010). In light of this conclusion, the line between semantics and pragmatics must be redrawn, and theorists must provide a systematic account of the pragmatic functions which allow context to influence the semantic composition process. Dan Sperber and Diederik Wilson (1995) provide a compelling solution to the latter problem by accounting for the pragmatic processes involved in semantic composition with a cognitive principle of relevance. However, though this account is grounded in a useful theoretical framework which is backed by empirical research in cognitive science and psychology, it faces critical problems of circularity and underdetermination. Alternatively, David Beaver and Jason Stanley (2023) lay the foundation for a new theoretical framework which defines meanings not as psychological states, but objective statistical correlations. This theory offers a unique way of accounting for the semantic composition process which avoids the problems of circularity and underdetermination. Furthermore, it offers an account of how large-scale changes in human behavior are influenced by individual speech acts. Nevertheless, it is not without serious problems of its own.

In chapter one, I argue that we should define 'what is said' according to a principle of accessibility; what is said is given by our intuitions, not some minimal semantic content. Furthermore, drawing from Recanati (2010), I argue that pragmatic factors have a top-down influence on the semantic composition of what is said. In other words, pragmatic factors affect

semantic composition in ways that are not mandated by linguistic material. This view amounts to a moderate contextualism which preserves compositionality.

Next, in chapter two, I demonstrate how Sperber and Wilson's relevance-theoretic framework can be applied as a top-down pragmatic process that enters the determination of what is said. Ultimately, this amounts to a cognitive principle of relevance which balances contextual effects and processing cost to arrive at an interpretation of what is said which is consistent with the principle. Again, though this theory is grounded in a well-established theoretical framework, it faces critical problems of circularity and underdetermination.

Finally, in chapter three, I adapt Beaver and Stanley's (2023) paradigm-shifting theoretical framework to offer an alternative account of the top-down pragmatic processes entering the determination of what is said. This amounts to defining meanings in terms of statistical correlations, or resonances. To determine what is said, hearers must be 'attuned' to the resonances of a speaker's utterances. Again, this theory avoids problems of circularity and underdetermination, whilst offering a compelling explanation of how individual utterances contribute to large-scale changes in human behavior. However, by focusing on behavior and not psychology, it neglects to provide an adequate explanation of how hearers interpret novel speech acts.

Looking forward, insights from both Sperber and Wilson's relevance theory and Beaver and Stanley's resonance-based theory should constrain and guide future research. By focusing solely on the psychological aspects of speech comprehension, Sperber and Wilson lose sight of the rich connection between speech practices and human behavior. On the other hand, by focusing on behavior, whilst neglecting to provide an explanation of the psychological processes underlying speech practices, Beaver and Stanley leave many aspects of the comprehension process unexplained. Hence, future research must bridge the gap between the psychological and behavioral. This paper is an attempt to synthesize the two separate research paradigms, taking insights from resonance and relevance theory while attempting to avoid the problems of both.

Chapter 1: What is 'What is Said'?

1.1 Introducing the problem

Paul Grice (1975) is responsible for popularizing the technical distinction between 'what is said' and 'what is communicated,' writing "In the sense in which I am using the term *say*, I intend what someone has said to be closely related to the conventional meaning of the words (the sentence) he has uttered" (pg. 44). Setting aside the ambiguity of the phrase 'closely related,' the point is intuitively clear. Suppose A utters (1) and B answers with (2).

- (1) Can I borrow your car?
- (2) The keys are on the kitchen table.

With (2), B answers A's question. A has B's permission to borrow their car. Nevertheless, nothing about what is said in (2) explicitly addresses A's question. What B said differs from what they communicated. This aligns with how Grice defines the distinction; clearly, there is no convention which connects the sentence expressed in (2) to the meaning which B intended to communicate with (2). Rather, the hearer must interpret the sentence expressed, in conjunction with extralinguistic aspects of the context, to infer what B was trying to communicate.

Put more succinctly, on the traditional Gricean picture what is said is some proposition P which can be combined with a set of propositions constituting the context to yield implicated propositions Q_1 through Q_n . P is taken to be conventionally related to the sentence uttered, whereas Q_1 through Q_n are not. This view does a great job of accounting for cases where the distinction between what is said and what is communicated is intuitively clear. However, when the view is extended to account for cases where the distinction is blurred, it leaves many questions unanswered. Suppose A utters (3).

(3) Pass me three screws.

What is said in (3)? In other words, what is the P from which some set of implicated propositions Q_1 through Q_n are inferred? Intuitively, (3) expresses a request for *exactly* three screws. However, B could respond to (3) by passing A eight screws. If A then responds by complaining that she asked for *exactly* three screws, B could argue that A only *said* she wanted a minimum of three screws, not that she wanted *exactly* three screws. There are numerous ways to account for this case, each of which can be plotted on the continuum between two diametrically opposing views in the philosophy of language—semantic minimalism, and contextualism.

As characterized by Claudia Bianchi (2010), semantic minimalism is that view that the meaning of an utterance is,

[...] closely related to the conventional meaning of the linguistic expressions employed in it and departs from that meaning only in cases of ellipsis, ambiguity and indexicality strictly understood (i.e. concerning only a small number of expressions such as true indexicals and demonstratives) (Pg. 2-3).

The view I want to criticize is not semantic minimalism, *per se*, but rather an extension of semantic minimalism which holds that the latter account of utterance meaning is an account of what is said. I call this account *literalism*. Assuming that the strictly linguistic meaning of (3) is something akin to *pass me at least three screws*, the literalist interprets the hearer who hands the speaker eight screws to be acting consistently with what the speaker said. If the speaker does manage to communicate that they want *exactly* three screws, then it is not because they *said* they want *exactly* three screws, but because they merely implicated it.

On the other end of the spectrum, Bianchi defines radical contextualism as the view that a sentence expresses a complete truth-conditional proposition only when it is used in a context where it is freely enriched by pragmatic constituents. He writes,

A sentence expresses a proposition only in the context of a speech act, when completed and enriched with pragmatic constituents that do not correspond to any syntactic element of the sentence (neither an explicit constituent, as in cases of syntactic ellipses, nor a hidden indexical present at logical form) and yet are part of its semantic interpretation (Pg. 6).

The radical contextualist argues that the literalist is wrong, the speaker of (3) literally said they want *exactly* three screws. What is said goes beyond what is encoded in the linguistic content of the sentence and is always partially determined by pragmatic factors. In other words, the sentence itself fails to fully determine what is said, and factors which are not mandated by linguistic material nevertheless affect semantic composition.

1.2 Support for Moderate Contextualism

The view I endorse here is akin to the moderate contextualism put forth in Recanati's (2010) *Truth-Conditional Pragmatics*. As opposed to radical contextualism, moderate contextualism is the view that though some sentences may express a complete proposition independently of any context, those sentences are the exception rather than the norm. Typically, what is said is partially determined by top-down pragmatic processes (i.e. pragmatic processes which are not mandated by linguistic material). Here, I discuss three strong reasons to adopt a moderate contextualist view. Firstly, rather than defining what is said as some inaccessible semantic content which plays a purely theoretical role, it takes our intuitive notion of what is said to be the explanandum of semantic theory. Secondly, it fits well with a plethora of cases in ways that opposing theories do not. And finally, it is easily defended against a common charge of non-compositionality.

1.2.1 Support One: The role of semantic theory

One way of arguing for the literalist position on what is said is to hold that the concept plays a purely technical or theoretical role, and (hence) is not meant to capture the intuitive distinction between what is said and what is communicated. By positing the existence of an inaccessible layer of meaning at a subpersonal level and simply calling it 'what is said', perhaps we can better account for the interpretation process that allows us to reliably converge on what a speaker intends to communicate. This view draws the line between what is said and what is communicated at the line between what is given by linguistic material and what must be inferred with extralinguistic information.

Indeed, there is good reason to posit the existence of some minimal linguistic content at a sub-personal level which plays a crucial role in the interpretation process. Denying this would arguably leave meaning completely unconstrained. If there is not at least some convention limiting the set of possible meanings of some expression, then it is unclear how the expression could be used to pick out a definite meaning. However, it is misleading to call this minimal linguistic content 'what is said', because it does not fit with what people intuitively define as

‘what is said’. When a speaker says, “Pass me three screws,” we do not consciously engage in an inference process from the minimal content, *pass me at least three screws*, to the enriched content, *pass me exactly three screws*. Hence, we should not say that what is said is the minimal content. Implicit here is the assumption that we should rely on something akin to Recanati’s availability principle (AP).

(AP) In deciding whether a pragmatically determined aspect of utterance meaning is part of what is said, that is, in making a decision concerning what is said, we should always try to preserve our pre-theoretic intuitions on the matter (Recanati, 1989).

AP forces us to start with the evidence and constrain our possible set of theories, rather than to interpret the evidence in accordance with our preexisting doctrines. Given AP, it is undeniable that what is said is deeply determined by contextual information. The intuitive meaning of what is said often goes far beyond what is given by linguistic material as evidenced by the cases in 1.2.2. In other words, it is typical for the P which is combined with a context C to yield implications Q_1 through Q_n to be enriched by contextual information that is not contained in the conventional meanings of the expressions used.

Thus far, I have defended moderate contextualism against a literalist position about what is said. Again, it is clear that what is said is more than what is minimally given by linguistic material. However, even given this assumption, there are multiple positions to occupy. In the rest of chapter two, I defend moderate contextualism against indexicalism, the view that context only enters into what is said when it is mandated by the linguistic material. This view argues that the type of rule-governed context sensitivity observed in explicit indexicals like ‘it,’ ‘I,’ or ‘that’ extends to seemingly nonindexical words such as adjectives, verbs, and nouns; top-down pragmatic processes—that is, pragmatic processes not mandated by the linguistic material—never enter into what is said. Contrastingly, moderate contextualism claims that top-down pragmatic processes often enter into the determination of what is said.

1.2.2 Support Two: Plethora of cases

In each of the following cases, taken from Bianchi (2010), the information represented in brackets is not part of the sentence uttered, and does not seem to be conventionally associated with the linguistic material.

- (4) I have nothing [appropriate to the occasion] to wear tonight.
- (5) Some [not all] children got stomach ‘flu.
- (6) Jill got married and [then] became pregnant.
- (7) Bob has [exactly] three cars.
- (8) Jack and Jill are engaged [to each other].
- (9) Tom hasn’t had breakfast [today].

These few cases are not cherry picked to prove a point; rather, they represent a widespread and pervasive phenomenon of semantic underdetermination. Furthermore, they provide compelling

evidence for moderate contextualism. The literalist's explanation would conflate what is said with the very minimal content which is inaccessible at the level of conscious awareness, and the indexicalist's explanation would posit the existence of covert indexical variables that mandate how contextual information should enter into interpretation. However, with many of the above examples, it is possible to imagine a context where the meaning of the utterance is not enriched with contextual information in the suspect locations. In other words, the contextual enrichment represented in brackets in the above cases seems optional—its possible to imagine cases where the extra information is not there. Take, for instance, (7). Suppose you and a friend are writing invitations to a car convention, but you do not want to invite anyone who has less than three cars. In such a case, (7) could be used in a way that is ambiguous between *exactly three* and *at least three*. So, the enrichment is optional (top-down) and not mandated by the linguistic material (bottom-up). I provide a more detailed account of this type of argument in what follows.

1.2.3 Support Three: Preserving Compositionality

A common charge against moderate contextualism is that by claiming that meaning is deeply context sensitive and influenced by top-down processes, the position fails to account for compositionality—the idea that the meaning of a complex expression is a function of the meaning of its component expressions. On the other hand, indexicalism fits intuitively with compositionality. Bottom-up processes are insulated and rule-governed; they fix the meanings of the terms which enter into composition, but top-down processes can be influenced by anything in the context and are not rule-governed in the sense that they are not mandated. Furthermore, top-down processes can influence how composition happens, because the context of prior words in the sentence can prompt pragmatic processes to affect later words. Nevertheless, I argue that despite their unpredictable nature and their resistance to formalization, top-down pragmatic processes pose no serious threat to compositionality. Hence, moderate contextualism is—contrary to popular belief—compatible with compositionality.

Formally speaking, compositionality is the idea that the interpretation of what is said involves two steps, the first being the application of a lexical rule (LR) which assigns a certain interpretation (m) to each simple expression (α), and the second being the application of a compositional rule (CR) which states that the interpretation of any complex expression ($\alpha*\beta$)—where ‘*’ stands for an arbitrary mode of combination—is a function (f) of the interpretation of its parts (Recanati 2010, Pg. 27-28).

$$(LR) I(\alpha) = m$$

$$(CR) I(\alpha*\beta) = f(I(\alpha), I(\beta))$$

The basic idea is that the meaning of any complex expression is fully determined by the meaning of its component parts and their mode of combination. Hence, to use Recanati's terminology, compositionality seems to imply that ‘top-down’ processes (i.e. optional pragmatic processes which are not mandated by linguistic material but driven by contextual factors) cannot enter into the interpretation of what is said. Rather, as literalism advocates, the only pragmatic processes

that enter into the interpretation of what is said are ‘bottom-up’ processes (i.e. pragmatic processes that are mandated by linguistic material) like those involved in the interpretation of indexicals (Pg. 30).

If compositionality implies that top-down pragmatic processes cannot enter into the interpretation of what is said, then there is reason to think that compositionality is false, since top-down pragmatic processes are often involved in the interpretation of what is said, as illustrated by the use of ‘danced’ in (8) and (9) (examples adapted from Recanati, Pg. 84-85).

(10) “John danced all night.”

(11) A: “Was John at the club last night?”

B: “Yes, John danced all night.”

In (10), the verb ‘John danced’ is location-indefinite (i.e. the location where John danced is not given), whereas in (11), ‘John danced’ means *John danced at the club*. To account for the semantic flexibility of verbs like ‘danced’ in (10) and (11) whilst preserving compositionality by relying exclusively on bottom-up processing, some theorists have argued for the existence of covert location variables which are encoded into the linguistic material. Recall, however, that bottom-up processes are mandated by the linguistic material, while top-down processes are optional. Thus, the existence of location-indefinite cases like (10) undermines the proposal that a bottom-up process is involved in the interpretation of ‘danced’, since if a bottom-up process was involved, the specification of where John danced would be necessitated by the linguistic material. So, either compositionality does not imply that top-down pragmatic processes cannot enter into the determination of what is said, or compositionality is false. Again, following Recanati, I argue for the former position.

Compositionality is compatible with semantic flexibility, and top-down pragmatic processes do enter into the determination of what is said. In other words, the fact that the meaning of any complex expression is fully determined by the meaning of its component parts and their mode of combination is compatible with the existence of top-down pragmatic processes. Semantic flexibility is accounted for by allowing the meaning of the component parts of an expression to be ‘modulated’ (i.e. determined by top-down pragmatic processes) with respect to the linguistic and extra-linguistic context. The modulated meanings of an utterance’s component expressions are what ultimately enters into the compositional process. Thus, according to Recanati, the compositional rule (CR) can be adapted to yield (CR’), which accounts for the influence of modulation on the compositional process.

$$(CR') I(\alpha*\beta) = f(g_1(I(\alpha)), g_2(I(\beta)))$$

In CR’, ($\alpha*\beta$) stands for a complex expression constituted by expressions ‘ α ’ and ‘ β ’, and the ‘g’s are “free higher-order variables ranging over available pragmatic functions (including identity, which gives us the literal case)” (Pg. 128). Essentially, the meaning of a complex phrase is a function of the meaning of its parts, but those parts can be modulated by top-down pragmatics

processes before they are combined to produce the larger expression of which they are a part. These pragmatics functions are represented by the quasi-variables g_1 and g_2 in CR'. In a given context, some specific pragmatic function modulates the meaning of (I(a)) and (I(β)) before—and potentially after—they enter into the compositional process.

Without getting too bogged down by details, (CR') should suffice to show that it is possible to provide an account of semantic composition which is compatible with top-down pragmatic processes. Metaphorically speaking, this view claims that the composition of a sentence is

[...] like the construction of a wall from sand-bags of different kinds. Though the size, structure, texture and contents of a sand-bag restrict the range of shapes it can take on, the actual shape it adopts in a particular situation depends to a greater or lesser extent on the shapes adopted by other sand-bags in the wall.

This is in stark contrast with the traditional view, which claims that semantic composition is like the construction of a wall from bricks of different shapes, where

[...] the result depends on the properties of the parts and the pattern of their combination. But just as each brick as each brick has exactly the same shape in every wall or part of a wall to which it is moved, so too each standard sense of a word or phrase is exactly the same in every sentence or part of a sentence in which it occurs (Cohen 1986 as quoted in Recanati 2010, pg. 31).

Setting aside the technicalities, these metaphors get at the heart of the idea behind the moderate contextualism I advocate for. The meaning of a word is not an insulated content that can be studied on a microscope. Rather, the meaning of a word is dependent on the words around it and the context in which it is uttered.

For present purposes, CR' provides an independently motivated schema or formalization of the way that context enters the semantic composition process. In the following two chapters, I examine two opposing theories of meaning—relevance theory and resonance theory— both of which fit naturally within the moderate contextualist framework advocated by Recanati. Using CR' allows me to show how both relevance theory and resonance theory can be applied as formal accounts of the interpretation process which allows people to determine what is said.

Chapter 2: Relevance as an Account of Semantic Composition

In this chapter, I introduce Sperber and Wilson's principle of 'relevance', a cognitive principle which constrains how human minds process information. This serves as a potential way of accounting for the top-down pragmatic processes involved in determining what is said. The theory builds from empirically supported theory of mind to give a plausible account of how the mind processes information. The principle of relevance itself, however, faces several serious problems.

2.1 Introducing Relevance

The book entitled *Relevance: Communication and Cognition* by Sperber and Wilson (1995) can be understood as an attempt to provide an explanatorily powerful account of the nebulous concept of relevance that kept Paul Grice awake at night in the mid-seventies. In the middle of his seminal paper, *Logic and Conversation*, Grice confesses,

Under the category of RELATION I place a single maxim, namely, 'Be relevant.' Though the maxim itself is terse, its formulation conceals a number of problems that exercise me a good deal: questions about what different kinds and focuses of relevance there may be, how these shift in the changed, and so on (Grice, 1975, p. 46).

In the classic paper, the notion of relevance is supposed to play an explanatory role; by assuming that speakers are adhering to a maxim of relevance, hearers thereby bridge the gap between what is said and what is communicated. The idea is that maxims like relevance act as premises in inferences which allow the speaker to infer Q_1 through Q_n from P in a context C . But, as Grice points out, it is far from clear what it means to assume that a speaker is adhering to a principle of 'relevance.'

To address Grice's concerns, Sperber and Wilson focus on two fundamental questions—(I) and (II).

(I) How do human minds determine which information is relevant? In other words, how do we determine which information is worth processing and which is not?

(II) How do hearers arrive at what a speaker intends to communicate with their speech acts?

The strategy is this: First, answer question (I). That is, provide a theory which is so powerful and broad that it predicts which of many stimuli in any given context will be most 'relevant' to some individual, where 'relevance' is the mind's way of ranking which information should be processed first. Next, demonstrate that the answer to question (I) is also an answer to question (II) by arguing that hearers lock on to the speaker's intended meaning by processing the interpretation of their utterance which has the highest degree of relevance. Let's examine the argument step by step.

To answer question (I), Sperber and Wilson start by assuming a theory of mind akin to Jerry Fodor (1975). This Fodorean framework (otherwise known as the language of thought hypothesis or LOTH) provides the perfect grounds for building up a theory of comprehension and is well supported by modern empirical research in psychology (see Quilty-Dunn, et. al., 2023). LOTH claims that the mind possesses a private language of discreet mental representations which are manipulated by the application of logical rules. Simple mental representations can be combined to produce complex mental representations, which, in turn, can

be decomposed into simple representations. Adopting this theory of mind, Sperber and Wilson write,

The fact that many central processes are inferential puts an important constraint on the conceptual representation system. Conceptual representations must have logical properties: they must be capable of implying or contradicting one another, and of undergoing deductive rules (Sperber and Wilson, 1995, p. 46).

In other words, assuming that central processes are inferential, it follows that our mental representations have certain logical properties (namely, truth-conditions) that allow them to enter deductive processes. Since there is good reason to believe that our thought processes *are* inferential—we constantly perform inferences in natural language which align with valid applications of logical rules such as conjunction elimination ($\Phi \ \& \ \Psi \vdash \Phi$) or disjunctive syllogism ($\Phi \ \text{OR} \ \Psi, \sim \Psi \vdash \Phi$)—there is also good reason to think that mental representations have logical properties.

So, as advocated by LOTH, the mind can be understood as a sort of information processing device which is constantly performing logical operations on the formal structure of mental representations. Furthermore, Sperber and Wilson argue, the mind is an *efficient* information processing device, aimed at improving its representation of the world in the most cost-effective way possible, presumably to increase evolutionary fitness. To do this, the mind combines new information with old information to yield improved representations. In sum,

[when...] interconnected new and old items of information are used together as premises in an inference process, further information can be derived: information which could not have been inferred without this combination of old and new premises. When the processing of new information gives rise to such a multiplication effect, we call it relevant (pg. 48).

Hence, a piece of information is relevant if it can be combined with stored (old) information to produce new information via inferential processes. Furthermore, Sperber and Wilson add, relevance is scalar—some information is more relevant than other information.

Now, to be more specific, relevance (R) is measure of the balance between the contextual effects of processing an assumption (C), and the cost of processing that assumption (P).

$$R = C \div P$$

The idea is that relevance can be understood as a sort of cost-benefit calculation—to determine whether some information is worth processing, first determine how much it will improve your representation of the world and then weigh that with the cost of processing it. If the cost far outweighs the benefit, the information is irrelevant. On the other hand, if the improved representation heavily outweighs the cost, then the information is relevant and worth processing. But what does it mean to say that information improves your representation of the world, and

how do you measure the extent to which it does so? Similarly, how do you measure the cost of processing a new piece of information?

Sperber and Wilson use the notion of *Contextual Effects* (C) to measure the extent to which a given piece of information improves the mind's representation of the world. Since our representation of the world plausibly consists of a set of assumptions—i.e. representations taken, by default, to accurately describe the actual world—the notion of contextual effects is defined in terms of assumptions. Ultimately, the number of contextual effects of any given assumption can be calculated by summing the number of implications that the assumption has in the context of processing it, where the context is defined as the subset of assumptions which are brought to bear on the processing of new assumptions¹ (pg. 15). Implications can be contradictions, where the new assumption contradicts with assumptions in the context, strengthenings, where the new assumption increases confidence in a previously held assumption, or weakenings, where the new assumption weakens the confidence in previously held assumptions. Assume a context of (14) through (17), where the brackets indicate the 'strength' or 'confidence'² of each assumption, and consider the assumptions represented in (12) and (13).

(12) Tom hasn't had breakfast today. (NEW ASSUMPTION)

(13) Tom is human. (NEW ASSUMPTION)

(14) If Tom hasn't had breakfast, Tom is hungry. [STRONG]

(15) Tom is hungry. [MEDIUM]

(16) If Tom hasn't had breakfast, Tom is sad. [STRONG]

(17) Tom is happy. [STRONG]

Here, assumption (12) has a large number of contextual effects; (15) is implied (and hence, strengthened from medium to strong) by (12), (14), and modus ponens; and the negation of (17) is implied (hence, weakening (17) from strong to medium) by (12), (16), and modus ponens. So, if the processing cost of (12) low, then it is highly relevant in the context of (14) through (17), because it has several contextual effects. On the other hand, (13) is irrelevant because it has no contextual effects.

Now, how is processing cost calculated? Admittedly, processing cost is difficult to quantify, but there are at least three key factors. Firstly, how logically complex is the

¹ Here, what counts as an implication is limited; introduction rules such as addition ($\Phi \vdash \Phi \text{ OR } \Psi$) or conditional introduction ($\Psi \vdash \text{IF } \Phi \text{ THEN } \Psi$) are strictly prohibited, because including them as contextual effects would imply that all assumptions have an infinite number of effects on the context. Addressing this problem, Sperber and Wilson conclude that "[...] introduction rules play no part in the spontaneous deductive processing information, the processing which our deductive device is designed to describe." Nevertheless, they also admit that "This is a substantial claim, and one that most people working on the psychology of deduction have been reluctant to make" (pg. 96). See Seuren (1987) for a more substantial commentary.

² 'Strength' or 'Confidence' is a non-representational measure of the credence of assumptions, comparable to the notion of 'accessibility' a term which Sperber and Wilson gesture at by writing, "[...] as a result of some kind of habituation, the more a representation is processed, the more accessible it becomes. Hence, the greater the amount of processing involved in the formation of an assumption, and the more often it is accessed thereafter, the greater its accessibility (Pg. 77).

assumption? Clearly, “Tom hasn’t had breakfast” is less complex than the longer assumption, “Tom hasn’t had breakfast, and he doesn’t like cheerios.” Hence, there is reason to think the processing cost of the first assumption is less than that of the second. Secondly, is the assumption contained in a larger chunk of memory? Sperber and Wilson speculate that, “[...] the smallest units which can be transferred from encyclopaedic memory³ to the memory of the deductive device⁴ are chunks rather than individual assumptions” (pg. 138). In other words, the mind stores memories in the form of interconnected chunks or bundles of information which consist of sets of assumptions whose activation consistently co-occurs. If an assumption is contained in a larger chunk, it requires less processing cost to access it. Finally, the third key factor involved in calculating processing cost is accessibility. How accessible is the context brought to bear on the processing of the assumption? Recall that the context is the set of assumptions brought to bear on the processing of a new assumption. For each new assumption that is processed, a context must be provided in which that assumption has some contextual effects. If the necessary context has never been processed or is not stored in a chunk of other assumptions, the context will be relatively inaccessible, and processing the new assumption will be very costly.

Now, consider (12) and (18), both of which are potential interpretations of (9), i.e. “Tom hasn’t had breakfast.”

(12) Tom hasn’t had breakfast today (NEW ASSUMPTION).

(18) Tom has never had breakfast in his entire life (NEW ASSUMPTION).

Assume that ‘breakfast’ means something akin to *food in the morning*. In normal contexts, I argue, (12) has more contextual effects than (18). Only if an abnormal context—for instance, a context where Tom was recently conceived, or a context in which he is not a human, but an android—is brought to bear on interpretation, will (18) have more contextual effects than (12). However, in normal circumstances, conjuring such a context would require a lot of processing power, making (12) more relevant to typical interlocutors.

2.2 Relevance as a Top-Down pragmatic Process

Above, I have shown how the principle of relevance determines the relevance of any given assumption in any given context. This answers question (I) by giving an account of how the mind determines which information is worth processing. But the problem addressed in this paper is not how the mind determines the relevance of a given assumption, but rather how the mind determines which of many potential assumptions a speaker is intending to communicate with their speech act in a particular context. Recall that Sperber and Wilson answer the latter question—that is, question (II)—by arguing that the answer to question (I) *is* an answer to

³ Encyclopedic memory is a store of factual information which constitutes the mind’s concepts. For instance, the encyclopedic entry of the mind’s *cat* concept contains information such as *animal that has claws* and *domesticated pet*.

⁴ The ‘deductive device’ is the central processing unit of the mind. It creates thoughts by performing formal operations on the mind’s mental representations.

question (II). The principle of relevance not only determines which assumptions are relevant, but it also determines which assumptions are intended by the speaker. Here, I elaborate on this argument before generalizing it to a formal account of semantic composition.

In chapter one, I briefly mentioned that relevance theory fits naturally with moderate contextualism. This is because relevance theory allows for the existence of incomplete assumptions. Rarely, Sperber and Wilson admit, do utterances express complete truth-conditional assumptions. Rather, utterances typically express ‘assumption schemas’—i.e. incomplete logical forms or expressions containing unbound variables (pg. 73)—which must be enriched before they express complete assumptions. This view might sound similar to indexicalism, but there is one key difference that makes it much closer to moderate contextualism; the linguistic material does not fully determine how the contextual variables will be filled with contextual information, and since there are often numerous ways that assumption-schemas can be developed into complete assumptions, top-down pragmatic processes must enter into semantic composition. Sperber and Wilson’s answer to question (II) helps clarify this view.

Sperber and Wilson’s answer to the question (II) is that “the right propositional form [i.e. the interpretation of what is said which aligns with the speaker’s intentions] is the one that leads to an overall interpretation which is consistent with the principle of relevance” (pg. 184). This answer is motivated by (19), where (at the purely linguistic level) there is no reason to prefer the *drinking-tube* interpretation of ‘straw’ to the *cereal-stalk* interpretation, and yet most hearers default to the former (Pg. 186).

(19) The child left the straw in the glass.

The reason that hearers default to the *drinking-tube* interpretation, Sperber and Wilson argue, is that “a child drinking from a glass is a stereotypical event which [...] is recorded in the form of a single chunk, stored at a single location in memory and accessed as a single unit” which constitutes a highly accessible context in which “the *drinking-tube* interpretation of [straw] can be processed at a minimal cost” (Pg. 186). Thus, the *drinking-tube* interpretation is more relevant—and hence, defaulted to by most hearers—because it can be accessed at a lower processing cost than the *cereal stalk* interpretation; in other words, the drinking-tube interpretation is more manifest than the cereal stalk interpretation. Presumably, in most contexts (C) the value of (P) for the *cereal stalk* interpretation (*c*) is much higher than the value of (P) for the *drinking-tube* interpretation (*d*), and since neither interpretation clearly yields a greater value for (E), the relevance of (*d*) is greater than the relevance of (*c*). However, this is not *necessarily* the case; it is possible to imagine contexts in which (*c*) would be more relevant than (*d*). For instance, in a strange context where a group of people are removing cereal stalks from glasses, (*c*) might be more relevant than (*d*) because the value of (E) is much higher.

Now, let’s examine how this kind of explanation applies to our chapter one example.

(3) Pass me three screws.

Recall two potential interpretations of (4); the intuitive interpretation, *pass me exactly three screws*, and the minimal interpretation, *pass me at least three screws*. In a construction or engineering context—i.e. in contexts stereotypically associated with utterances like (4)—the intuitive interpretation has many contextual effects, because it increases the credence of many plausible assumptions; that holding more than three screws at a time is difficult, or that the speaker is trying to complete a task which obviously requires exactly three screws to complete. Presumably, it would take more processing cost to conjure an unusual context in which the speaker wants a minimum of three screws than it would to conjure a context where the speaker wants exactly three screws. Hence, the principle of relevance acts as a top-down pragmatic process which is not mandated by linguistic material, and which prevents the processing of the unusual interpretation.

Having examined how the principle of relevance applies to semantic composition in particular cases, let's now see how this generalizes to a formal theory of composition. A first attempt at implementing relevance theory into a systematic account of the semantic composition process involves simply substituting the 'g's in (CR') with a function (R) which takes as argument a context (C) and some expression (α), yielding as output the most relevant interpretation of (α). Call this the relevance theoretic compositional rule (or RTCR1).

$$\begin{aligned}(\text{CR}') \quad I(\alpha * \beta) &= f(g_1(I(\alpha)), g_2(I(\beta))) \\ (\text{RTCR1}) \quad I(\alpha * \beta) &= f(R(I(\alpha)), R(I(\beta)))\end{aligned}$$

According to one interpretation of RTCR1, semantic composition starts when some expression (α) has been identified by a specialized module which automatically turns speech sounds into phonemic representations. Next, the default semantic interpretation of (α) is given by ($I(\alpha)$) and fed as input into (R), which yields the most relevant interpretation of ($I(\alpha)$) in (C). Finally, the same process is applied to (β) to yield ($R(I(\beta))$), which is combined with ($R(I(\alpha))$) according to some mode of combination '*' to yield what is said.

Notice that the processing involved in (RTCR1) presumably proceeds in a stepwise fashion; ($R(I(\alpha))$) is processed before ($R(I(\beta))$). But this cannot be the case, for it would confine the principle of relevance to the interpretation of simple expressions, preventing it from being applied (as intended) to complex expressions. For instance, if (RTCR1) was true, then it would be impossible to compare the relevance of (20) and (21).

(20) John ate rocky road for breakfast.

(21) John ate rocky road for breakfast, and Thelonious Monk is an exquisite pianist.

But this is absurd, for complex expressions do have varying degrees of relevance. Thus, by reductio, (RTCR1) is false.

To avoid the above problem and implement relevance theory into a systematic account of the semantic composition process, the compositional rule would have to be something more like (RTCR2).

$$(RTCR2) I(\alpha*\beta) = R(f(I(\alpha), I(\beta)))$$

According to (RTCR2), the meaning of a complex expression is a function of the context, and of the semantic meaning of the expression's component parts. The semantic meaning of the component parts is given by $f(I(\alpha), I(\beta))$, and is—according the availability principle (AP)—inaccessible to introspection. This semantic content typically takes the form of an assumption schema, rather than a complete assumption, for the meaning of most utterances is deeply underdetermined by linguistic content. Thus, the principle of relevance is applied to the assumption schema yielded by $f(I(\alpha), I(\beta))$ to produce a complete propositional form based on the context. This propositional form constitutes what is said.

One interesting contrast between (CR') and (RTCR1) and (RTCR2), is that the latter two formulas replace the higher order variables represented by the 'g's in (CR'), with the lower order function R, which yields the most relevant interpretation of the expressions it operates over. While (CR') posits the existence of a multitude of pragmatic functions, one of which is selected in a given context, (RTCR1) and (RTCR2) hold that there is a single universal pragmatic function which applies in all contexts—namely, R.

2.3 Circularity and Underdetermination

In this section, I present two theoretical difficulties with the principle of relevance—circularity and underdetermination. These problems do not mean that relevance theory should be rejected, but rather that the *principle of relevance* fails to achieve its intended aims. By building from the language of thought hypothesis, *Relevance theory* provides a promising and theoretically useful account of how the mind processes information, laying the groundwork for a powerful account of semantic composition. However, the principle of relevance—i.e. the proposal that the mind is always working to maximize relevance—makes no testable predictions and tells us little about how the mind determines what is said.

2.3.1 Circularity

To measure the relevance of an assumption, the mind must first find a relevant context to bring to bear on interpretation. So, the relevance of an assumption is dependent on the context it is processed in, but the relevance of said context is dependent on the assumption. Arguably, this does not create any serious problems when the assumption is fixed, that is, when there is only a single assumption whose relevance is being measured. In such a case, the relevance of the assumption is a measure of how much processing cost it takes to find a context in which it has a high number of contextual effects. However, in cases where the principle of relevance is supposed to constrain which element of a set of potential assumptions is the most relevant way of developing some assumption schema, an inevitable problem of circularity arises.

Suppose there are three assumptions under consideration (a_1, a_2, a_3) and three potential contexts (c_1, c_2, c_3). To determine which assumption is worth processing, the mind must determine which assumption leads to the greatest amount of contextual effects for the least

amount of processing cost. Hence, the principle of relevance could be understood as a two-place function which takes an assumption and a context as arguments, and outputs a relevance value. For instance, suppose $R(a_1, c_1) = 1$, $R(a_2, c_1) = 0.5$, and $R(a_3, c_1) = 0$. Given different contexts, the relevance values change; suppose, for instance, $R(a_1, c_2) = 0$, $R(a_2, c_2) = 0.5$, and $R(a_3, c_2) = 1$.

The above way of understanding relevance gives a specific relevance value for each assumption in a specific context. But we need to determine which context is the most relevant to determine which context to bring to bear on interpretation. So, let's flip the arguments in R , and reconceptualize the function as providing a relevance value for each context given a fixed assumption. So, for instance, $R(c_1, a_1) = 1$, $R(c_2, a_1) = 0.5$, and $R(c_3, a_1) = 0$. And, $R(c_1, a_2) = 0$, $R(c_2, a_2) = 0.5$, and $R(c_3, a_2) = 1$.

In sum, the principle of relevance is circular and provides no constraint on which assumption is the most relevant because the relevance of a given context depends on which assumption is chosen, and the relevance of a given assumption is dependent on which context is chosen.

2.3.2 Underdetermination

Another serious problem with the principle of relevance is that to determine how many contextual effects an assumption has in a given context—and hence, how ‘fruitful’ it will be to process it in that context—the assumption must be processed in that context. For you cannot be sure how many contextual effects an assumption will have in a context until you actually process each of those contextual effects. This position is self-defeating, because the whole point is to use the principle of relevance to avoid expending unnecessary effort by processing irrelevant assumptions.

Recognizing this worry, Sperber and Wilson suggest that the mind employs a heuristic to estimate the processing cost of new assumptions, writing that,

[...] other things being equal, stronger assumptions have greater contextual effects. Similarly, a variety of factors may make it possible to predict how the processing effort needed to perform particular pairs of tasks would compare. For instance, processing more information in the same context, or the same information in a larger context, involves a greater effort. People can take advantage of these comparative abilities in trying to maximize the relevance of the information they process (Sperber and Wilson, 1995, p. 131).

Thus, the hope is that heuristics (i.e., ‘comparative abilities’) will limit the number of assumptions whose relevance value must be determined; however, without specifying those heuristics, the notion of relevance fails to limit the assumptions which must be processed. In other words, relevance theory fails to do what it was intended to do—namely, explain how

hearers limit their interpretations of what is said to those that are intended by the speaker—shifting the explanatory burden to some vague heuristic.

So, the outcome for the principle of relevance is grim. Nevertheless, there is promise to the relevance-theoretic account of how the mind processes mental representations by performing formal operations on their logical structure.

Chapter 3: Resonance as an Account of Semantic Composition

In this chapter, I provide an alternative account of semantic composition which, like the relevance-theoretic account, also fits with the moderate contextualism put forth in chapter one. Unlike the relevance-theoretic account, however, the account put forth in this chapter defines meanings in non-psychological terms; the meaning of a term is not the psychological state of grasping that term, but the statistical correlation of the use of that term with features of the context. This account, which I call the *resonance-theoretic* account of semantic-composition, is a way of extending Beaver and Stanley's theory of meaning to work as an account of the determination of what is said. This account avoids the problems of circularity and underdetermination plaguing the relevance-theoretic account but faces its own problems of novelty and mystery. Ultimately, I hope to show that the two diametrically opposed accounts can solve each other's problems, leading to a unique theory of interpretation which bridges the gap between meaning and behavior, and between psychology and sociology.

3.1 Meanings as Resonances

Though both books address many of the same problems, Beaver and Stanley's (2023) *The Politics of Language* is motivated by drastically different concerns than is Sperber and Wilson's *Relevance: Communication and Cognition*. Since both books approach the same problems from radically different perspectives, there is much to be gained from a comparison of the two. It should also be emphasized that *The Politics of Language* is a somewhat radical attempt to create a paradigm shift in the philosophy of language. Hence, it is best understood in the context of more traditional work with which it contrasts.

The traditional way of tackling the question 'what is meaning' is to examine the psychological states and processes involved in *grasping* the meaning of a term. Consider Donald Davidson's deeply influential truth-conditional semantic theory, which states that the meaning of an expression is the conditions under which it is true (see Larson & Segal (1995) and Davidson (1967)). One virtue of truth-conditional semantic theory is that it provides a plausible and explanatorily powerful account of how language users are able to interpret novel sentences.⁵ The meaning of novel sentences is theorized to be a function of the truth-conditions of primitive

⁵ It is important to note that Davidson's truth-conditional semantic theory makes no claim about the bearers of truth-conditions. Nevertheless, the way I am describing the view may seem to imply that truth-conditional semantics defines truth-bearers as psychological entities such as mental states or propositions. This is not Davidson's view, but a natural way of building upon Davidson's theory which aligns with Fodor's language of thought hypothesis.

sentences in the user's language. For instance, the meanings of the English sentences on the left sides of the biconditionals in (22) and (23) is represented by the conditions on the right sides, expressed in some metalanguage (typically, this language is understood to be a system of mental representations).

(22) "Snow is white" iff snow is white.

(23) "Apples are red" iff apples are red.

Thus, assuming the language-user possesses (22) and (23) as theorems of their language, they can interpret any complex sentence such as (24), (25), or (26), each of which is built up from the primitive theorems.

(24) "Snow is white, and apples are red" iff [snow is white] and [apples are red].

(25) "Snow is white, or apples are red" iff [snow is white] or [apples are red].

(26) "If apples are red, then snow is white or apples are red" iff if [apples are red], then [snow is white, or apples are red].

This theory gives a plausible account of how language-users can interpret an infinite number of novel sentences using knowledge of a finite number of theorems. Perhaps in virtue of its explanatory power, truth-conditional semantic theory is typically taken for granted in most contemporary discourse in the philosophy of language, *Relevance: Communication and Cognition* being a salient example.

A contrasting way of tackling the question of 'what is meaning' is to examine the things in the world which are actually referred to by our terms. In other words, perhaps we should define the meaning of a term not as the psychological state of *grasping* that term, but rather, as the things which we attempt to *refer* to in virtue of being in the psychological state of grasping. Hilary Putnam (1973) is famous for advancing this externalist perspective on meaning, arguing that the fact that we can be wrong about what our terms refer to is evidence that "'meanings' just ain't in the head!" (pg. 704). For instance, to take Putnam's example, suppose that I incorrectly refer to an elm tree as a "birch" tree in the presence of my arborist friend. Suppose also that my friend corrects me, saying "that's an elm tree, not a birch tree." My friend would be correct, and I would be wrong. My psychological state of grasping does not define the term. Rather, the meaning of terms is their referents, and we revise our concepts in light of what we learn about the referents. Hence, Putnam concludes that all linguistic communities exemplify a *division of linguistic labor*. In other words, each linguistic community

[possesses] at least some terms whose associated 'criteria' are known only to a subset of the speakers who acquire the terms, and whose use by other speakers depends upon a structured cooperation between them and the speakers in the relevant subsets (pg. 706).

This view provides a compelling explanation of the meaning of scientific terms, whose meaning is defined by a small set of specialists. In these cases, the terms in question are defined ostensively—the meaning of the term is *that* (pointing to the object)—and the concepts

associated with the terms are subject to change based on empirical evidence. Putnam's examples tend to focus on natural kind terms, or other technical terms defined by experts. But how does it apply to ordinary terms which are not clearly defined by a small subset of our linguistic communities? It is not clear how the latter view can account for ordinary language.

Beaver and Stanley's resonance-theory (as I have chosen to call it) can be understood as a radical overhaul of Putnam's externalist theory of meaning. Like Putnam's externalism, resonance-theory rejects the view that meanings are mere psychological contents and endorses the idea that linguistic communities are involved in fixing the meaning of terms. Unlike Putnam's externalism, however, resonance-theory is intended to apply to all terms, not just the kinds described by a subset of experts. Furthermore, whereas Putnam—like most contemporary philosophers of language—assumes that mental concepts can be understood as discrete *contents*, Beaver and Stanley explicitly reject this view. They argue that the commonplace⁶ container/conduit metaphors are deeply misleading, writing that though the word 'content' has been somewhat divorced from the notion of containment in technical talk about meaning,

[...] it is well to remember that talk of words as something like boxes with *content* or trays that bear meaning might easily lead any of us quite unconsciously into a simplistic way of thinking, whereby meanings are treated as fixed entities that can be carried in or by an expression (Beaver and Stanley, Pg. 29).

This container metaphor manifests itself in the disciplinary divide between philosophy of language, which deals with how we arrive at a speaker's meaning, and epistemology, which deals with the judgment we make upon that meaning. Because it has become commonplace to think of words as boxes of meanings that can be unpacked, and because when opening a package, we first determine what is inside before making a judgment as to whether we would like to accept the contents, philosophers have tended to assume that the process of determining what something means can be studied separately from the process of making a judgment about said meaning. But this assumption leads to several problems, each of which can potentially be avoided by viewing communication through a different metaphorical lens which compares meanings not to containers and contents, but to resonances and harmonies.

Idealizations can be useful for understanding and theorizing about the world, but they can also lead us astray by averting our attention from useful information. As Beaver and Stanley convincingly argue, there is good reason to think that the deeply entrenched container/content metaphor for communication has led philosophers astray by averting attention from the connections between speech practices and human behavior. By implying that the processes of comprehension and integration are independent, the container metaphor abstracts from the form of the original message being processed, focusing instead on the pure information *contained* in

⁶ See Lakoff & Johnson (1980) for an exposition of the ways in which the metaphor of meanings as contents and communication as the conduit for the delivery of those contents permeates our everyday conversation and hence, our common way of thinking about language and communication.

the message. And by causing us to think of meanings as discreet propositions, the container metaphor prevents theorists from providing insight into how language can produce emotional reactions or incite action (pg. 122). As J.L. Austin famously acknowledged, communication is not the exchange of sterile propositions or sentences on a blackboard, but actions which enact states of affairs (Austin,). Hence, if the right conditions hold, saying “I do” during a marriage ceremony enacts a marriage, and the commands of a military officer enact war and cause soldiers to kill. The idea that meanings are isolated mental contents prevents theorists from exploring the rich effects that words have on human behavior.

Perhaps Recanati’s (AP) from chapter one can offer useful wisdom here.

(AP) In deciding whether a pragmatically determined aspect of utterance meaning is part of what is said, that is, in making a decision concerning what is said, we should always try to preserve our pre-theoretic intuitions on the matter (Recanati, 1989).

At the heart of AP is the idea that theorists should first examine the evidence and then craft their theories to fit with the evidence, instead of interpreting the evidence in a way that fits with the theory. Arguably, many philosophers of language have been guilty of the latter sin, interpreting linguistic data in a way that fits with the notion that meanings are discrete mental contents instead of rejecting the content theory to better account for non-propositional aspects of meaning such as force. In line with this perspective, Beaver and Stanley define meanings as resonances, and resonances as statistical correlations.

The resonances of an action are the features of the context whose probability has increased given that the action has occurred. Put more succinctly,

An action in the extension of a practice is a (positive) resonance of a feature of context to the extent that an occurrence of the action changes the probability of that feature (positively) (Beaver and Stanley, Pg. 43).

Consider how this view could apply to the resonance or meaning of the expression “ouch.” Because the probability that the speaker is in pain given that they uttered “ouch” is greater than the probability that they are in pain in otherwise, being in pain is a positive resonance of each use of the word “ouch.” Similarly, the probabilities that the speaker is alive, human, and an English speaker, each increase (to greater and lesser degrees) given the use of “ouch”. Formally, the strength of the resonances of an action are given by the difference between the probability⁷ of each feature given the action, minus the probability of the feature. This can be represented in mathematical terms, where the strength of the resonance is given by the following formula: $p(\text{feature} \mid \text{Action}) - p(\text{feature})$.

⁷ Here, ‘probabilities’ are defined in frequentist terms; they are nothing more than measurement of the frequency at which events co-occur.

A few terms must be clarified here. What is a ‘context’ and what is a ‘practice’? Beaver and Stanley define the context as a *centered-possible world*, or “[...] the way things are, but relativized to parameters that are significant for a conversational interaction” (Pg. 43). For instance, contexts are centered at a particular date, time, and location allowing them to contain the set of things happening right now, two hours before/after now, and things happening next door. ‘Practices,’ on the other hand, are characterized as actions that are *habitual*, and *rule governed*. Practices are habitual because they must be established through repetition. By randomly deciding to put walnuts in my oatmeal instead of almonds, I am not performing a practice of putting walnuts in oatmeal. On the other hand, by eating oatmeal, I am performing the practice of eating oatmeal; the performance is habitual and learned through repetition, and there is a recognizable distinction between what does and does not count as an extension of the practice. Practices are rule governed because they are recognizable—observers can recognize similarities across all instantiations of the practice—and there are criteria which determine whether something counts as an extension of the practice.

To unpack, consider how resonance-theory contrasts with the views previously presented in this chapter. Unlike truth-conditional semantic theories which define meanings in terms of psychological states, the definition of resonance makes no reference to psychological states. It is possible for psychological states to be resonances of actions, but it also possible for behaviors, practices, and states-of-affairs to be resonances. Psychological states play no special role. Furthermore, like Putnam’s externalism, resonance theory allows that meanings are external to the mind. But unlike Putnam’s externalism, resonance theory goes further to claim that all terms have resonances (i.e. meanings) independently of our psychological states.

Whereas previous theories have been implicitly bound by the conduit/container metaphor, resonance theory breaks free, accounting for the ways in which meanings are non-categorical, non-propositional, and even political. Consider an utterance of (27).

(27) “I majored in philosophy as an undergraduate student.”

It seems reasonable to presume that among the many positive resonances of (27) are (28) through (32).

- (28) The speaker of (27) knows something about the ideas of Plato and Socrates.
- (29) The speaker of (27) has gone to college.
- (30) The speaker of (27) speaks English.
- (31) The speaker of (27) knows that $\Phi \rightarrow \Psi$, $\Psi \models \Phi$ represents an invalid argument.
- (32) The speaker of (27) dislikes relativism.

Resonance is non-categorical in the sense that it is scalar; whether something is or is not a resonance of an action cannot be represented in binary because resonance is a gradient notion. (28) is likely a very strong resonance of (27), whereas (31) and (32) are—I might conjecture—weaker resonances of (27). Resonance is non-propositional because anything (i.e. not just

propositions) can be a resonances. For instance, the proposition, *the speaker of (27) went to college*, is not a resonance of (27). Rather, it is the *fact* that the speaker of (27) went to college that is a resonance of (27). Finally, resonance is political in the sense that the resonances of our actions can be ideologies, identity affiliations, and practices. Again, to engage in conjecture, the probability that the speaker of (27) possesses some democratic ideals is a resonance of (27). If a hearer is attuned to these resonances, then they might implicitly come to recognize the political beliefs of the speaker.

In sum, resonance theory can be understood as a more radical version of Putnam's externalist theory of meaning. According to resonance theory, meanings are not isolated contents in our head, but social entities defined by the community; by using terms in specific contexts, people slowly craft the resonances of those terms. Since there is no limit to the number of features of the context whose probability is increased with the use of a term, the meaning of a term goes far beyond what it is typically taken to mean. Furthermore, since it would be practically impossible for anyone to recognize all the resonances of the commonly used terms in their linguistic community—they would have to be omniscient to know every feature of the context whose probability increases given the use of each term—no one fully grasps the meaning of any commonly used terms. This is, perhaps, a welcome consequence of resonance theory, but it also points to a larger question. If resonance is a purely external, non-psychological notion, then how or by what mechanism do interlocutors become *attuned* to resonances?

3.3 Interpretation as Attunement

Beaver and Stanley define attunement to a practice as follows:

An agent is **attuned** to a practice to the extent that their state and behavior predictably evolve in accordance with its presence in the agent's context (pg. 80).

The key terms in this definition are *predictability*, *accordance of state and behavior* and *presence in an agent's context*. Let's briefly consider each of these, in turn. Beaver and Stanley do not provide an airtight definition of predictability, and in fact point out that doing so is inherent difficulty because, "[...] we are now talking not simply about whether two things co-occur, but about whether a complex behavior is responsive to something's presence" (pg. 70). Nevertheless, for present purposes, let us think of predictability of a state and behavior in terms of whether members of the agent's community would be capable of recognizing that the state and behavior are likely given that some practice has occurred. By performing certain actions associated with a practice, people who are engaged in the practice or who are attuned to the practice may recognize that I am performing the practice. Moving on, *accordance of state and behavior* can be defined as some complex relation between the presence of a practice and the agent's way of acting in the world. For instance, in the presence of the practice of driving, I stop at red lights, stay in my lane, and try not to hit other cars. Finally, *presence in the agent's context* can be understood in two ways; either I am an active participant in the practice, or am external to the practice. For instance, being in the presence of the practice of breakdancing is different from

actually being engaged in the practice of breakdancing, and both may affect my behavior in different but *predictable* ways.

Now, having roughly outlined what it means to be attuned to a practice, it is natural to ask *how* we come to be attuned to certain practices—specifically, the present question is how we become attuned to speech practices. Since the use of words in a language is habitual and rule-governed, the use of words is a practice. Hence, the question of how we become attuned to the meanings or resonances of the words in our language is (broadly speaking) just the question of how we become attuned to the resonances of a practice.

To address the latter question, Beaver and Stanley begin by assuming a roughly associationist theory of mind. Though they explicitly state that association is just one of many mechanisms through which the mind makes sense of the world, their explanation of how we become attuned to resonances is founded on associationism. They write,

Just as Hume considered that thoughts can become associated with each other when the sense stimuli that trigger the thoughts regularly co-occur, so we will consider under what conditions attunements can become associated with, or dissociated from, each other. On the view we have developed, people associate a word with its resonances. That is, the resonances of communicative actions just are associations, albeit associations created and perceived collectively by the members of a community of practice (pg. 138).

So, it is in virtue of the minds propensity to create associations between events that frequently co-occur that I have come to associate use of the word “ouch” with the presence of the speaker’s state of pain. The community shares these associations, and it is partly due to these shared associations that words have their resonances. Furthermore, people often embody their attunements; when a speaker says “ouch”, hearers often flinch, empathizing with the speaker’s state of pain. Hence, it is possible to be attuned to the attunements of others by recognizing that the probability of them having certain attunements is increased by the presence of their behaviors.

3.4 Resonance as a Top-Down pragmatic Process

Now, let us finally return to the problem which was introduced in chapter one; how do hearers determine what is said? The relevance-theoretic account answers this question by holding that hearers default to the interpretation of what is said which maximizes relevance. On the other hand, resonance-theory answers the question by holding that the resonances of a speaker’s utterance provides a sufficient constraint on the set of potential meanings by predicting which interpretations the hearer is likely to be attuned to. Here, I examine this alternative account of the determination of what is said.

Let’s consider our previous examples. Recall (19).

(19) The child left the straw in the glass.

According to the relevance-theoretic account, the reason hearers tend to default to the *drinking-tube* interpretation of (19) is that, unlike the *cereal-stalk* interpretation, it has a significant number of contextual effects in a stereotypical context that can be accessed with minimal effort. Contrastingly, on the resonance-theoretic account, the reason that hearers tend to default to the *drinking-tube* interpretation is that, unlike the *cereal-stalk* interpretation, it is a strong positive resonance of (19). In other words, in most situations where an utterance like (19) is used, the speaker tends to be referring to a drinking-tube. Assuming hearers are attuned to this resonance, they will default to the *drinking-tube* interpretation.

Recall (3).

(3) Pass me three screws.

Similar to the above example, on the relevance theoretic account the reason hearers tend to default to the *exactly three screws* interpretation of (3) is that unlike the *at least three screws* interpretation, the *exactly three* interpretation has a significant number of contextual effects in easily accessible contexts that are stereotypically associated with utterances like (3).

Contrastingly, on the resonance-theoretic account, the reason that hearers tend to default to the *exactly three* interpretation is that it is a strong positive resonance of (3). Rarely are utterances like (3) used in situations where the speaker wants *at least three* screws. Assuming hearers are attuned to this resonance, they will default to the *exactly three* interpretation.

Note that, like relevance-theory, resonance theory is consistent with the moderate contextualist view that top-down pragmatic processes enter into the compositional process. Indeed, resonance can be understood as a strange sort of top-down pragmatic process, which determines what is said in a statistical sense. What is said changes depending on the context, because the use of word in one context is a different event type than is the use of the same word in a different context. Hence, the word will have different resonances, and hence different meanings, in different contexts.

Now, before attending to glaring objections to this view, consider how it avoids the problems of circularity and underdetermination facing relevance theory. As demonstrated in section 2.3.1, the principle of relevance leads to an inevitable circularity because the relevance of a given context depends on which assumption is chosen, and the relevance of a given assumption is dependent on which context is chosen. The concept of resonance is an objective, non-psychological measure, and hence, faces no such problem. Either the expected interpretation is also the one with the highest degree of resonance (in which case, resonance-theory makes a correct prediction) or it is not, and resonance theory makes a false prediction. Hence, there is no need to engage in a circular process of searching for a relevant context to determine the relevance of an assumption and then determining the resonance of the assumption to determine the relevance of each context.

On the other hand, resonance theory also avoids the problem of underdetermination facing relevance theory, but the problem admittedly returns in a different form in section 3.5. As demonstrated in section 2.3.2, relevance theory ends up relying on unspecified heuristics and comparative abilities to determine the most relevant interpretation of what is said. But given that the explanatory burden weighs on these undefined heuristics, the semantic composition process is left completely underdetermined. Contrastingly, resonance theory does provide predictions as to which interpretations of what is said will be most relevant. It is undeniable that different features of the context are positively correlated to a given action or utterance to different degrees. Resonance theory predicts that the strongest correlations are what determine what is said.

3.5 Novelty and Mystery

Now to tend to the objections. The problem of novelty is that some actions have never been executed before. Hence, those actions do not have any resonance; there is no probability change in any feature of the context given that the action has been performed. One way of addressing this problem is by allowing features of the context to be resonances of abstract *types* of actions, not just specific instantiations. Furthermore, the problem could also be addressed by allowing attunements to guide, but not fully determine, the interpretation process. But this leads to the second serious objection to the resonance-theoretic framework—mystery. The concept of attunement tells us little about how the mind actually processes speech acts.

3.5.1 Novelty

Consider (33).

(33) Ben's friend, Tom, said that he ate a decadent breakfast at a small diner on planet Saturn.

Presumably, this is the first time (33) has ever been used. Hence, it should not have any resonance. But it nevertheless has a clear meaning. How does the mind interpret completely novel speech acts?

One way of accounting for novel sentences like (33) on the resonance-theoretic framework is by allowing that resonance be applied to abstract types of actions, and to be measured in a limited domain. Let's start with the first word of (33), 'ben'. In contexts where the referent of a proper name is unclear, there is an increased probability that the speaker intends the hearer to supply a random indelible human referent. Hence, one meaning or resonance of the action of saying 'Ben' is that the speaker intends to conjure an indelible referent. Now we are in a domain where the speaker has just conjured an indelible referent. In domains where the speaker has just conjured a referent, the action of using the word 'friend' has the resonance of referring to a person in a specific 'friend-relation' to the referent. Now we are in a domain where the speaker has conjured an entity called 'tom', in a friend-relation to some indelible referent, 'Ben'. This type of explanation could be carried through the whole sentence. The meaning of the sentence ends up being the resonance of the whole action of using the sentence in a context. This

action cannot be characterized in simple terms as it constitutes an extremely complex nexus of practices.

This way of accounting for the problem of novelty leads to a view of interpretation where the meaning of an action is a complex type that constrains interpretation but does not fully determine it. Though (33) has likely never been used before, sentences of the same type as (33) have been used. And so, if you allow types of actions to have resonances which are, essentially, complex grammatical practices, then, in a restricted sense, (33) does have a resonance.

Now, it should be emphasized that resonances do not compose. At least, not in any traditional sense of compositionality. However, actions do compose, and complex actions can be built from primitive actions or practices to produce complex actions. By ‘compose’, I mean that some actions are defined by the occurrence of sub-actions. Writing a check, for instance, is constituted by an act of signing one’s name with a pen and writing a dollar amount. Each of these sub-actions compose to produce the complex action. In some cases, these complex actions constitute practices. Hence, we can think of (33) being composed of the smaller sub-actions of uttering (34), (35), (36), and so on.

(34) Ben’s

(34) Ben’s friend

(36) Ben’s friend, Tom

[...]

(n) Ben’s friend, Tom, said

Each of these complex actions has resonances, and so do each of the individual parts. The meaning of the whole can be understood as the resonances of the whole complex action, which not only includes the entire utterance, but also other aspects of the context.

3.5.2 Mystery

Assuming that words *can* have abstract types as resonances and that those abstract types *can* constitute complex grammatical practices, how do we come to recognize them? Again, the resonance theoretic answer is that we learn to be attuned to those resonances through association. Simply put, by being immersed in the presence of complex grammatical practices, we automatically start to develop the sorts of attunements necessary to perform them. This can simply be attributed to the mind’s propensity to create associations between events. But this view seems unsatisfying. When two things become so deeply associated in the mind that the correlation between their occurrence is near 1, we might call that association a rule. Does this view sound familiar? Relevance theory accounts for how the mind determines what is said by appealing to a complex set of rules that govern how the mind processes information. Hence, the resonance-theoretic way of dealing with the meaning of novel assumptions essentially collapses back into a structured, rule-governed account of how the mind processes information. The more strongly associated two types of mental representations are, the more evidence we have for the

existence of a law, doing causal work behind the scenes. Especially when the associations we are talking about involve abstract types of events, as those cannot be simply subsumed under association of even tokens or instantiations.

It might be argued that Beaver and Stanley go too far in rejecting the traditional content-based theories of meaning. By focusing on the ways in which the container/conduit metaphor misleads us, they lose sight of its explanatory power, rejecting theories of mind which have the potential to complement their view in positive ways. Instead of adopting a computational theory of mind and using it to complement the idea of resonance, Beaver and Stanley opt to assume an outdated associationism which provides no explanation of how the mind processes information.

Conclusion

In chapter one, I argued for a moderate contextualist perspective on the semantic composition of what is said. This position is partially motivated by Recanati's availability principle, which states that we should adhere to our pre-theoretical intuitions when determining what is said. Both Sperber and Wilson's and Beaver and Stanley's theories fit roughly within the moderate contextualist framework I advocated for, but the latter remains true to the spirit of the availability principle, whereas the former does not. By defining meanings in terms of resonances, Beaver and Stanley do not arbitrarily abstract away from the connections words have to human behavior. On the other hand, using abstraction arguably allows relevance theory to provide a more adequate theory of how the mind processes information.

In chapter two, I present Relevance theory, which constitutes both an account of how the mind represents and processes information, and an account of how it determines which information is worth processing (i.e. which information is relevant). This account is grounded in, and hence builds from, an explanatorily powerful and empirically vindicated theory of mind. But the principle of relevance faces serious theoretical challenges; it fails to fully determine what is said and make testable empirical predictions.

Finally, in chapter three, I present Resonance-theory. This is my own adaptation of Beaver and Stanley's radical theory of meaning, which I created to address the interpretation of what is said. In line with Putnam (1973), Resonance theory defines meanings not as psychological states, but as statistical correlations. This allows for 'meaning' to be understood very broadly. Furthermore, this theory provides a more predictive account of what is said, thereby addressing some of the problems facing relevance theory. However, Resonance theory also faces serious theoretical challenges. It does not explain how meanings are processed and composed to produce complex representations in the mind and hence cannot account for novel utterances.

Ultimately, I hope to have shown that there is significant promise to combining relevance theory, which constitutes a psychological account of how the mind processes information, with resonance theory, which constitutes a more sociological account of meaning and human

behavior, to yield a more powerful account of how hearers determine what is said. Relevance theory provides a plausible account of how the minds represent and process information—this constitutes the groundwork needed for an account of how sentences get their meanings from their parts while allowing context to enter the semantic composition process. On the other hand, Resonance theory provides a way of determining what a speaker means in a particular situation. The scope of possible practices that the speaker could be using is determined by the resonance that their complex actions have in contexts of abstract types.

Admittedly, the latter part of this paper has been somewhat vague about exactly how resonance theory should be fleshed out and combined with relevance theory. This is an open question, and in the spirit of the availability principle, it would be unwise to provide simple abstractions which miss out on the breadth of linguistic and behavioral data available. Much more work must be done, but it is clear that any adequate account of the determination of what is said should explain both the psychological processes involved in interpretation, and the sociological processes involved in determining the resonances or meanings that words and practices carry.

Citations

- Austin, J. L. (1961). *Performative Utterances*. In John Langshaw Austin, *Philosophical Papers*. Oxford, England: Clarendon Press.
- Bianchi, Claudia (2010). *Contextualism*. Handbook of Pragmatics Online.
- Beaver, David & Stanley, Jason (2023). *The Politics of Language*. Princeton University Press.
- Davidson, Donald (1967). *Truth and meaning*. Synthese 17.
- Fodor, Jerry A. (1975). *The Language of Thought*. Thomas Y. Crowell Company, Inc.
- Grice, Paul (1975). *Logic and Conversation*. In Donald Davidson, *The logic of grammar*. Encino, Calif.: Dickenson Pub. Co.. pp. 64-75.
- Lakoff, George & Johnson, Mark (1980). *Metaphors we live by*. Chicago: University of Chicago Press. Edited by Mark Johnson.
- Larson, Richard K. & Segal, Gabriel M. A. (1995). *Knowledge of Meaning: An Introduction to Semantic Theory*. MIT Press.
- Putnam, Hilary (1973). *Meaning and reference*. Journal of Philosophy 70.
- Quilty-Dunn, Jake ; Porot, Nicolas & Mandelbaum, Eric (2023). *The best game in town: The reemergence of the language-of-thought hypothesis across the cognitive sciences*. Behavioral and Brain Sciences 46.
- Recanati, François (2010). *Truth-Conditional Pragmatics*. Oxford, GB: Oxford University Press.
- Recanati, François (1989). *The Pragmatics of What is Said*. Mind and Language 4.
- Seuren, P. A. M. (1987). *How relevant?: A commentary on Sperber and Wilson 'Précis of relevance: Communication and cognition.'* Behavioral and Brain Sciences, 10, 731-733.
- Sperber, D., & Wilson, D. (1995). *Relevance: Communication and cognition* (2nd ed.). Blackwell Publishing.