© Copyright 2011 Nathalie Morasch

UNDER WHAT CONDITIONS ARE TWO UTTERANCES PEFORMANCES OF THE SAME WORD?

BY

NATHALIE MORASCH

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Philosophy
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2011

Urbana, Illinois

Doctoral Committee:

Professor Garry Ebbs, Chair
Associate Professor Steven Wagner, Co-Chair
Professor Timothy McCarthy
Professor Peter Lasersohn

ABSTRACT

Starting from the fact that people sometimes use the "same" words to talk about a given topic, I want to clarify what word-sameness comes to in those uses. I will adopt an epistemic framework, in which words are primarily instruments that render the inter-subjective transfer of knowledge possible.

In the course of my dissertation I refine Kaplan's notion of words to propose an account that occupies the middle ground between (social) anti-individualism and the kind of individualism that individuates a speaker's words without input from the speaker's linguistic community. I make the case that speakers keep track of the various performances of a given word *w* via a *mental register*. According to my proposal, the conceptions the speaker comes to associate with her mental register over time may play a role in whether we ought to identify her idiolectal word *w* with the public word *w'*. I will argue that in the end we must leave it up to the speaker's own (informed) judgment whether she interprets her word *w* as repeating a particular public word *w'*.

According to Kaplan the individual speaker's word *w* is referentially bound to the public word *w'* through her intentions to repeat *w'*. I stress that the intention to repeat her own words *w trumps* the intention to repeat the word *w'* produced by some other speaker. I attempt to solve Kripke's Paderewski puzzle and problems of self-knowledge by arguing that the speaker cannot be wrong about how she keeps track of her own words.

ACKNOWLEDGMENTS

It was one of Dr. Gary Ebbs's courses on external world skepticism that got me fascinated with Putnam's "The Meaning of 'Meaning'." From this seed my interest in philosophy of language broadened and subsequent courses led by Dr. Ebbs, Dr. Steven Wagner, Dr. Timothy McCarthy and Dr. Peter Lasersohn firmed the roots of my absorption in this subject. I would like to thank Dr. Lasersohn for encouraging me to explore the overlap of philosophy and linguistics. Thanks to Dr. McCarthy for reconnecting the mosaics in my head resulting from Dr. Ebbs's scrupulous attention to detail. Dr. Ebbs, I am grateful for your meticulous, plentiful and always insightful corrections. With patience you guided me in focusing my curious mind. One couldn't wish for a better advisor! My indebtedness to Dr. Wagner far surpasses the academic realm. Thank you for getting lost in time and spending many "over-caffeinated" hours at "Paradise Café" with me. I am still learning from your courage of facing the world with open, but optimistic eyes.

Thank you, Kyle, for ruthlessly correcting me when I am wrong. It would have been very lonely during my first years in the United States without you as a genuine friend and much more.

Liebe Eltern, vielen Dank für eure Unterstützung selbst in Bestreben die für euch nicht leicht zu akzeptieren waren.

Thanks to Paul for loving also the philosopher within me. You always attentively listen to and help me refine yet another thought experiment.

TABLE OF CONTENTS

CHAPTER 1: UNDER WHAT CONDITIONS ARE TWO UTTERANCES PERFORMANCES OF THE SAME WORD? PROPOSAL: PHYSICAL RESEMBLANCE 1

CHAPTER 2: UNDER WHAT CONDITIONS ARE TWO UTTERANCES PERFORMANCES OF THE SAME WORD? PROPOSAL: RESEMBLANCE OF COGNITIVE CONCEPTIONS. 25

CHAPTER 3: UNDER WHAT CONDITIONS ARE TWO UTTERANCES PERFORMANCES OF THE SAME WORD? PROPOSAL: APPROPRIATE INTENTIONS. 47

CHAPTER 4: UNDER WHAT CONDITIONS ARE TWO UTTERANCES PERFORMANCES OF THE SAME WORD? PROPOSAL: RETROSPECTIVE JUDGMENTS. 97

CHAPTER 5: ARE INDIVIDUAL SPEAKERS RELIABLE WHEN IT COMES TO INDIVIDUATING THEIR OWN WORDS? TEST-CASE: KRIPKE'S PADEREWSKI PUZZLE. 138

CHAPTER 6: ARE INDIVIDUAL SPEAKERS RELIABLE WHEN IT COMES TO INDIVIDUATING THEIR OWN WORDS? TEST-CASE: SLOW SWITCHING. 179

BIBLIOGRAPHY. 211

CHAPTER 1

UNDER WHAT CONDITIONS ARE TWO UTTERANCES PERFORMANCES[1] OF THE SAME WORD? PROPOSAL: PHYSICAL RESEMBLANCE

1.0 Introduction

The transmission of *knowledge* across speakers and intra-subjective reasoning rely on our ability to express *one and the same* thought from one occasion to the next. With "same thought" I mean something of the following: You tell me: "Syrus is a dog." Days later I reason: "If Syrus is a dog[2], then Syrus needs daily exercise." I recall your utterance: "Syrus is a dog." So I conclude: "Syrus needs daily exercise." The validity of my *modus ponens* reasoning depends on the respective elements of my utterances expressing the *same thing.*[3] For example, the consequent expressed by the last part of my first utterance, and the conclusion expressed by my last utterance must *say the same thing*. Moreover, your original utterance and my report of it must say the same thing. I can't rely on information[4] offered by other people if I don't, at the very least, *understand* what their utterances state.[5]

---

[1] The term "performance" is taken from a forthcoming paper from Lepore and Hawthorne entitled "On Words."

[2] A young and healthy dog, that is.

[3] See also Fine (2009): p. 1.

[4] I am using a very narrow notion of "information" here; with "information" I mean the content of an utterance.

[5] I am assuming here that we can theoretically target sentence meaning independently of the pragmatic meaning of an utterance. I will not be dealing with problems of contextual enrichment or implicature. Some utterances merely provide a "semantic skeleton", which needs to be fleshed out via contextual enrichment such as to make the respective utterance propositional (Recanati (2004): chapter 1). Consider the following utterance by John, a student getting ready to host a party: "Every beer is in the fridge". John clearly doesn't mean to say that every beer *in the universe* is in the fridge. There are multiple candidate propositions:

P_1 Every beer *we will serve later* is in the fridge.
P_2 Every beer *we bought at the corner store* is in the fridge.
P_3 Every beer *for the party* is in the fridge.

Starting from the fact that people sometimes use the "same" words to talk about a given topic (i.e. the individual Syrus or dogs in general), I want to clarify what word-sameness comes to in those uses. I will explore what notion of words is pre-supposed by our linguistic communicative exchanges. In this epistemic framework, words are primarily instruments or tools that render the inter-subjective transfer of knowledge – about the co-occupied environment (conceived broadly) – possible.[6]

In my above reasoning about the individual Syrus, I rely on knowledge acquired via *testimony*. I cannot employ a statement *S* of yours in an argument, unless I, at the very least, *understand S* – which presupposes that I understand the words you are using in the respective utterance.[7,8]

P_4 Every beer *in the kitchen is* in the fridge.
The problem with multiple candidates is that even if John had uttered: "Every beer is in the fridge" with any particular proposition in mind (which he need not), the audience can still understand the utterance without entertaining that same proposition (Buchanan & Ostertag (2005)).
Grice considers the following utterance: "Jones has beautiful handwriting and his English is grammatical" Grice (1961). The speaker may have intended to convey any of the following propositions:
P_1 Jones is/was a poor student.
P_2 Jones doesn't have the intellectual virtues required for graduate work.
P_3 Jones should not be admitted to the graduate program.
P_4 I am not in the position to say anything stronger.
P_5 Jones is not prepared for graduate work.
Again, the audience may have understood the original utterance without having entertained any particular proposition of the list above**.**
I will therefore concentrate on utterances with as little contextually variant meaning as possible. Even my conclusion: "Syrus needs daily exercise" is contextually variant. With the proper name 'Syrus' I designate a *specific* dog of that name, not some dog or other named 'Syrus'. When I specify 'daily exercise' I don't have in mind a stroll around the block, I have in mind 3-5 miles of brisk walking.
[6] According to Evans "communication is *essentially* a mode of transmission of knowledge" (Evans (1982): p. 310). This, he explains "should debar us from certifying anything as successful communication which is *intrinsically* incapable of conveying knowledge" (ibid.: p. 320).
Other works in which this framework is endorsed: Heck (1995); Goldberg (2007a); Ebbs (2009).
[7] We must also assume that I understand the compositional structure of my utterances and yours, that is, I understand what each contributes to the meaning of the utterance as a whole.
[8] I will leave other epistemic conditions for the reliable transfer of knowledge via testimony aside and focus merely on the linguistic aspects.

I will take for granted that each speaker knows the "semantic content" of their own words, that is, I understand what my (present) expressions state[9] and am entitled to state biconditionals of the following kind:

(a) *x* satisfies my word '<u>dog</u>' if and only if *x* is a <u>dog</u>[10,11]

I could then attempt to express the "meaning" of your words in terms of my own. This allows us to theoretically circumvent the provision of a general account of meaning, or semantic content. For example, the meaning of your word 'dog' I can state thus:

(b) *x* satisfies my word '<u>dog</u>' if and only if *x* satisfies your word '<u>dog</u>'[12,13]

If your word 'dog' and my word 'dog are one and the same[14], then (b) is true.[15] Finally, I can *express* what your word 'dog' means:

(c) *x* satisfies your word '<u>dog</u>' if and only if *x* is a <u>dog</u>[16]

Suppose we have a longer conversation about the origins of the dog and how the domestication of the latter is inherently connected to our own evolutionary anthropology. You will utter

---

[9] Burge (1988).
[10] Tarski (1994) was the first to propose a definition of truth in terms of satisfaction conditions.
[11] Even though I here propose to define sentence meaning in terms of satisfaction conditions (which is controversial), nothing in my account depends on such a disquotational approach. If, one prefers to talk about sentence meanings using propositions, this is compatible with my account.
[12] If you used the word 'dog' in such a way that it is satisfied by all and only fish, I could not extend a *judgment of sameness of satisfaction* to your word 'dog' as described above. I would have to say instead: "*x* satisfies my word 'dog' if and only if *x* satisfies your word 'fish'."
[13] Ebbs explains how such intra-subjective word-individuations are learned, yet typically non-deliberative; they constitute nonetheless *judgments* because we will reject such communicative assumptions such as (b), if we come to think of them as wrong (Ebbs (2009): p. 5).
[14] As I will discuss in this chapter, phonetic and graphic identity are not sufficient conditions for two speakers to be using the *same* word. If you only know the word 'bank' denoting a financial institution and all I am familiar with is the word 'bank that denotes the slope along the river, then I *cannot* state the satisfaction conditions for your word 'bank' thus:

x satisfies my word 'bank' if and only if *x* satisfies your word 'bank'.

[15] Usually when two speakers belong to the same linguistic community, they are able to state satisfaction conditions of each other's words (as described above). Parts of my dissertation deal with isolated exceptions to this normality.
[16] See also Ebbs (2009): chapter 4.

several tokens[17] "dog_1", "dog_2", "dog_3" … "dog_n" and I will produce a multitude of tokens "$dog_{1'}$", "$dog_{2'}$", "$dog_{3'}$" … "dog_m" of my own. In order for us to have a *successfully communicate*[18], I need to be able to do more than merely "link" an individual word token to one of your word tokens. I must be able to keep track of my own tokens as instantiating a unified word-*type* and also identify your individual tokens as the *same* word-type 'dog'. So, in (a) through (c) we are actually referring to words as *types*.

Let's take a closer look at (a), which is not as trivially true as it seems at first. If the two underlined tokens constitute performances[19] of the same word, then I stated something true.[20] We may realize that the two word tokens are spelled the same, but that doesn't suffice for them to constitute tokens of the same word type. Consider the following bi-conditional:

(d) *x* satisfies my word '<u>robin</u>' if and only if *x* is a <u>robin</u>[21]

If in (d) the first underlined word token is an instance of the British English word 'robin' and the second underlined word token is an instance of the American English word 'robin', then the bi-conditional is false. It is intuitive to assume that we simply don't make such mistakes with regard to classifying our own word tokens. The European Robin and the American Robin are distinct species of birds.

---

[17] Peirce used the terms 'type and 'token' to indicate a distinction between two senses of the word 'word'. Although there is only one definite article, 'the', in English, there are apt to be about twenty instances of it on a single page (Peirce, Charles Sanders (1931-58)).

[18] That is, have the kind of conversation where the transfer of knowledge is possible.

[19] When I talk about Aristotle the shipping magnate and utter: "Aristotle…", then I thereby performed the name 'Aristotle'. If, at a different occasion I talk about Aristotle the philosopher and utter: "Aristotle…", then I thereby performed the name 'Aristotle'. Those names sound and look very similar, but they do not constitute performances of the same word. If I utter: "'Aristotle' contains nine letters", then I also performed the name 'Aristotle' – I mentioned it. Perhaps use and mention are different *modes* of presentation.

[20] The biconditional can hold true, without the used word being the same as the mentioned word. For example, I could say: "*x* satisfies my word 'urchin' if and only if *x* is a hedgehog" and thereby utter a true biconditional. Whenever the words in question are synonyms of one another, the biconditional is true.

[21] Example taken from Ebbs (2009): pp.114-115; see also Putnam (2001): p. 22.

Words are types instantiated in utterances as tokens. Ebbs explains how most theories of meaning subscribe to a two-step word individuation process:

(U) Two word tokens *t* and *t'* are of the same word type if and only if
(i) *t* and *t'* are each tokens of the same orthographic or phonetic type, and
(ii) facts about the explanatory-use[22] of *t* and *t'* determine that they each have semantic values and that their semantic values are the same.[23]

It is agreed that tokens of the word 'bank' denoting a financial institution and 'bank' denoting the slope of a river are not instances of the same word-type, even though they share the same orthographic or phonetic type. It is thus tempting to conclude that there must be some *additional* facts that divide the tokens in question into separate types. It is these facts that Ebbs tries to capture with the expression 'explanatory-use': the kinds of facts that supposedly determine the meaning or satisfaction conditions of a given token (i.e. the cognitive content of the speaker, causal-historical facts, the social environment of the speaker or the physical environment of the speaker). [24]

In the present chapter, I will show that there is no intrinsic relationship between word tokens that would allow us to define an orthographic or phonetic type – beyond a perceived relation. As a result, the individuation process collapses into step (ii).[25]

1.1 Proposal: all and only tokens of a given word type share a common "form"

We communicate our thoughts through physical word tokens, or marks.[26] Marks are physical objects such as inscriptions in ink, the composite sound coming out of the mouth of a speaker,

[22] This terminology is introduced by Ebbs (2009).
[23] Ebbs (2009): p.112; Ebbs rejects the token and explanatory-use model of words (U).
[24] Ebbs (2009): p.6.
[25] We could attempt to reword (i) into something of the following: t and t' are each *perceived* as tokens of the same orthographic or phonetic type by the appropriate judges. Of course, much will depend on who is entitled to make such evaluations. The multiplicity of ways in which various tokens of a given type can be realized in are still going to constitute serious challenges for (U).

hand signals, raised dots, ripples in the sand – just to name a few. Some of those marks, i.e. a blob of ink that I accidentally shook out of my fountain pen, don't represent anything; others are replicas[27] of pictures or meaningful strings of characters such as words or numerals – and eventually utterances. Can we specify what all and only marks of the word 'dog' have in common? Do they perhpas all share the same distinct shape, form or pattern?

A word can be realized in a variety of symbolic systems: The word 'dog', for instance, can be physically conveyed through a visual (i.e. a grapheme, hand signal, short-hand symbol, or Gödel Number), auditory (i.e. a phoneme or sequence of Morse Code) or tactile sign (i.e. a symbol in Braille).[28] The grapheme and phoneme for the word 'dog' don't have anything in common, capable of recognizing them as a unified pair distinct from tokens of the word 'cat', except that they are tokens of the word type 'dog' and not 'cat'. Not only do we *not* see a common form among all and only word tokens of the word type 'dog', we couldn't even come up with *some* shared equivalence relation.[29] Defining the word types in terms of a *resemblance* relation among the tokens of the type comes full-circle; the only distinctive resemblance between the utterance 'dog' and the written inscription 'dog' is their membership in the same word type.[30] In fact, they

---

[26] 'Mark' is a technical term adopted from Goodman (1976) and also found in Ebbs (2009). Goodman may in turn have adopted it from Peirce, who writes: "For a "possible" Sign I have no better designation than a *Tone*, though I am considering replacing it by "Mark". […] An Actual Sign I call a *Token*; a Necessitant Sign a *Type*" (Peirce (1908, 1998): p. 480). Goodman's technical term 'mark' seems to be better translated as Peirce's 'token' than Peirce's 'mark'.

[27] 'Replica' is a technical term adopted from Goodman ((1976): chapter 4), who in turn may have adopted it from Peirce: "Every legisign [i.e. word] signifies through an instance of its application, which may be termed a *Replica* of it. […] Each single instance of [the single word "the"] is a Replica" (Peirce (1903,1998): p. 291. Goodman's nominalist notion of replica is, however, missing in Peirce. You can find Goodman's notion of marks developed further by Ebbs ((2009): chapter 4).

[28] As far as I know we don't have a symbolic system that relies on our gustatory or olfactory senses.

[29] With the exception of perhaps mental word tokens?

[30] For those who wonder whether we couldn't say that two marks are replicas of the same word type only if they have the same denotation, recall that the ink spill mark looking like this: 'dog' doesn't denote dogs. Only once we determined that a mark indeed is a word (or other linguistic entity) do we know that it has a semantic value.

appear so radically distinct, that we couldn't even rely on some sort of family resemblance relation among marks of a given linguistic type across distinct symbolic systems.

Maybe there is some *abstract* form that categorizes all the marks into word types, but we don't *see* that form *in* all the tokens. How an abstract form or pattern is supposed to single out its tokens by itself is puzzling, at best?[31,32] If what "glues" the tokens into a type is the abstract pattern instantiated in all of them, but we don't notice that instantiation, then we have an explanatory problem. It is not sufficient for the word type 'dog' to merely exist. This type is supposed to facilitate our intra-subjective reasoning and the inter-subjective exchanges of information. Therefore, unless we can somehow come to know of it and reliably recognize its instantiations, the abstract form cannot explain those processes.

1.2 Proposal: all and only tokens of a given word type resemble each other *within a symbolic system*

Couldn't we at least group the tokens of the word 'dog' within each symbolic system in terms of resemblance, such that the following strings:

i) dog

ii) DOG

iii) *Dog*

are all similar to each other relative to the Latin Writing Script? Given that we probably haven't yet exhausted all the possible symbolic systems, we won't be able to create a complete list of the

[31] Ebbs (2006): p. 115.

[32] It is also puzzling how we could, in principle, have knowledge of abstract objects, due to their acausal nature (Wetzel (2000): p.368). At the same time we take for granted that we have knowledge of mathematical entities (such as numbers and circles). Therefore, unless we can explain the relevant ontological difference between abstract mathematical objects and abstract linguistic objects, we cannot reject the latter to do any theoretic work on the basis of epistemic concerns, without rejecting the former.

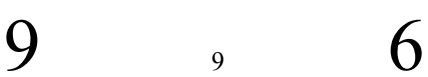
representative categories. We could begin by compiling all the existing ones (Braille, American Sign Language, International Sign Language, Writing System in Latin Alphabet, Spoken Language…). Resemblance would then "merely" have to group the tokens within each symbolic system.

Since everything resembles everything *in some respect or other*, we have to distinguish *degrees* of resemblance. Let's specify a set of features (i.e. shape, color, size, orientation, etc.) on the basis of which to categorize marks. Within a selection of marks, the two that display the most overlap in features are also the most similar to one another.[33]

Consider the following string of marks:



The second, third and fourth marks are all replicas of the first with one feature changed: size, then color and finally orientation. In a context neutral setting, those marks are all equally similar to one another and thus any grouping of the above marks bear claim to forming a class. For example, equal shape and coloration – subset 1:



Or equal shape and size – subset 2:



Or equal shape and orientation – subset 3:

[33] Goodman also considers this suggestion and rejects it (see Goodman (1970): p.26).

9 9 9

If we had asked actual *users* of those marks – people who use those marks as characters – most of them would have picked subset 3 as the one that selects a *relevant* type. Some features we more *habitually* pay attention to than others. Individuating numerals is something we often have to master in order to communicate numbers to one another. If we are on a tight budget, it matters whether a given piece of meat costs $6 or $9 per pound. As it turns out the change in color and size are irrelevant modifications (as long as the result is legible) – given the context of individuating marks as replicas of characters, or numerals more specifically.

Whether a subset groups together marks of an *interesting* type, depends on whether we have a use for the type in our ordinary activities.[34] One may object that this doesn't mean that only those types that matter to human activities should be of importance in *general.* However, if all subsets of the original series form a *type*, then we fail to mark a useful *distinction*. At least in this example, the disinterested observer: nature, doesn't group the tokens into types for us.

A person, a dog and a dust particle are all *extended in space*. Which two members of this set are more alike than the third? You may think that of course, the dog and the human are more like each other as contrasted with a dust particle. The dog and the person are both capable of perception; they are both mammals, which implies they have a heart, can breathe in oxygen through air, the females give birth to live off-spring which are fed with breast-milk… There are, however, important characteristics shared by a person and a dust-particle (which dogs lack): the

[34] Sometimes we classify individual (people) into types, which has significant socio-political consequences. Consider the following monologue from the movie *A Single Man* directed by Tom Ford: "There are all sorts of minorities. Blonds, for example [pauses] people with freckles. The minority is only thought of as one when it constitutes some kind of threat to the majority – a real threat or an imagined one. […] Minorities are just people, people like us […]".

person and the dust particle are alike because they are tailless, furless, not four-legged, unable to bark, unable to track via sense of scent, lacking curiosity for excrements…[35] In terms of taxonomies relevant to human activities a person is clearly more like a dog than a dust particle. In order for resemblance to be a relation capable of drawing contrasts, we must pick *some*, but not all, shared properties as salient. Salience is in the eye of the beholder – an individual or group of individuals. It is *not* an intrinsic feature of the objects under consideration.[36] Which similarities are of importance to our community relies on the kinds of activities we engage in. Picking up on the salient groupings of tokens is *learned* through participation in those activities. When we say that marks $\mathbf{m_1}$ and $\mathbf{m_2}$ are similar to one another, we mean that $\mathbf{m_1}$ and $\mathbf{m_2}$ are similar to each other with regard to *respect R* – the presence of certain properties *we* care about (for one reason or other). This reduces the similarity relation to the presence (or absence) of the property or properties we regard as salient to the type – thereby rendering similarity itself superfluous.[37]

1.3 Proposal: all and only types of a given word type resemble each other *in some salient respect* within a symbolic system

Recalling our example of numeral individuation you may think: that was easy! So marks replicating the numeral '9' all share the same *orientation* and (geometric) *shape*. But what do you mean with 'shape'? You go on to propose a geometric construction for the ideal-numeral-

[35] I failed to come up with features that are *present* in the person and in the dust particle, but not in the dog. Sometimes individual marks are unified by the *lack* of a trait.
Consider the series:



The first and the second are more similar, because they lack the horizontal line extending from the top.
[36] "But how could a pattern or form settle which respects are relevant? The idea of an abstract form or pattern that somehow by itself singles out its tokens is puzzling, at best" (see Ebbs (2009) p. 115).
[37] Goodman (1970): p. 27.

nine-shape: let's first draw a circle and divide it by drawing the North-South[38] diameter through it. Next we "discard" the West[39] segment thereby created (that is, the left arch and the diameter get erased). For our last step we draw a circle with half the radius of the original circle attached to the top end-point of our arch such as to touch two (the maximum amount of) points along it.[40] After that, no rotations allowed! This construction describes the shape that all marks replicating the numeral nine have in common.

Of course actual physical replicas of the numeral '9' *deviate* from the "ideal" geometric shape constructed above. To begin with, geometrical lines have no width, or thickness, while all actual physical replicas that we recognize as the numeral nine must have width, or we couldn't possibly perceive them. How could an ideal numeral-nine construction relate to actual numeral nines in physical space? Clearly it is not the case that *we* have in mind or are somehow *guided* by something like the above construction rule when we write a replica of the numeral nine. Even if we were actually guided by such a homogenous construction rule, we would all have to interpret the rule homogenously. Just looking at how people actually draw the numeral nine, we realize that either we aren't guided by the same rule, or we interpret it differently. Some people start their numeral nine on the top (beginning to draw the small circle first), others on the bottom. Then again, some draw a small circle and attach a *straight* line. That means we are forced once more to look at the end-products – the marks – and categorize them in accordance with some resemblance relationship.

[38] North-South would be the diameter parallel to the y-axis.
[39] West would have to be defined relative to the x-axis.
[40] I am only proposing this as one of several possible ideal mathematical numeral nine-constructions.

As mentioned before, it is puzzling how an abstract form could *somehow by itself* specify which "deviant" physical marks constitute appropriate copies of it – relevant to *our* purposes.[41,42]

1.4 Proposal: *relative to a particular context*, all and only types of a given word type resemble each other in some salient respect within a symbolic system

Let's take a closer look at actual character tokens – do tokens of a given character share any (limited amount of) properties? Consider the series:



Figure 1.1[43]

In terms of intrinsic features, the first two marks are more similar to one another. Nevertheless, it is the first and the second mark that are considered replicas of the same character type. This means (intrinsic) similarity between marks is neither a necessary nor a sufficient condition for them to replicate the same character.

Let's observe the following strings next:

---

[41] Ebbs (2009): pp. 114-120.

[42] Invoking abstract entities to account for the grouping of marks as characters makes us vulnerable to long-standing – but controversial – epistemic objections:

> [T]he truth values of our mathematical assertions depend on facts invoking platonic entities that reside in a realm outside of space-time. There are no causal connections between the entities in the platonic realm and ourselves; how then can we have any knowledge of what is going on in that realm? And perhaps more fundamentally, what could make a particular word like 'two', or particular belief state of our brains, *stand for* or *be about* a particular one of the absolute infinity of objects in that realm?

If we cannot explain our knowledge of geometric shapes, then we cannot explain our knowledge of characters in terms of geometric shapes. This leaves us at a loss elucidating why all replicas of the numeral nine share the same *shape* and how this shared shape facilitates communication and reasoning. Field (1980): p.1.

[43] This example is taken from Goodman (1976): p. 138; the order of the letters has been changed.

ISABELLE

language

I0 + 20 = 30

Figure 1.2

The first mark of the above three strings of marks are identical (except for their location in space and their environment). In the first string, the first mark is a replica of the capital letter 'I'. In the second string that same mark is a replica of a different letter, the small letter 'l'. In the last sequence the same mark doesn't stand for a letter at all, but is a replica of a numeral. The identification of the mark as a replica of a particular character depends on the marks that follow it (and what characters those marks replicate) and its location in the sequence it is a part of. The three marks in question are ambiguous. Although indistinguishable in terms of intrinsic features, they all replicate *different* character types. Notice how the marks in question could all pass as "standard"[44] replicas of distinct characters. Even physical identity (that is, constituting an exact copy) is not a sufficient condition for character-hood. If we had encountered the first marks of the three strings of figure 1.2 above in isolation, we would have been at a loss of how to individuate them in terms of characters. Imagine a person draws '**I**' on a piece of paper. Only the producer of that mark would know what character she intended to replicate, or whether she perhaps just drew a line with no character in mind (unless she sincerely shares that information

[44] A standard replica is one with a high degree of recognizability by competent users of the character in question. We cannot define 'standard' as very similar to the ideal pattern or form, since we have yet to find a working definition of 'similar to' and we don't know how an abstract pattern could be similar to a physical token.

with us). According to Kaplan we cannot individuate marks as replicas of characters without taking into account the intentions of the producer.[45]



Figure 1.3[46]

In figure 1.3 the first and the last mark were written by different authors – authors A and B respectively. Both marks are replicas of the numeral '4'. If the middle digit was written by B, then it is clearly a '9'. If the same mark was produced by A, then it could also be a '4', albeit a somewhat botched '4'. The individuation of the mark as a replica of a given numeral sometimes depends on the style of handwriting. Figure 1.3 shows furthermore that some marks may deviate from the standard way of writing a character. The middle mark may be a standard replica of the numeral '9', but it is not a standard replica of the numeral '4' (regardless of authorship).

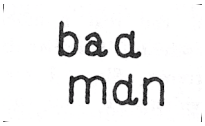


Figure 1.4[47]

In figure 1.4, the last letter of the first word and the second letter of the second word constitute identical marks. Nevertheless, each mark is a replica of a different letter type. Letter individuation is dependent on the progression of the string of letters ('bad' and 'man' are words

[45] Kaplan (1990); see my chapter three on further discussion of Kaplan's notion of words.
[46] This example and the accompanying figure are taken from Sas (2006).
[47] This example is taken from Goodman (1976): p. 138.

of our vocabulary, while ‘baa’ and ‘mdn’ are not meaningful letter sequences in English[48]), as well as the style of handwriting or choice of font. Which of the following two styles:

i) **a**

ii) **a**

is the small letter ‘a’ written more like? In contrast to the first example, the mark in question is not a standard replica of the letter ‘**a**’ or ‘**d**’. The confusion arises partly, because (even assuming a certain style or font) the mark falls in-between what would count as a “standard” replica of an ‘a’ or ‘d’. It is the *context* that settles (partly) whether the borderline instance is a replica of the letter ‘a’ or ‘d’.

1.5 Added hurdle: accounting for replicas constituting mistakes

Not only do we have to account for one and the same mark replicating different characters, depending on context, we also have to allow for the possibility that a mark is considered in-between the standard replica of two different characters. In addition, some tokens are misspelled, mispronounced or otherwise mis-“shaped” versions of the linguistic type. If I accidentally hit the ‘D’ key, instead of hitting the ‘A’ key when typing the word ‘man’, such as to produce: ‘mdn’, then the second letter of that word doesn’t replicate a highly deviant ‘a’; instead, I misspelled the word ‘man’. It is only by considering ‘mdn’ a token of the word ‘man’ that we can account for its mistaken spelling – otherwise it constitutes just a meaningless string of letters.

[48] Of course it is possible that the second mark in figure 1.4 is simply a misspelled version of ‘man’. Perhaps the writer got distracted and was thinking of a loved one whose name starts with a ‘D’. The mistake may even go unnoticed. But if you showed her the result, she would say: ‘Oh no, I spelled ‘man’ with a ‘d’ in the middle!’. She wouldn’t say: ‘This is how I write my ‘a’ letters’. I think those retrospective judgments are of philosophical importance. I will discuss this in a later chapter.

Foreigners often pronounce words with the wrong inflection. I used to consistently pronounce 'reciprocal' with an accent on the 'o', instead of the 'i'. The correct pronunciation simply wouldn't stick. This "insignificant" deviation from the norm would confuse my interlocutors (unless they were close friends already used to my quirky pronunciations) to such a degree that they were unable to recognize my phonetic production as a replica of 'reciprocal'. I think it accords with our pre-theoretic intuitions that I still managed to use the English word 'reciprocal', even though competent speakers didn't recognize it as such. Kaplan agrees:

> I would claim that a *mis*pronunciation of a word is an utterance (a *Pronunciation*, if you will) of *that* word. […] It's just that the subject can't come anywhere close to giving the word a standard pronunciation.[49]

There is a surprising lack of homogeneity in (acceptable) word-pronunciation even within an individual's idiolect. In some contexts, perhaps while reading out loud or speaking to somebody who is less familiar with the English language, I may carefully enunciate all six syllables of the word 'extraordinary'. At other times, when speaking very quickly or perhaps when slightly tired, I will shorten it to merely four syllables.[50] Some English words have more than one correct spelling such as: 'medieval', which has two alternative spellings: '*mediaeval*' and '*medieval*'.[51] As long as we could reduce each pronunciation or spelling to a very limited list of options – that is, one or a few more strings of phonemes or letters – these variations shouldn't constitute a problem. However, this means we theoretically depend on phoneme or letter types. We have yet to offer an account of letter types; phoneme types are even harder to account for.[52] As if matters were not already complex enough, what is considered *correct* spelling may change depending on the medium. In a mobile phone text-message, the pronoun 'you', for instance, is

[49] Kaplan (1990): p.105.
[50] Wetzel (2006).
[51] This example is taken from Janssen & Visser (2004).
[52] Wetzel (2009): pp.65-68.

often written: '*u*' – although this would be considered a mistake in most other written media. Similar variations can be found with respect to phoneme types which are sensitive to class- and regional-dialects and need not even be constant throughout an idiolect.[53]

1.6 Final Proposal: relative to a particular context, all and only types of a given word type resemble each other in some salient respect within a symbolic system *as recognized by a computer program*

Although intrinsic resemblance between two marks is neither a sufficient nor a necessary condition for those marks to replicate the same character type, surely resemblance – even if it is the kind we habitually pay attention to (resemblance in a particular respect) – must *guide* us somehow in the individuation process. If each replica of a given letter were *radically* different from its next, how could we ever individuate words? Communication would be unimaginable. The individuation of marks as replicas of certain types depends on the general purpose of the activity in which the mark was produced, as well as more narrowly construed contextual factors within linguistic reproduction: such as "style" of characters and the individuation of the surrounding characters and the placement of the mark within a meaningful string of symbols. Certainly the naïve snapshot approach failed: even when it comes to the identification of marks as replicas of characters, the snapshot that captures the intrinsic features of the mark is not sufficient (nor necessary) information. We furthermore need *extrinsic* information that sets the mark within a *meaningful context* of a given practice. Not all tokens of a given character type share the *same* shape, but they display family resemblances. Can we render those family

[53] Wetzel (2006).

resemblance features transparent with the help of mathematical formulas and thereby program computers to do the individuating work for us?

Computer science is individually tackling those issues with programs such as: spell-checker, predictive text generation, optical character recognition and speech recognition – just to name a few.[54] Although optical character recognition for a typewritten text is achieved with very high success rates, there is presently no computer-automated method that recognizes handwritten characters with acceptably low error rates[55] (unless the dictionary is very restricted, as for instance, when individuating marks entered on the amount line of personal checks). Computer scientists are challenged by (i) regional differences in conventions of how to write or draw letters using the same alphabet or graphemes and (ii) variations across individuals (see figure 1.3). In order to overcome regional and author-specific variations some methods identify the author first by relying on features such as: (average) mark width, (average) mark height (distinguishing between marks replicating small or capital letters), relation between width to height, average distance between individual marks within a string of marks, distance between individual strings of marks, average writing slant and angle…[56] Significantly higher recognition rates are achieved by employing author recognition techniques.[57]

Spell-checker programs (as used in Microsoft *Word*) are able to identify word tokens as misspelled word types and suggest a correction. Such programs compare word tokens against a dictionary of correctly spelled words. For example if I mistype 'word' as 'wor', my spell-checker will offer the following options: 'work', 'word', 'worn', worm', 'wore', 'war', 'woe',

[54] See Kukich (1992); Golding, Golding & Roth (1999); Sas (2006).
[55] "First, the richness of the linguistic structures that must be represented results in extremely high-dimensional features spaces for the problems. Second, any given target concept depends on only a small subset of the features, *leaving a huge balance of features that are irrelevant to that particular concept* [emphasis added]" (Golding & Roth (1999)).
[56] Ibid.
[57] Ibid.

'worry'. Those options are ranked in terms of (i) how commonly the word type occurs in the language generally, or on the basis of frequency of use in documents on the respective personal computer or in the particular text, as well as by (ii) the degree of (letter) deviation between the misspelled sequence and the proposed one and (iii) the distance of the corrected letter(s) in the proposed word on the keyboard from the ones that were erroneously typed. There are two major problems with this approach: (a) if the misspelled string matches the correct spelling of a *different* word, the spell-checker won't be able to identify the mistake[58] and (b) in many cases the program is not able to link the string to only *one* word, it merely matches it to a small selection of distinct words (sometimes not even words of the same grammatical type).

In order to avoid the aforementioned problems more sophisticated models that allow for context-sensitive spelling correction are being developed.[59,60] The context features tested are word tokens near the target word (that is, the word type with correct letter type order), as well as patterns of parts of speech around the target word. In addition to a dictionary, the computer is provided with confusion sets containing homophones {whether, weather}, {hear, here}, {there, their}…, and typographical errors {form, from}, {desert, dessert}, {maybe, may be}.[61] Examples of useful features for the first confusion set {whether, weather} include:

(1) 'cloudy' within ±10 words

(2) ____ 'to' followed by a verb

If (1) is encountered, then it is highly likely that 'weather' was intended. If (2) occurs, then 'whether' is probably the target word. Of course, we will have to deal with tricky cases such as: 'I don't know *whether* cloudy *weather* is to be expected'. Major improvements are being

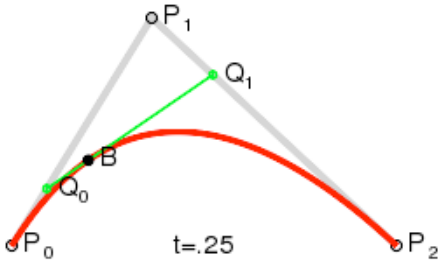
[58] Such errors account for anywhere from 25 to over 50% of observed spelling errors (Kukich (1992)).
[59] Golding & Roth (1999).
[60] Microsoft Word 2007, like several other text-writing programs, already makes use of such a context sensitive spelling program.
[61] Golding & Roth (1999).

achieved in word individuation for spell-check purposes. The question is ultimately whether we can identify a limited number of features that can, for each context and for all confusion sets, select the target word. The more features and confusion sets we add in order to increase completeness (the distinction of all words in all "actual" confusion sets) the more computationally challenging becomes the task. Since increasing the amount of features to test for, also adds irrelevant features for the distinction of the words of a given confusion set in a particular sentence[62], I wonder whether completeness will eventually interfere with soundness (the identification of the target word). Notice, however, that we already perform some of the classification work for the computer, by hitting specific keys on the keyboard (or by selecting from the symbol menu). Whenever I want to produce a token of the letter 'd' I have to press the 'D' key. The production of character tokens through a computer is standardized.

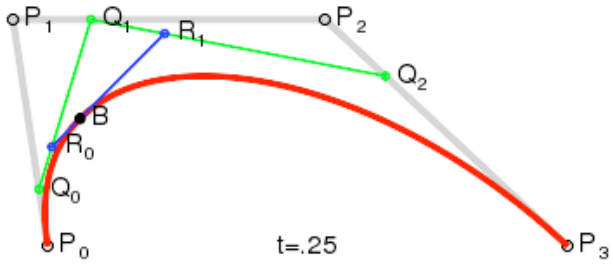
Will we be able to combine programs that correct spelling mistakes (such as spell-checker or predictive word-generating) with optical character recognition for handwritten text, or speech recognition such as to avoid having to input marks produced by humans and thus first make those marks recognizable to the computer?

Computer generated characters of a particular font use Bézier curves to model smooth curves that can be scaled indefinitely. It should be theoretically possible to generate infinitely different deviations of the "geometric" numeral nine construction described above with the help of Bézier curves – altering the circles into various curved lines. This would account for all "ideal" numeral nine shapes.[63]

[62] See footnote 48.
[63] This constitutes what Wetzel calls a disjunctive shape theory (Weztel (2009): p.64-65). We may find something that all the numeral nine token construction Bézier functions have in common.



Construction of a quadratic Bézier Curve [64]



Construction of a cubic Bézier curve[65]

Human handwritten characters are never perfectly *smooth* (many points will deviate from the trajectory of the curved or straight lines attempted), sometimes they are partially *incomplete* (see first and second mark in figure 1.3; see third letter of the first word in figure 1.4), or contain unintended *extensions* (see second letter of the second word in figure 1.4). These are probably all deviations that can be eliminated through statistical functions. Computers can be programmed to find the center of a human pupil, which are only *approximately* circular.[66] This means overcoming (primarily) the following problems: (a) pupils never have a perfectly smooth circumference, (b) some pupils have "dents" interrupting the circle and (c) sometimes part of the pupil is hidden by the lower eyelid and the eye-lashes[67] – all problems we also encounter when individuating characters. If such "correction" operations can be performed on an approximate circle, it should in principle be possible to do the same for marks approximating ideal geometric character shapes. Next, the computer would check which of the Bézier defined shapes the mark matches best statistically.

This is only a sketch of how we may program computers to recognize handwritten marks as character types. However, we cannot claim that the set of all the tokens of the letter 'd' have nothing in common that distinguishes them from tokens of other letter types, while at the same

[64] See URL: http://en.wikipedia.org/wiki/Bézier_curve
[65] See URL: http://en.wikipedia.org/wiki/Bézier_curve
[66] Moore & Haslwanter (1996).
[67] This is relatively rare, but happens when people have droopy eyelids, or the picture of the pupil is taken when it is dilated.

time computers can be programmed to reliably[68] decipher hand-written text. It remains to be seen whether the completeness of the program will interfere with the soundness of its results.

1.7 Insurmountable hurdle: some tokens remain unrecognizable

There are two basic problems that no computer will be able to overcome. If a new font (significantly distinct from the presently existing fonts) were to be invented[69], the computer program may or may not catch it, depending on whether its statistical "correction" calculation can match it to one of the existing fonts or letter shapes. Furthermore, an entirely new symbolic system could be invented. Pupils in school may decide to communicate information to each other via a new code based on closing and opening of the eyelids. Perhaps this code catches on and finds good use in hospitals with paralyzed patients.[70] If this code specified a sequence of eye blinks as equivalent to the word 'dog', then this would need to be recognized. Again, there is no way to account for this development at present. Whether with or without the help of a computer, we cannot identify what word type a mark produced in a future symbolic system replicates. That mark may already exist today (people do blink their eyes), but we simply couldn't know that this mark will come to mean 'dog'. However, if we want to reduce a word type to a set of physical marks (relative to symbolic system and context), then this set should contain all past, present and *future* marks. Otherwise, the word type 'dog' would change constantly with the production of new marks.[71]

---

[68] 'Reliably' here means that the program must do as well as humans or better.
[69] Wetzel suggests the invention of a new font as a problem for "disjunctive shape" theories (Wetzel (2009): p.65).
[70] I believe that the present technique only allows patients to say 'yes' and 'no'.
[71] Ebbs (2009): pp. 118-119.

A computer furthermore lacks access to speaker intentions.[72] A person may produce the following mark on an otherwise unmarked piece of paper:

The manufacturer of that mark, may have intended to write 'Idea', but didn't get to complete that word. This would mean that the above mark replicates the capital 'I'. No computer program, no matter how sophisticated, may tell us whether that person intended to write the capital letter 'I', the small letter 'l', the numeral '1' or whether that mark replicates no character at all – it's just a line of specific width and length. However, we wouldn't do much better than the computer (unless that person informed us of her intentions).

1.8 Conclusion

As demonstrated in this chapter we weren't able to *reduce* homonymic word types to sets of physical word tokens (marks). Even if my imagined computer program worked, it would have to be based on mathematical formulas, which at the very least manipulates symbols, that is, character types, not tokens![73] We may have character individuation reduced to manipulation of a few types: a quantifier, negation, conjunction, zero and one – thereby elucidating linguistic types, even if we didn't reduce them to physical marks. Goodman foresaw our failure: "The

[72] Unless computers science were to find out that the intention to produce a certain character is causally linked a measurable brain-state.

[73] Wetzel argues: "Formalism is, roughly, the philosophy of mathematics that holds that math is not "about" anything, least of all numbers, sets, and spaces; it is the mere manipulation of symbols […]. […] The problem is that these "concrete signs" cannot be construed as physical tokens if Hilbert [one of its main proponents] is to derive the mathematics he wants" (Wetzel (2009): p.56).

letter-classes of our alphabet […] are established by tradition and habit; […] defining them would be as hard as defining such ordinary terms as 'desk' and 'table'."[74]

We weren't able to define the homonymity relation beyond a *perceived* one.[75] Of course perceiving that tokens *t* and *t'* are of the same physical word type does not make it so. In order to rigorously define homogeneity, we would need an account for phoneme and letter *types* – which we do *not* have.

Given the complexity of factors to be taken into account for correctly individuating words, the ease and success rate with which we do so is surprising! In most cases we know (without reflection) what character or word a given mark replicates – even when the mark was produced by other language users. Even though sometimes only the producer of a given mark recognizes what character or strings thereof the mark replicates.

Starting from the fact that people sometimes use the "same" words to talk about a given topic, I want to clarify what word-sameness comes to in those uses. My first attempt towards explicating the notion of *words* operative in our communication of knowledge began by asserting that words are types instantiated in utterances as tokens. Ebbs explains how most theories of meaning subscribe to a two-step word individuation process:

(U) Two word tokens *t* and *t'* are of the same word type if and only if
  (i) *t* and *t'* are each tokens of the same orthographic or phonetic type, and
  (ii) facts about the explanatory-use of *t* and *t'* determine that they each have semantic values and that their semantic values are the same.[76]

Our failure to account for step (i) means that we will have to individuate words entirely through step (ii) or switch to a *different* framework altogether.

---

[74] See Goodman (1976): p.138.
[75] It is quite possible that cognitive science may discover what brain state we are in as perceivers when we pick up on homonymity. Even if that were possible, it still wouldn't make the relation one that obtains between words.
[76] See Ebbs (2009): p.112; Ebbs rejects the token and explanatory-use model of words (U).

CHAPTER 2

UNDER WHAT CONDITIONS ARE TWO UTTERANCES PERFORMANCES OF THE SAME WORD? PROPOSAL: RESEMBLANCE OF COGNITIVE CONCEPTIONS

2.0 Introduction

In order to individuate words for the purpose of drawing valid inferences and reliably including information based on other speaker's testimonies, we must be able to individuate our own words and those of others across time and varying contexts. Words are types instantiated in utterances as tokens. Ebbs explains how most theories of meaning subscribe to a two-step word individuation process:

> (U) Two word tokens *t* and *t'* are of the same word type if and only if
> (iii) *t* and *t'* are each tokens of the same orthographic or phonetic type, and
> (iv) facts about the explanatory-use[77] of *t* and *t'* determine that they each have semantic values and that their semantic values are the same.[78]

Assuming that our goal is to individuate words in a way relevant to knowledge transmission (i.e. through testimony), we came to realize in the previous chapter that it is not some resemblance relation between physical tokens or marks that groups the tokens into word-types.[79] In this chapter I will show that the resulting collapse of (i) and (ii) begets insurmountable challenges for semantic individualism.

[77] Facts about explanatory-use are the necessary and sufficient conditions that supposedly determine the meaning or satisfaction conditions of a given token (Ebbs (2009): p.9).
[78] Ebbs (2009): p.112; Ebbs rejects the token and ex-use model of words (U).
[79] There are notions of *word* that are independent of semantic value – those relevant to copy editing, for instance. Those notions, however, do not fulfill my explanatory goals.

2.1 Individualism defined

For the scope of this chapter I define 'individualism' as the thesis that any word-token *t* is individuated *independently* of correlations with the speaker's linguistic community.[80, 81] The kind of individualism I have in mind conceives of the semantic value of a word token *t* in terms of the *cognitive content* the individual speaker associates with *what individual or group of objects t denot*es (in conjunction with external physical facts about the speaker's environment)[82].

The comparison of individual conceptions is inherently problematic. We clearly don't have access to the totality of conceptions other speakers associate with their tokens 'dog'. I am not even sure whether I have first-person access to all of my own dog-related conceptions. I certainly cannot come up with a definition that captures all the *essential* features I believe to constitute dog-hood. It seems that we must think of these cluster of beliefs related to a given word as *idealized* aggregates of beliefs:

> [T]heories of meaning for terms such as "clutch" and "water" in particular agents' language [are] idealized constructions, i.e. as ideally complete specifications of their "concepts," which hypothesize on the basis of ideally complete evidence […] all the beliefs that an agent associates with each such term in his language. Of course, we cannot make such idealized specifications

[80] Individualism, so understood, must be distinguished from the stronger thesis of *internalism*: any word-token *t* is individuated independently of the world external to the individual who produced *t*. For individualists it is permissible to take into account information about the individual's physical environment (except for brain states of *other* individuals). Sometimes individualism is understood to make a stronger claim, namely, that content (or what utterances express) is determined *only* by factors internal to individual agents. Bilgrami labels this stronger version of individualism 'internalism' and views himself as a proponent of individualism, but not internalism (Bilgrami (1992): p.65).

[81] This kind of individualism is primarily defined in contrast to *social* externalism, which Burge defines as follows: "The main idea of [… anti-individualism] is that the natures and correct individuation of many of an individual person's intentional, or representational, mental states [,words] and events commonly depend in a constitutive way on relations that the individual bears to a wider social environment" Burge (2007): p. 151.

[82] Bilgrami's externalist constraints enter into the individuation process as follows: externally determined concepts of an agent are fixed by consulting the respective *indexically formulated utterances* in conjunction with relevant beliefs within her complex network of beliefs (Bilgrami (1992): p.5).

> explicitly because we usually cannot list all the beliefs an agent associates with these terms of his language.[83]

It is somewhat nebulous what these idealized constructions amount to. Of course, a speaker can draw on memories of experiences involving other speakers (especially for complex terms such as 'discrimination'). It is the *lack of linguistic deference* to other speakers that renders the conceptions individualistic. For example, we could not include intentions such as: I am repeating my neighbors word 'German Shorthaired Pointer'.[84] Assuming that we have some way of determining individual conceptions, let's proceed.

Given the primary role contents of belief are to play in psychological explanations, accounting for individual variations in understanding is a primary motivation for the individualist account:

> [A]n individual's representational content depends (partly) on a web of inferential connections with other representational contents. The idea is that the constitutive conditions for understanding a concept cannot outrun the network of inferences that the individual can draw.[85,86]

In order to make sense of the inferential patterns of the speaker individualism attempts to derive word-type identity from individual conceptions.[87,88,89]

Weigand stresses that understanding is a *process* of *coming* to an understanding[90] – she is thereby suggesting that understanding is not a binary phenomenon, but one that admits of

---

[83] Bilgrami (1992): p. 12.
[84] Perhaps I heard that she is intending to adopt a dog of that breed and am conveying this to some other neighbor using that same expression – even though I am not quite certain what constitutes a member of that particular breed (I merely know that they are hunting dogs).
[85] Burge (2007): p. 176.
[86] Begby (forthcoming) describes the motivation thus: "the range of thoughts that a person can entertain by reference to the concepts at that person's cognitive disposal".
[87] A speaker who associates one set of beliefs with Paderewski, the politician, and another set of beliefs with Paderewski, the pianist, may very well believe that 'has musical talent' can be predicated of one Paderewski, but not the other (even though there is actually only one Paderewski). Such an individual is not able to see that 'Paderewski has musical talent and it is not the case that Paderewski has musical talent' constitutes a logical contradiction. See chapter five for further discussion.
[88] "One has failed to capture the inferential patterns which give rise to an agent's behavior. And this [malaise] is diagnosed as flowing from a general source: an agent's conceptions are being ignored in the attribution of externalist content" Bilgrami (1992): p.19.
[89] See chapters five and six for further discussion.

gradations[91]. Differences in understanding, she argues, result from the interlocutors inhabiting *different* cognitive backgrounds.[92] Utterer and interpreter inhabit different physical, social and mental worlds, such that common ground is almost never really common.[93] Comprehension is the *creation* of meaning resulting from a *gradational overlapping of meaning*, individual cognitive contents that is.[94] Pagin proposes in a similar vein:

> A communicative event (between speakers that have propositional attitudes) is *successful* just if the content of the hearer's thought is sufficiently *similar* [to] the content of the speaker's thought. Sufficient similarity as here intended is an equivalen[ce] relation[…].[95]

According to individualism no two people are likely to have the *same* aggregate of conceptions associated with the tokens of a given word-type, thus making it difficult to explain how we could share the same word-type intra-subjectively.[96] Consequently, a similarity relationship between these aggregate conceptions will have to confirm the appropriate groupings of tokens produced by different speakers into word-types. Your word-token 'dog' and my word-token 'dog' constitute tokens of the same word-type because we both share relevantly similar conceptions with the respective token we produced. I will argue in this chapter that there are too many other tokens with which some speakers associate *equally similar* conceptions (i.e. 'hound', 'dingo' and 'gray wolf')[97]. We would need a method to filter out these tokens *without* relying on community norms that provide us with adequate similarity standards grouping physical forms into the relevant types. My arguments of the previous chapter deprive individualists of a way to exclude

---

[90] Weigand (1999).
[91] The frequency of *incomplete* understanding plays an important role in Burge's anti-individualism (Burge (1979, 2007).
[92] Weigand (2000).
[93] Verschueren (1999).
[94] Ferbezar and Stabej (2008).
[95] Pagin (2006): p. .
[96] Bilgrami (1992: p. 10), himself a proponent of this view, admits this to be true of the individualist notion of concepts.
[97] For those wondering why the conceptions associated with 'dog' could be similar with 'gray wolf', let's not forget that people can be under-, as well as, mis-informed about a topic.

these competing candidate tokens on the basis of their physical features. Our failure to account for word-*forms* begets insurmountable challenges for individualist conceptions of words.

2.2 *Intra*-subjective word individuation

Does the relationship between the (totality of) *conceptions* a speaker associates with what her tokens *t* and *t'* denote determine whether *t* and *t'* constitute performances of the same word? We would have to show that whenever the respective conceptions *overlap* in the appropriate manner, *t* and *t'* constitute performances of the same word (and *vice versa*).

Imagine a physics student uttering u_1:

"*x* satisfies '<u>space</u>' if and only if *x* is <u>space</u>."

These are the kinds of things we can also hear her say: "Space is Euclidean, that is, two parallel lines can be extended indefinitely and will never intersect." Some time goes by and she now sincerely affirms: "Strictly speaking, space is non-Euclidean …" At some later time that student utters u_2:

"*x* satisfies '<u>space</u>' if and only if *x* is <u>space</u>."

We would say that the student learned[98] something that *radically* altered her conception of space, but the subject kept track of the same topic throughout her learning process: "Participants [of a discussion]", according to Burge, "commonly regard their object-level thoughts (thoughts about, say, chairs) as undergoing correction in the course of the inquiry."[99]

Given that the conceptions associated with what 'space' denotes in u_1, *radically* diverge from those associated with what 'space' denotes in u_2, it would seem that according to anti-

[98] Explaining how we can *learn* from others is one of the central motivation for anti-individualism. See Ebbs (2002).
[99] Burge (1986, 2007): p.261.

individualism the respective tokens would have to constitute performances of different word (types). Consequently, her utterances u_1 and u_2 are not stating the same thought[100] – which is counterintuitive.

At other times, speakers can be under- or mis-informed about a given topic. Putnam confesses not being able to distinguish an elm from a beech tree.[101] He may utter u_3:

"*x* satisfies '<u>elm</u>' if and only if *x* is an <u>elm</u>."

as well as u_4:

"*x* satisfies '<u>beech</u>' if and only if *x* is a <u>beech</u>."

Since Putnam's conceptions associated with what his tokens 'elm' and 'beech' denote respectively are identical, it follows (from individualist assumptions) that u_3 and u_4 express the same (thought) – whereas Putnam would say that his utterances u_3 and u_4 express different *satisfaction* conditions. It is tempting to argue that, since the tokens in question do not instantiate the same shape or form, the question of word-sameness doesn't even arise. This assumption would significantly simplify the task, since we would only have to test tokens of the same form for word-sameness. I argued in the previous chapter that we were unable to group tokens in terms of resemblance relations between physical tokens or marks. As a result, this simplification is not available to us.

Individualism leads to counterintuitive results when it comes to grouping tokens intra-subjectively. What I have not yet considered is expanding the theory with an intentional-clause: if a speaker *S* intends to replicate a token *t* (which *S previously produced*) and thus produces

[100] With 'thought' I mean the content of what an utterance expresses.
[101] Putnam (1975): pp. 226-227.

token *t'*, then *t* and *t'* are of the same word type.[102] It seems to me that those intentions would have to be part of the resources available to an individualist, since they are (so to speak) "inside the individual's head." Consider my uttering u_5:

"Syrus is a dog."

And while pointing at the same individual, I utter u_6:

"Syrus is a dog."

We don't want to rely on the somewhat "flimsy" assumption that as long as the last token of the utterance u_5 and the last token of the utterance u_6 were produced within a relatively short timeframe, the corresponding conceptions associated with what each token denotes are likely to be similar. Even if between uttering u_5 and u_6, I learn some substantial new facts about dogs (i.e. that dogs are domesticated animals) we would want u_5 and u_6 to express the *same belief*. We can account for this identity by arguing that I intended to replicate the last token of the first utterance with the last token of the second utterance.

In my example, Putnam did *not* intend for 'elm' to replicate 'beech' (or *vice versa)*, thus utterances u_4 and u_3 express *different* (thoughts). Individuals are supposed to have privileged ("full") knowledge of their own intentional states.[103] If a speaker's utterances express her thoughts, then she must be able to group her tokens into word types. Without assuming this ability, we cannot explain how individuals are able to draw valid inferences from their own utterances. I only grant individualists to draw on intentions with regard to *intra*-subjective sameness of word types. The same cannot be done for *inter*-subjective sameness of word types,

[102] Kaplan applies this intentional proposal also to inter-subjective tokens, which I will discuss in the chapter three.
[103] I will explore this issue in chapters five and six.

since this would introduce semantic deference to *other* speakers – a violation of individualist principles!

My notion of *thoughts* is intricately linked with how I conceive of words. The latter are instruments that render the inter-subjective transfer of knowledge possible. Thoughts are what utterances express. They are what speakers attempt to communicate through the use of words.

2.3 *Inter*-subjective word-individuation

Do individual conceptions play an important role in inter-subjective word individuation? In the previous section we concluded that individual conceptions don't determine whether two tokens *t* and *t'* are tokens of the same word type. Perhaps resemblance of individual conceptions can explain under what conditions we understand utterances of other speakers.

Sometimes individual conceptions are understood in terms of paradigms.[104] Suppose James's paradigm "dog" is a German Shepherd (a guard dog), while Sarah's paradigm "dog" is a Chinese Crested (a toy dog).[105] Let's go so far as to say that the cognitive concept that James associates with his word 'dog' is *importantly different* from the cognitive concept Sarah associates with her word 'dog' – there is little overlap between James's conception *dog* and Sarah's conception *dog*. The Chinese Crested is hairless, weighs about 8 pounds on average and doesn't bark. The German Shepherd, in contrast, weighs about 85 pounds on average and, as most of us know, is covered with brown, yellow, black or white fur and likes to bark (hence the frequent use as a

[104] Rosch, Eleanor (1975); Bilgrami, Akeel (1992): p. 25 – just to give some examples.

[105] The other day at the dog park I made an interesting observation of how the paradigms that we associate with a kind can influence how we (sometimes falsely) individuate that kind. A family with two kids of ages about three and five walked in with their Great Dane. The kids were unfazed by the rough treatment of their Dane, even when they accidentally got pushed over. But when some smaller dogs came to participate in the chasing games, they ran away screaming. Their father had to explain: "It's okay, those are dogs too". It seems that the children couldn't fathom that animals so much smaller than their own dog could be dogs also.

guard-dog). Under what conditions are Sarah and James talking about the same topic when they use their respective words 'dog_S' and 'dog_J'? The problem for the proponent of individualism seems to be that they *cannot* explain how Sarah and James ever could be using the same word, if indeed the individual conceptions associated with their respective utterances are *significantly* different.[106]

Of course there are examples where speakers clearly talk past one another. Putnam invites us to imagine somebody, let's call him Oscar, uttering: "This is a tiger!"[107] – while pointing at a snowball. Jane, a friend of Oscar's, would describe tigers as follows: "a tiger is a yellow and black striped feline…" Putnam declares that there wouldn't be much use in Jane's talking to Oscar about tigers (using her word 'tiger'). In this case, individualism produces the intuitive result! The conceptions Jane associates with what 'tiger' denotes are *radically different* from those Oscar associates with what his 'tiger' tokens denote. We need some theory that distinguishes cases like this tiger thought-example from the case above concerning the words 'dog_J' and 'dog_S'. I don't think ranking conceptions in terms of degrees of discrepancy will turn out to be of much explanatory value – given the possibility of words denoting conceptually similar things (such as the American Robin and European Robin).

Let's recall Putnam, who confesses not being able to distinguish an elm from a beech tree.[108] Let's assume that since this coincidence suits his theory of language, he makes a point of never consulting an arborist or encyclopedia to fill in this knowledge-gap**.** That is, the conception

---

[106] Loar objects to paradigms as picking out referents for similar reasons: "A non-social causal theory requires non-socially-mediated causal-reference relations. I doubt that we have conceptions of such relations […] consider the simple fact that there is no such thing as *the* natural kind to which paradigms of a predicate like "dog" belong: they are dogs, members of the family Canidae, mammals. A person's term "dog" does not acquire a natural kind as its reference merely by pointing to paradigms" Loar (1985): pp. 128-129.

[107] We are assuming that this snowball doesn't even have the "shape" of a tiger.

[108] Putnam (1975): pp. 226-227.

Putnam associates with his tokens replicating 'elm' across time remain unchanged. I earlier imagined Putnam uttering u_3:

"*x* satisfies '<u>elm</u>' if and only if *x* is an <u>elm</u>."

as well as u_4:

"*x* satisfies '<u>beech</u>' if and only if *x* is a <u>beech</u>."

We solved this problem by arguing that Putnam didn't intend to replicate the last token of the first utterance u_5 with the last token of the second utterance u_6. This solution only gets us so far, since an analogous problem can occur with *inter*-subjective word individuations. As clarified earlier, intentions to repeat *another* speaker's token introduce deference conditions that are not available to individualists without violating their own theory.

Suppose Quine knew exactly how to determine whether a given tree is a beech or an elm. When he utters u_7:

"*x* satisfies '<u>elm</u>' if and only if *x* is an <u>elm</u>."

as well as u_8:

"*x* satisfies '<u>beech</u>' if and only if *x* is a <u>beech</u>."

those utterances u_7 and u_8 lay equal claim to expressing the same (thought) as Putnam's utterance u_4 within the individualist framework. This conclusion is counterintuitive.

Putnam's conceptions associated with 'elm' are equally similar to those Quine associates with tokens *of* 'elm' and 'beech' respectively; yet it is with Quine's token 'elm' that we want to group Putnam's 'elm' – into a unified word-type. Again, what we would like to say is that Putnam's token 'elm' is of the same orthographic or phonetic type as Quine's token 'elm' (and *not* of the same orthographic or phonetic type as Quine's 'beech') – but we don't have the necessary resources to do so. Even though such a shared "form" would *not* suffice to explain why u_3 and

u_7 express the same thought it would at least *narrow down* the relevant set of tokens against which to compare Putnam's tokens to considerably.

If all we have to go by are the physical characteristics of the respective individual tokens, substantial similarity may be completely lacking. Suppose Putnam knows how to spell the word 'elm' (in Latin Script), but is clueless about how to pronounce it (which of course is not actually true). Quine, on the other hand, only ever utters the word 'elm', but has no idea of how to put it in writing (which is of course also counterfactual). We could take this one step further and ask a competent "hearer", reader, speaker and writer of the English language to spell Quine's utterance and try to "match" it with Putnam's orthographic production. Clearly, this can go wrong on either end. The "translator" may be unable to understand Quine's utterance, or read Putnam's handwriting. At this point it truly seems that we have run into an impasse, because we won't receive any useful information from the comparison of the physical features of the tokens under consideration.[109]

It is so tempting to argue that the question of whether Quine's word 'beech' and Putnam's word 'elm' are one and the same word (in the sense of *word-sameness* that allows for the transfer of knowledge) *doesn't even arise*, since they don't even instantiate the same word *form*. As argued previously, this line of reasoning is not available to the individualist.

This time intentions – as they did for *intra*-subjective token individuations – can't come to the rescue, since this would render individual conceptions *entirely* superfluous from an explanatory point of view. Assume speaker S_1 produces a linguistic token *t* and speaker S_2 produces *t'* with the intention to repeat *t*. It is possible that in some circumstances S_2 merely intends to repeat the word-form. For instance, when children are first introduced to cursive handwriting, they are

[109] See Chapter 1 for further discussion.

slowly trained to adhere to the similarity standards of their respective community relevant to reproducing a written copy. The focus is at first on how the letters and words *look*, not what they mean. However, in conversations where the interlocutors are using the respective tokens in order to discussing a single topic, S_2 wouldn't merely commit to repeating the word-form, but also to what that word expresses (in her linguistic community). Imagine Putnam were to ask Quine: "You keep mentioning elms and beeches. I am not sure how to distinguish them from each other. Can you tell me more about those trees?"[110]

In conclusion, resemblance among the relevant individual conceptions doesn't group inter-subjective tokens into word types and will thus not explain under what conditions we are able to understand each other's utterances.

The individualist will grant my observation and embrace that within their framework neither u_7 nor u_8 has any chance of expressing Putnam's utterance u_3, since u_7 and u_8 are too fine-grained in their respective contents.[111,112] It is perhaps better to moderate my observation: given the *epistemic* explanatory goal that I set forth, the individualist approach fails to adequately group the tokens into types.

In addition, individualists believe in a *social conversion* of individual conceptions regarding a given topic of conversation (i.e. dogs).

> There is a common underlying explanation of both deference and the exceptions to deference which makes no appeal to Burge-like considerations [such as socially shared *public* concepts and words]. More obvious commonsense sociological explanations, i.e. more pragmatic and historical explanations of deference, can be introduced instead of Burge's explanations, such as: if we wish to be understood and understood without strain, we will by and large […] use words as others do; if we have been brought up in the same social environment and our words and beliefs

[110] This doesn't mean that speakers couldn't fail in their attempt to repeat tokens produced by other speakers. I will discuss this problem in chapters three and four.
[111] Bilgrami's asserts that "at the meaning-theoretic level, […] concepts are very fine-grained and they are hardly ever shared by people" Bilgrami (1992): p. 11.
[112] Thanks to Ebbs for pointing this out to me.

> have been formed in it, our meanings will tend to converge and especially so if we want to bring others up in it.
>
> These alternative explanations belong to commonsense sociology and have no relevance to philosophy in the way Burge insists on.[113]

Only once we already have grouped the tokens into word-types intra- and inter-subjectively can we make sense of individualist conceptions converging. A social conversion must also be driven by effects within individuals – even if those individuals don't entirely understand what they are motivated by. What Bilgrami seems to have in mind here is something along the following lines: early on we get trained to call dogs 'dog', or red things 'red' – otherwise other agents won't respond to our utterances in the manner we need or expect them to. Again, this seems to pre-suppose that we can split the intention to repeat the word-form from the intention to use the word as it is used by *other speakers* – thereby introducing an anti-individualist ingredient of deference.

How would Putnam's conceptual apparatus be *singling out* tokens of Quine's word 'elm' and the behavior associated with *those* tokens as evidence from which to derive conceptions associated with the topic labeled 'elm' (within Putnam's cognitive system)?

Why not thus single out the behavior associated with other tokens? On what basis can the individualist even group (almost) all and (almost) only tokens of Quine's word 'elm'? The suggestion of a social conversion only makes sense, once we already assume adequate groupings of tokens into word-types.

Bilgrami's intuitions about the role of belief content in psychological explanations lead him to *reject* the intuition that Putnam and Quine are using the same word in *both* conversations. He argues that at

> the meaning-theoretic level, the concepts are very fine-grained and they are hardly ever shared by people.
>
> But this does not matter since it is not *these* "concepts," so thought of, which go into the contents that explain action. Action-explanation always takes place at a much more local level

[113] Bilgrami (1992): pp.79-80.

> than the meaning-theoretic level. Here, the entire aggregate of beliefs that an agent associates with "water" are not all relevant. One distils out of the aggregate of resources provided by the meaning-theory only those beliefs that are relevant to the action-explanation at the local level. Thus, if you and I are both drinking some substance from the kitchen tap because we want to quench our thirst with the cheapest available drink, we may *in this locality* both be attributed the same content: "… that water will quench thirst." […] The specification of your local water-concept in the attribution of content which explains your behavior does not contain your chemical beliefs. […] Hence, although our idiolects are never likely to be the same for any single concept, in many localities we may nevertheless share many contents. There will obviously be other localities in which we will not share contents because in those localities we *will* find it necessary to use your chemical beliefs from the overall pool in order to explain your action.[114]

I interpret Bilgrami to be saying – in the terminology I used so far – that he admits individual conceptions that different speakers associate with their words to be similar in some contexts (or localities), but not with regard to others. That is, from the "pool of resources" of meaning-theoretic concepts, the superfluous fine-grained beliefs are "distilled out" such as to leave us with the action-explanation content for each speaker's words or statements *at that locality*. If those contents are identical, supposedly communication succeeds in facilitating the coordination of behavior.

On what grounds do we identify Putnam and Quine's beliefs at the local level? Ebbs similarly objects:

> [Bilgrami's] idea that beliefs attributed locally are "selected" from beliefs attributed at the meaning-theoretic level pre-supposes that *some of the beliefs attributed at the meaning-theoretic level are also attributed locally*. But his own theory undermines this presupposition.[115]

On the one hand Bilgrami writes how his interest in meaning-theoretic concepts is entirely secondary – it is "merely a gesture acknowledging that a selection has to be a selection from somewhere." His real interest is in *local* contents. He continues: "the "locality thesis" dissolves the very idea of content composed of context-invariant concepts [thereby rendering his] account of intentionality […] genuinely radical."[116]

---

[114] Bilgrami (1992): p. 11.
[115] Ebbs (1998): p. 619.
[116] Bilgrami (1992): p. 12.

How can context invariant meanings ground the selection of local contents, while at the same time the existence of context invariant meanings is denied? Without grouping tokens into orthographic and phonetic word-types, individualists cannot account for any continuity in a given topic within or across contexts.

2.4 Resemblance is context sensitive, but semantic content is context *in*sensitive[117]

Even if we could separate steps (i) and (ii), that is, offer a robust account of orthographic and phonetic word types, insurmountable challenges would remain. Individual cognitive conceptions may "overlap" with regard to some conversations, but not all, while speakers interpret themselves as using the *same* words (associated with those cognitive conceptions) throughout those conversations.

What makes comparing cognitive conceptions associated with various tokens difficult is that relative to one context (and purpose) the conceptions under consideration may be similar, while they are significantly distinct relative to some *other* context.

Our human activities are so varied that which features are salient is *not* constant from context to context. Goodman remarks: "[I]n the case of ordinary actions, the principle of classification varies with our purposes and interests."[118]

I repeatedly catch myself thinking about some random group of things that a subset of that group is *de facto* (independently of my or anybody's point of view) a type because the members of that subset are *de facto* more similar to each other than to members outside that subset. It

[117] Semantic content of proper names and kind terms, that is!
[118] Goodman (1970): p.22.

takes a good amount of distancing myself from habitual categorizations, to see alternative groupings of types.

Consider the following three objects:

a. an aspirin tablet

b. an ibuprofen tablet

c. a pack of macadamia nuts

On a first look it seems that the aspirin and the ibuprofen tablet clearly form one class: they are examples of over-the-counter pain medications. In fact, sometimes people use 'aspirin' as a general term for relatively mild over-the-counter drugs that reduce pain, in which case it denotes either of the following: ibuprofen, acetaminophen or aspirin tablets (and maybe others). Dog-owners, in contrast, will, if it comes to something their dog may have consumed, put the ibuprofen tablet and macadamia nuts into one category. While moderate amounts of aspirin are harmless to dogs[119], even one ibuprofen tablet may cause a dog's death. If consumed by a dog macadamia nuts can lead to kidney failure, in the same way that ibuprofen does. Whether a. and b., or b. and c. form a relevant category – based on a relevant resemblance relation – depends on the context of the conversation. Are you looking for a headache remedy or the cause to your pet's vomiting?

What could constitute the apt standard for semantic or conceptual resemblance? We use words as means to communicate our thoughts to others successfully. Sometimes it is suggested that conversations are successful when they run "smoothly", that is, when it is in virtue of the

[119] Although adult aspirin tablets ought not to be given to very small dogs (less than 15 pounds). Then again, we wouldn't consider the same dosage to humans weighing less than that either.

words that they are exchanging that the interlocutors successfully *co-ordinate their behavior.*[120,121]

In the following section I will imagine two conversations. One turns out to be successful according to that standard, the other not. Sarah believes 'aspirin$_S$' refers to aspirin tablets. James, on the other hand, believes 'aspirin$_J$' to be a generic term for over the counter pain medications (including ibuprofen, aspirin and acitaminophen tablets).

Conversation 1:

Sarah: "Do you have some aspirin?"

James: "Yes, I have some aspirin in my pocket"

James hands Sarah an *ibuprofen*. Sarah swallows the tablet and her headache disappears.[122,123]

---

[120] Quine proposes a behavioristic approach to translation. Since stimulus meaning cannot be assumed the same across speakers for a given observation sentence, Quine abandons inter-subjective likeness of stimulus meaning for less technical success conditions: "[T]here is nothing in linguistic meaning beyond what is to be gleaned from overt behavior in observable circumstances. […] Success in communication is judged by smoothness of conversation, by frequent predictability of verbal and non-verbal reactions, and by coherence and plausibility of native testimony" Quine (1992): pp. 38-43.
Paul suggests that communication succeeds if senses are similar enough to allow speaker and hearer to co-ordinate their future actions. Paul (1999) pp. 161-162.

[121] I am deliberately sketching this notion of communicative success. If it fails to work while being painted in broad brush-strokes, why would it fare better once we are picky about the details?

[122] Pagin considers a similar case based on Putnam's Twin Earth thought-example.

> [S]uppose we have two different concepts of water, *water$_1$* [H_2O] and *water$_2$* [XYZ]. […] Jill associates the concept *water$_1$* with the mass term 'water', whereas Jack takes it to express *water$_2$*. Jack says to Jill:
> (1) I have a glass of water in my hand.
> By the identity standard [word-sameness], communication has failed. […] The question is whether this matters to the success of the actual communicative event. There is, I submit, a pretty strong intuition to the effect that it doesn't.

See Pagin (2006).

[123] Heck presents us with a similar example. He envisages an North American and a British speaker to talk about pies, each using their own word 'pie'. As it turns out in England a proper pie is required to have a pastry top, whereas in the United States it is optional. Heck observes that we would consider this to be a successful communicative act, regardless of the distinct dialects in use – that is, word-sameness is not necessary for communicative success (Heck, Richard (2006) p. 87).

Sarah and James *successfully* coordinated their behavior and their respective words 'aspirin$_S$' and 'aspirin$_J$' played an important role – they were the "primary" carriers of the information exchanged. Given the overall context and purpose of the conversation, the *overlap* between Sarah's conceptions regarding what 'aspirin$_S$' denotes and James's conceptions regarding what his word 'aspirin$_J$' denotes was *substantial*.

Conversation 2 (occurring a little later in the day):

Sarah: "Could I have another of those aspirin? Dakota, my dog, is suffering terribly from arthritic flair-ups today."

James: "Here, this is my last aspirin tablet."

When James passes on the bottle of ibuprofen, Sarah catches a glimpse of the label and reads 'Ibuprofen'.

She scolds: "James, ibuprofen is not aspirin. You can't give ibuprofen to a dog!"

James: "I'm sorry Sarah. I didn't know that."

In conversation 2, the distinct conceptions Sarah and James associate with their respective word 'aspirin$_S$' and 'aspirin$_J$' lead to failed co-ordination of behavior. If Sarah hadn't caught a glimpse of the label, she may have administered human pain medication to her dog and thereby possibly killed it.

Conversation 1 brought about successful co-ordination of behavior, while conversation 2 didn't. According to the above-mentioned criteria of conceptual resemblance, we would have to argue that Sarah and James shared the same word in conversation 1, but not so in their later exchange. We cannot say that Sarah and James shared the word in question in conversation 1, but failed to do so in conversation 2. Sarah would judge that her respective tokens were tokens of the same word type. James would say the same about his tokens that are meant to be

repetitions of his word 'aspirin$_J$'. We cannot take the ability to reliably group intra-subjective tokens into word types away from individuals or we introduce all sorts of self-knowledge problems (which individualism is supposed to be able to avoid).[124] Individualists cannot allow for speakers to be wrong about individuating their own words. If a speaker were to be wrong about how she relates her own words to each other, then the constitutive conditions for understanding the corresponding concepts would *outrun the network of inferences* that she can draw – hence threatening the individualist's psychological explanatory goals.

If we accept that Sarah and James are respectively using tokens of the same word types 'aspirin$_S$' and 'aspirin$_J$' in conversations 1 and 2 (in accordance with their own judgments), then 'aspirin$_S$' and 'aspirin$_J$' *cannot* constitute replicas of the same word-type with regard to conversation 1, but replicate *distinct* word-types with regard to conversation 2.

The individualist ought to propose that Sarah and James have to be *generally* successful in their conversations when using their words 'aspirin$_S$' and 'aspirin$_J$'. Of course, the mis-communication may happen to never *actually* become apparent. We therefore must consider a wide-range of *potential* conversations between Sarah and James. In some cases the conversation may not go "smoothly" at all if Sarah is using the word 'aspirin$_S$' to talk to medical experts. Imagine a group of pediatricians discussing the amount of *aspirin* safe to administer to an undernourished two year-old infant. Although Sarah knows how to effectively use the word 'aspirin$_S$' in day-to-day exchanges, she lacks the necessary knowledge to participate in an exchange among experts. The possibility that either interlocutor may be under- or mis-informed about a given topic can in and of itself cause a "bumpy" conversation (or interfere with the co-ordination of behavior), but shouldn't therefore be an indication that communication failed in the

[124] Bilgrami (1992): pp.18-19.

sense that the relevant tokens were not tokens of the same word type. Demanding general success in co-ordination of behavior across a variety of situations introduces far too stringent demands on word-sameness. In that case, speakers would barely ever replicate each other's tokens (as tokens of the same word type relevant to the communication of knowledge).

The anti-individualist would argue that so far Sarah and James's example was under-described. In order to know whether they are using the same word 'aspirin', we need to know whether they are using the same public word. There is a generic word 'aspirin' that matches James's use of '$aspirin_J$' and there is a brand'-name 'Aspirin' that matches Sarah's use of '$aspirin_S$'[125]. The generic 'aspirin' and the brand name 'Aspirin' constitute *distinct* words, in the sense that they cannot be interchanged in a conversation without changing the topic (they cannot be truth-functionally interchanged in statements to the content of which they contribute non-vacuously[126]). According to the anti-individualist then, Sarah and James are (probably) using distinct words, if there is a parallel generic and specific use of the word-"form" 'aspirin' in their linguistic community. The individualist and anti-individualist analysis thus far agree.

In a linguistic community, however, where no generic term '$aspirin_G$' exists (and thus we cannot translate James's word '$aspirin_J$' as '$aspirin_G$'), the outcome could be quite different. James may simply be *mistaken* about whether ibuprofen is *aspirin*. Imagine the generic word '$aspirin_G$' didn't exist and somebody claimed of an ibuprofen tablet "this is aspirin". We would

[125] At least this is true in Germany. In the United States the brand name in question is 'Bayer Aspirin'.
[126] Some may want to argue that the word 'aspirin' is ambiguous. I am analyzing words as tools in our inter-subjective transfer of thoughts (about the external world). The content of the thought changes depending on whether I am thinking of Aspirin the brand, or more generally of over-the counter pain medication. I will thus have to argue that ambiguity pertains to word-forms. The word-form 'aspirin' is ambiguous between two words '$aspirin_G$' and '$Aspirin_B$'.
My view is in line with Kaplan's: "[On] my conception, there are two phonographic words "base" (meaning '*low*' and *bottom*), not, as the orthographic conceptions would have it, a single word with two meanings" Kaplan (1990): p.100.

simply say that the person has the false belief that ibuprofen is aspirin.[127] So, whether the linguistic exchange is successful (in the sense relevant to whether they understood each others assertions), ultimately depends on how we individuate James and Sarah's words – not the other way around. If James hands Sarah an ibuprofen tablet (which is *not* what Sarah asked for), the conversation itself may still be successful (in both imagined linguistic exchanges). James sincerely conveys to Sarah that he believes that what he handed her was *aspirin* (Aspirin) – but his belief is false. It was his mistaken conceptions associated with 'aspirin' that almost endangered the life of Sarah's dog[128], not his use of a different word (a word Sarah is unfamiliar with).

Public language accounts provide us with a *reference point* against which to measure individual token productions as "matching" standard productions (i.e. spoken and written productions). Do other speakers *recognize* my mark 'dog' as a replica of their word 'dog'?

2.5 Conclusion

With the collapse of steps (i) and (ii), individualism must explain why two word tokens *t* and *t'* are of the same word type in terms of resemblance features between the conceptions associated with what *t* and *t'* denote respectively. I have demonstrated that the individualist framework

[127] In extreme cases if the person is completely unwilling to be corrected in her use of the word in question, we may have to say that she is not using a word that we can find in our own vocabulary.

[128] We find similar objections that demonstrate that individual conceptions are not sufficient for word-individuation. See also Burge (1979); Putnam (1975), Kripke (1972) (although Kripke's criticisms are targeted at descriptivism in general).
Burge's arthritis patient perhaps shows that they aren't necessary. It depends on whether we consider the patient and the doctor's conceptions regarding their respective words 'arthritis' as different or roughly similar. If we allow this kind of discrepancy between individual conceptions to still relate to the same word, we will run into problems distinguishing conceptions that relate to words denoting similar kinds of things (i.e. elm versus and beech trees).

cannot accomplish this goal – if we are interested in words as communicative tools through which we convey *knowledge* to each other.

CHAPTER 3

UNDER WHAT CONDITIONS ARE TWO UTTERANCES PERFORMANCES OF THE SAME WORD? PROPOSAL: APPROPRIATE INTENTIONS

Niemand würde viel in Gesellschaft sprechen, wenn er sich bewusst wäre, wie oft er die anderen missversteht.
-- Johan Wolfgang von Goethe, Die Wahlverwandtschaften

No one would talk much in society if they were aware how often they misunderstand others.
-- Johan Wolfgang von Goethe, Elective Affinities

3.0 Introduction

Kaplan's "Words," lays out the groundwork for a *new* theory. It lures us away from the resemblance-based model and offers an alternative framework based on intentionality. Most objections against Kaplan's proposal show it to be in need of refinement, but are not targeted at its basic assumptions. My analysis is primarily focusing on how Kaplan answers the question: under what conditions are two utterances performances of the same word? His answer is simple: if a speaker produces token *t'* with the (appropriate) *intention* to repeat token *t*, then *t* and *t'* belong to the same repetition word-tree – they both constitute performances of the same word.[129] Since we can clearly imagine cases where a speaker fails in her attempt to repeat *t*, Kaplan owes us an account of what constitutes an *appropriate intention* such that within his framework we can make sense of these repetition failures.

[129] Kaplan (1990) rejects the type token distinction altogether and would object to the above formulation.

3.1 Kaplan's new notion of words

With resemblance relations eliminated from the theoretical repertoire, Kaplan replaces those failed criteria of word-identity with individual speaker intentions. Resemblance is degraded to solely providing us with clues as to which intentions a speaker may have had – given that we cannot access other speaker's intentions directly.

> The identification of a word uttered or inscribed with one heard or read is not a matter of resemblance between the two physical embodiments (the two utterances, the two inscriptions, or the one utterance and one inscription). Rather it is a matter of intrapersonal continuity, a matter of intention: Was it *repetition*? We depend heavily on resemblance between utterances and inscriptions (using *resemblance* here not to mean matching of physical characteristics but of their appearance as we look and listen) in order to divine these critical intentions. If it sounds like "duck", it probably is "duck". […] It is [intention] that decides the matter.[130]

Tokens *t* and *t'* are of the same word type, if one was produced with the intention to repeat the other. Sometimes intermediate tokens (finite in number) t_1" through t_n" link *t* and *t'* in the described manner. For example, *t* is produced as a repetition of t_1", while t_1' is produced as a repetition of *t'*. All tokens *t*, *t'* and t_1" belong to the same performance continuum *T* – that is, they all constitute performances of the same word (in the thick sense of word-sameness). Kaplan calls the resulting complex 'repetition trees'. Any subsequent tokens t_2"…t_n" produced with the same intention to constitute repetitions of the tokens *t*, *t'* or t_1" in turn become nodes of repetition tree *T*.

> Common currency names (and other common currency words) are not abstract constructions, they are natural objects. Not physical objects, though most will have physical embodiment at many places and times. And not mental objects, though most will have mental embodiment (an oxymoron?) at many places and times.
>
> One might think of them as trees. Stemming out from their creation, with physical and mental segments […].[131]

[130] Kaplan (1990): p. 104.
[131] Ibid.: pp. 116-117.

Kaplan thus ascribes a *constitutive role* to the (appropriate) speaker intentions. Those word "trees" are shared across speakers.

The type-token model is abandoned in favor of a naturalistic notion of words: the stage-continuant model.

> I propose a quite different model according to which utterances and inscriptions are stages of words, which are continuants made up of these interpersonal stages along with some more mysterious *intra*personal stages.
> […] The identity of a common currency word lies in its continuity, both interpersonal and intrapersonal.[132]

In the last chapter I wondered how we know that Putnam's word 'elm_P' is the same word as Quine's word 'elm_Q'. Putnam was assumed to be significantly under-informed about elms (not being able to tell a beech from an elm tree) while Quine was ascribed arborist-like knowledge of elms. According to Kaplan's theory, Putnam's inscription 'elm' [t_1] and Quine's (verbal) utterance 'elm' [t_2] are stages of the same continuant (or repetition "tree"). It is Putnam's *intention* to *repeat* Quine's utterance t_2 that makes Putnam's token t_1 belong to the same repetition tree as Quine's token t_2. Quine's token t_2 in turn reaches all the way back through a repetition chain to the original token that gave birth to our word 'elm'.

Kaplan offers us a strange hybrid version of words: they are natural objects but can be embodied at several places at the same time – a characteristic usually reserved for *abstract* kinds. Wetzel raises an objection against Goodman's attempt to eliminate abstract semantic objects for a nominalist theory, which equally applies to Kaplan's proposal.[133] She challenges the nominalist to explain (without reference to types) why the following statement is true:

(i) 'Paris' consists of five letters.

[132] Kaplan (1990): p.98.
[133] Wetzel, Linda (1999); Weztel, Linda (2009): chapter 5.

Not all nodes of the 'Paris' repetition tree contain five letters.[134] Apparently some outdated performances looked (or sounded) like this: 'Parrys' or 'Pareiss'[135] and some are misspelled: 'Pari', 'Pariss', 'Pairis'… Kaplan would have to argue for either of the following:

A_1 A correct repetition of 'Paris' consists of five letters.

But then what does it mean to be a correct repetition? Doesn't that suggest that we are copying something – a standard form? Kaplan doesn't accept this "myth".

A_2 A repetition of 'Paris', which is generally recognized as a correctly spelled repetition consists of five letters.

But who is to judge? Clearly, not all mistaken repetitions of 'Paris' constitute a slip of the pen. Some people simply don't know how to correctly spell 'Paris' and would constitute bad judges. Since Kaplan is a public language proponent, he could allow individuals to defer to the relevant experts with regard to the correct spelling of 'Paris'.

A_3 A repetition of 'Paris', which is recognized as a correctly spelled repetition by the relevant experts consists of five letters.

Sometimes a person's handwriting is simply illegible. That person's inscription of 'Paris' will not be recognizable as one replicating 'Paris'. Let's suppose we already have a correctly spelled inscription 'Paris' as a starting point. If the producer of the repetition token views herself as having correctly copied the sample standard token, then the copy will consist of five letters. The sample standard tokens may be found in reputable dictionaries or were otherwise produced by relevant experts.[136] To some this may not be as robust of an analysis of (i) as they were hoping. Given our explanatory goal, individuating words for the purpose of drawing valid inferences and

[134] I am setting aside here that we also don't have an account for letter types.
[135] Wetzel (2009): p. 95.
[136] Those tokens cannot constitute slips of the pen.

reliably including information based on other speaker's testimonies – we don't need a more robust account of correctly spelled tokens.

A consequence of Kaplan's ontology is that words don't exist outside the repetition trees. Although in the course of our existence we can say a whole lot, this means that there will only be finitely many words. "There being only finitely many [performed] sentences", according to Wetzel, "entails, among other things, that although A and B might be sentences, their conjunction, (disjunction, equivalence, negation) need not be."[137]

This violates some key assumptions in logic. If sentences A and B are individually provable in system S with axioms a_1, a_2, $a_{3...}$, then so should their conjunction. Something that doesn't exist is not provable. If the conjunction of A and B is never performed (in the past, present or future, that is, never part of a repetition tree), then the conjunction of A and B is not provable in S, although A and B are individually.

Quine recognizes the problem and offers the following solution:

> A more humdrum reason for supposing that the propositions outrun the eternal sentences could be that for many propositions the appropriate eternal sentences, though utterable enough, just happen never to get uttered (or written [or otherwise performed]). […]
>
> *Prima facie* the answer is that a sentence is not an event of utterance but a linguistic form that may be uttered often, once, or never;

We have shown in the first chapter that we cannot reduce word or character types to physical events. Quine goes on:

> But we must not accept this answer without considering more precisely what these linguistic forms are. If a sentence were taken as the class of its utterances, then all unuttered sentences would reduce to one, viz., the null class; they might as well not exist so far as propositions are concerned, for all distinction lapses among them. […] We can take each linguistic form as the *sequence*, in a mathematical sense, of its successive characters or phonemes. A sequence a_1, a_2,…, a_n can be explained as the class of the n pairs $<a_1, 1>$, $<a_2, 2>$…, $<a_n, n>$. […] We can

[137] Wetzel, Linda (2000): p. 367; Weztel intends this to be an objection against Goodman's nominalism, not Kaplan's notion of words.

> still take each component character a_i as a class of utterance events, there being here no risk of non-utterance.[138]

By adopting an ontology that includes sets or classes of signs or characters Quine renounces a strict "Goodmanesque" nominalism.[139,140]

Kaplan stresses that the notion of 'word' he introduces is one of natural objects, so I don't think Quine's solution is available to Kaplan. Perhaps Kaplan would argue that he is only interested in *actual* proofs and *actual* inferences – that he is not interested in such formal relations.

A legitimate move in the German language game is to "glue" together several nouns such as to produce a new noun. For example: The German word '*Blumen*' (plural of 'flour') and the word '*Topf*' ('pot') and '*Erde*' ('soil') can be added together to a new meaningful string of letters: '*Blumentopferde*'.[141] In principle, there could be such a combination of nouns '*a*' and '*b*' that if '*ab*' were uttered '*ab*' would be instantly meaningful to a competent German speaker. Let's say this noun '*ab*' never actually is performed. Therefore, Kaplan would have to deny that '*ab*' is a word of the German language. The problem with this objection is that we could never actually provide a counterexample. As soon as it is uttered or only thought of, it ceases to be an unperformed word.

---

[138] Quine (1960): p. 194-195.

[139] In an earlier paper written with Goodman, Quine renounces all abstract entities (including classes) he admits: "we cannot say in general, given any two inscriptions, there is an inscription long enough to be the concatenation of the two" (Goodman & Quine (1947)). This means that Quine doesn't think abstract objects *in general* problematic, but only renounces propositions, predicates etc. Even though signs also have an acausal nature, we certainly seem more *confident* in our interaction with signs, then in our identification of propositions.

[140] Wetzel and Ebbs endorse Quine's solution. Wetzel, Linda (1993); Ebbs (2009): pp. 119-120.

[141] '*Donaudampfschifffahrtsgesellschaftskapitänshutsfederfarbe*' is one German noun. It roughly means: the color of the feather of the hat of the captain of the Donau shipping company. With sufficient creativity this noun can be extended (almost) indefinitely.

Kaplan could argue that these German words could be treated on a par with compound words. A speaker is thus linked to the repetition tree of '*Blumentopferde*' if she is linked to each component basic word: '*Blume*', '*Topf*' and '*Erde*'.[142] Especially, since Kaplan doesn't require correct pronunciation or spelling, I doubt that without dashes and space marks, it will be clear how to break down compound words into their component parts. The word '*Blumentopferde*' also contains the plural of 'horse': '*Pferde*'. The compound word thus may also be read as denoting a type of horse (although, admittedly in this case 'Blumento' is not a meaningful component and neither is 'Blumento-Pferde' a meaningful word). I may be very good at repeating the components '*Blumen*', '*Topf*' and '*Erde*', but when I see the word '*Blumentopferde*' and repeat '*Blumento*[pause]*pferde*' I clearly *failed* to repeat the compound word adequately. I take it to denote a specific type of horses, rather than soil for flowers.

3.2 Are Kaplan's performance standards[143] too lenient?

Cappelen accuses Kaplan's of not introducing appropriately stringent performance standards:

> A proponent of the sufficiency thesis holds that in order to produce a token of a certain word it is sufficient for the producer of the token to be in some intentional state at the production time. According to Kaplan, the intentional state is that of standing in the repeating relation to a previous interpersonal stage of a common currency word.[144]

He goes on to invite us to imagine something of the following: a person attempts to write the letter 'l', but he ends up merely with a squiggle looking something like this:

[142] Thanks to Ebbs for bringing this option to my attention.
[143] The term 'performance standard' is taken from a forthcoming paper from Lepore, Ernest and Hawthorne, John entitled "On Words,".
[144] Cappelen (1999): p. 94.



Figure 3.1[145]

Kaplan acknowledges "the difference in phonographology, the difference in sound or shape or spelling, can be just about as great as you would like it to be."[146]

Let's not forget that according to Kaplan, we attempt to reproduce our own standardized replicas of a given word or character. As a result, it is up to the individual to impose standards on her replicas. I find it hard to believe that anybody would *in good faith* claim the above squiggle to be a replica of the character 'l' (unless we extend our symbolic system such that the above mark is to count as a replica of the mark 'l').[147] Then again, Georges Braque, while developing cubism, must have said to himself while looking at a painting with the eyes, nose, hair-line and mouth scattered somewhat randomly throughout the canvass "yes, this is a face!"

[145] Cappelen (1999).
[146] Kaplan (1990): p. 101.
[147] We put our novice language producers through rigorous training. Recall the endless repetitions of the letter '*A*' every student has to draw and carefully press in-between the (pre-printed) triple lines until the teacher is satisfied with the output. This mimicking process beings much earlier when toddlers play around with names. Sofia: 'Nadli?' – the adult: "No Sofia, it's N-A-T-A-L-I !" Still, on order for the adult to even suggest an improved pronunciation of her name, she must recognize the attempt as one performing her name.



Georges Braque "Woman with Guitar" (1913)

Even in the painting entitled "Woman with Guitar" we recognize, after some observation, a face. If we were looking at the original painting with a friend who simply couldn't recognize the face, we would talk them through the process of discovery. "Look here are the eyes, and this is the mouth …" Hopefully, we can thus bring our friend to see the face in the picture.

Perhaps the producer of the squiggle could explain to *us* how he wants us to look at the mark as a performance of the letter 'l'. He can't just stipulate: "It is a replica of the letter 'l' because I intend it to be!" It also isn't acceptable for him to tell us that during the process of production he had been bumped accidentally. Recall my mispronunciation of the word 'reciprocal', where I placed the accent on the letter 'o', rather than the letter 'i'. Although at first my interlocutors couldn't recognize my utterance as a replica of the word 'reciprocal', friends of mine who acted as "translators" were able to make them hear my replica as one of 'reciprocal'. With some

guidance, my interlocutors were able to hear my 'recipr<u>o</u>cal' as a performance of the familiar 'rec<u>i</u>procal'.[148]

A replica of a character or word must be intentionally produced in *good faith*. Additionally, the outcome has to meet the standards of the *performer*. I assume that if these conditions are met, the audience can be *brought to recognize* the mark as a replica of a character or word – even if it takes a little explanation by the performer or more charitable observers. Kaplan demands that the mark has to be produced with an *appropriate* intention. Perhaps part of what makes an intention appropriate is that the speaker wants to communicate successfully with her audience, which brings about that her performances are (in principle) *recognizable* to others.[149] This is why resemblance, or more precisely the perception of resemblance, still plays an enabling role in practice – albeit not an explanatory role in theory.

While Cappellen's last attack on Kaplan intended to demonstrate that intentions are not sufficient to repeat, his next objection attempts to show that intentions are not necessary in the creation of the word- or character-replica. He envisions finding a piece of paper displaying the following marks:

CAN YOU SPARE A QUARTER?

Cappellen invites us to imagine that this mark was produced as a result of an accidental spilling of ink. Perhaps a cat was playing with a fountain pen and the above mark was the end result.

---

[148] It is surprising how good we are at individuating "deviant" words. What the following passage states is incorrect – the alluded study was not conducted in Cambridge – but our ability to understand each word is surprising nevertheless:

> Aoccdrnig to a rscheearch at Cmabrigde Uinervtisy, it deosn't mttaer in waht oredr the ltteers in a wrod are, the olny iprmoetnt tihng is taht the frist and lsat ltteer be at the rghit pclae. The rset can be a toatl mses and you can sitll raed it wouthit porbelm. Tihs is bcuseae the huamn mnid deos not raed ervey lteter by istlef, but the wrod as a wlohe.

From: http://www.mrc-cbu.cam.ac.uk/people/matt.davis/Cmabrigde/
Most people report that they can read the above passages without much difficulty.

[149] See also Ebbs (2009): pp. 124-125.

Next, in this thought-example, Cappellen holds this very piece of paper up to random passers-by in order to collect change. He does so without ever letting a word slip out of his mouth or physically altering the found piece of paper.

I suppose Cappellen wants us to reason as follows: He communicates his (hypothetical) desire (or need) for some change via that piece of paper. At that point the piece of paper clearly states the question 'Can you spare a quarter?' Cappellen did *not produce* the string of marks in question, so he did *not produce* the meaningful string of words. Therefore, it seems the question must have already existed on the piece of paper when Cappellen found it.

An analogy with art-pieces constituting found objects should be illuminating. In the movie "Ghost World" the main character Enid finds a historical poster-advertisement from a fast-food chain depicting a grotesquely caricatured black man. Enid (without claim of having designed or produced the poster) displays it as her art-piece in an exhibition. She doesn't physically alter the poster; she merely alters its location and contextual setting. Enid's intention is to provoke disgust in the observer as a response to the overt display of racism on the poster and thereby cause the audience to reflect on our blindness to present expressions of racism – which are less overt and thus more resilient. Placing the poster in a meaningful context and picking-up on its importance is Enid's contribution to the piece of art. Sometimes advertisement posters can be considered works of art, but this one lacked any artistic value. When did the mere poster become a piece of art? Enid did *something* to it to cause the transformation, without physically altering the poster. It is her invitation of reflection via the poster – placing it in a *new* context – that changes its status. The same goes for the marks on the piece of paper. Cappellen produced the question, without physically altering the marks on the piece of paper, by placing it in the

appropriate communicative context. Just because Cappellen (in the thought example) didn't produce the physical marks doesn't mean that he didn't produce the meaningful string of words.

If the analogy with art-pieces composed of found objects doesn't speak to you, allow me to present my own linguistic thought example that reveals how Cappellen's assumptions lead to ridiculous conclusions. Suppose I am missing a comma in a sentence that I was just in the process of writing, when my pen runs out of ink. I then walk downstairs into the kitchen to select a single kernel of rice from the Basmati-bag. Carefully, I place the rice-corn where a comma is needed (at the correct location and with the right orientation) and glue it onto the page. We would certainly recognize the rice-corn as a comma (albeit an unusual one). The rice-corn existed as a physical object, a mark, prior to my using it as a replica of a comma. It is absurd to claim that it already was a replica of a comma prior to my using it as such. In that case I would be storing a whole bag full of commas in the kitchen![150,151]

Cappellen may simply reject my supposition that a rice corn could (under the right circumstances) replicate a comma. It looks like a comma, without "properly" replicating a comma – or so he might argue. He would maintain that my audience is interested in what I wrote (rather than in my ability to produce character-tokens) and thus is willing to fill-in the missing elements of my written string of words – the comma.[152] I don't deny that we use

[150] See also Ebbs's objections (Ebbs (2009): p. 124); Alward calls for a more charitable interpretation of Kaplan than Cappellen is willing to extend (Alward (2005): p. 177).

[151] Independently of Cappellen's objections to Kaplan's notion of words, Putnam rejects the idea that any "physical object can, in itself, refer to one thing rather than to another". He insists: "nothing physical has 'intentionality', save as that intentionality is derivative from some employment of that physical thing by a mind" (Putnam (1981) p. 2).

[152] "What these cases show is that charity in interpretation often extends to the classification of ink marks, sounds etc. In some cases we overlook the fact that the speaker didn't produce a token of the word she tried to produce. We know which words she tried to produce and since we are interested in what she has to say (rather than in her ability to produce word tokens), we give the speech act the interpretation we would have given it had the speaker managed to produce a token of the word we know she tried to produce" (Cappellen (1999): p. 95).

contextual clues to figure out what word or string of words the speaker intended to produce. But I disagree with the comma having been missing altogether and instead having been added by the observer. The rice corn was merely harder to recognize than a standard comma! The idea that words have to be produced with standard materials (i.e. with ink on paper) is not tenable, once we consider the already accepted alternatives: words carved in wood, traced in sand, formed out of modeling clay and so on. We don't want to argue that a rice-corn intrinsically constitutes a comma and we don't want to strip it of its potential semantic content either. Our constitutive intentions turn the rice-kernel into a comma!

3.3 Considering standard anti-individualist thought-examples

According to social anti-individualism, speakers can be partly wrong or under-informed about what one of their words denote.[153] Sometimes the speaker has only minimal information about what a word denotes. I use a variety of *unlabelled* spice mills for coriander, mustard seeds, pepper and cumin. My husband is able to identify the pepper, but he starts guessing on the content of the other mills. So, if a recipe asks for coriander, he will consult me: "which of the jars contains coriander?" In this case he at least knows that the word 'coriander' refers to a spice. He even knows that it comes from the same plant as the fresh herb cilantro. When it comes to picking out a coriander sample from *similar* seeds, however, he struggles. He is partly under-informed about what 'coriander' denotes. Assuming that the author of the recipe knows what she is talking about, it is her word 'coriander' that he attempts to repeat to me.[154,155] This

[153] "Participant commony regard their object-level thoughts (thoughts about, say, chairs) as undergoing correction in the course of the inquiry" Burge (1986, 2007): p. 261.
[154] That is, he is assuming that she didn't make up a new word 'cilantro$_{NEW}$' that denotes some ingredient other than cilantro. He assumes that the author of the respective recipe is using her word 'cilantro' to denote cilantro.

case is best viewed as one where my husband is under-informed about a given topic[156] and lends support to Kaplan's notion of words.

In one of the standard anti-individualist[157] thought-examples, Burge conceives of a patient who has been diagnosed with arthritis in both knees. Let me call this patient 'Alice'. She is overall competent in the English language. At one of her follow-up visits she conveys to her doctor: "I believe that arthritis has spread from my knee to my thigh." In the appropriate circumstances she would utter sentences containing the term 'arthritis' (non-vacuously):

'I have arthritis in my knees'.

'Certain aches are characteristic of arthritis'.

'It is better to have arthritis than cancer of the liver'.

'There are various kinds of arthritis'.

'Arthritis often afflicts the elderly'.

The latter three utterances are true, if Alice's word 'arthritis' denotes arthritis. Prior to examining the patient, the doctor informs Alice that arthritis is a condition which couldn't afflict any other parts of the body, but the joints. Alice is surprised as well as relieved and goes on to ask about what could be wrong with her thigh.[158,159] How should we interpret Alice's word

---

[155] Putnam presents us with a similar thought-example (Putnam (1975): pp. 226-227).

[156] Burge elucidates: "A primary impetus for my discovering the thought experiments was recognizing how many words or concepts I went around using which I found, on pressing myself, that I did not fully understand. I came to realize that this was not just a personal weakness. It was part of the human condition, at least in complex societies" (Burge (2007): p. 175).

[157] Burge defines anti-individualism thus: "The main idea of [… anti-individualism] is that the natures and correct individuation of many of an individual person's intentional, or representational, mental states and events commonly depend in a constitutive way on relations that the individual bears to a wider social environment. […] The differences in the social environment bear on the meanings of words and the ways words are connected through social chains to their subject matters" (Burge (2007): p. 151).

[158] Burge makes it part of the description of his thought experiment that the patient "thinks falsely that [she] has arthritis in the thigh". Burge, *using* the patient's term 'arthritis', begs the question as to whether the patient has thoughts about arthritis. I presented a paraphrased version of Burge's thought experiment,

'arthritis'? It seems most "natural" to describe the exchange between Alice and her doctor as one where Alice *learned* from her doctor that arthritis cannot afflict muscle material.[160] If Alice's word 'arthritis' denotes arthritis, we are in a position to say that her utterance "arthritis has spread from my knees to my thighs" expresses her belief that arthritis has spread from her knees to her thighs – which is *false*. Alice *revised* her mistaken conceptions about arthritis based on what her doctor told her. This presupposes that Alice and her doctor share the same word 'arthritis', although she is severely mistaken about what that word refers to. She also grasps the concept *arthritis* although she isn't able to grasp it fully – her conceptions that she associates with what 'arthritis' refers to considerably diverge from the concept *arthritis*.

Notice, that if you think that we shouldn't interpret Alice's utterance "Arthritis has spread from my knee to my thigh" as an utterance about arthritis, then we also cannot interpret her "true" utterances[161] as beliefs about arthritis. I think this would simply not coincide with how we generally interpret other speakers' words in practice.

Alice learned that she doesn't have arthritis in her thighs since this condition doesn't occur in parts of the body other than the joints. In fact, this is how Alice would interpret her own epistemic experience. Alice would not view herself as having learned a *new word*, but instead as having acquired a *new belief*.

Burge explains:

> In the course of the dialectic, we *stand corrected*: we recognize ourselves as convicted of *mistakes*, not merely infelicitous strategies for communication. We come to know something that characterizes empirical entities […] Usually, all participants begin the discussion without being able to give a precisely correct normative characterization [norms of identification regarding the

---

the description of which doesn't assume that the patient's utterances in question are about arthritis (Burge (1979, 2007)).

[159] I once witnessed a patient asking her physical therapist whether the pain in her quad is caused by arthritis. This reinforced my belief that Burge's thought experiment isn't far fetched.

[160] Ebbs (2002).

[161] 'True utterances': The utterances that are true if they are spoken by a competent 'arthritis'-user.

> empirical entities]; all or most would make minor errors in attempting to do so. But this does not entail that any lack object-level thoughts expressible, by the rest of us, with the term whose meaning is in question. If it did, *most people would have few if any object-level thoughts so expressible – an absurdity* [emphasis added]. Participants commonly regard their object-level thoughts (thought about, say, chairs) as undergoing correction in the course of the inquiry. They stand corrected on substantive matters.[162]

It should not be overlooked that Alice would in retrospect (post-correction) interpret herself as having used the same word throughout the conversation *and* as having repeated the doctor's word 'arthritis' – that is, she would interpret herself as having employed the same word 'arthritis' as the doctor.[163] Perhaps Alice is merely viewing herself to be repeating the mark 'arthritis' but is not bound to use it to denote arthritis. Burge anticipates this objection[164] and reminds us how Alice is open to *corrections* by the doctor regarding the reference of 'arthritis'[165] and wants to *learn* more about arthritis from him (in addition to being cured from whatever condition is causing the pain in her thigh). Burge stresses how *Alice would regard herself* as having had a *false* object-level thought. That is, a thought that pertains to the condition arthritis, not *merely* the word 'arthritis':

> The subject may maintain that [her] reasoning did not fix upon words. [She] may be brought up short by metalinguistic formulations of [her] just-completed ruminations, and may insist that [she] was not in labels. In such cases, especially if the reasoning is not concerned with linguistic issues in any informal or antecedently plausible sense, attribution of an object-level thought content is supported by the relevant evidence, and metalinguistic attribution is not.[166]

---

162 Burge (1986, 2007): pp. 260-261.

163 When we begin to identify a phenomenon it is not only the laymen, but also the experts who associate mistaken and incomplete conceptions with it. If conceptions individuate words, then given how little even experts know about arthritis – it would follow that *nobody* can talk about that condition – whatever it turns out to be – using our *present* word 'arthritis'.

164 Burge (1979, 2007): p. 122.

165 Burge clarifies: "We may imagine a cast, ragged network of interdependence, established by patterns of deference which lead back to people who would elicit the assent of others. […] To put it crudely, a person counts as among the *most competent* if he or she would be *persuasive* to other competent speakers in the use and explication of the language. The point about persuasion is fundamental" see Burge (1986, 2007): p. 259.

166 Burge (1979, 2007): p. 127.

If Alice were merely repeating a mark, she wouldn't need to yield her beliefs regarding what her word 'arthritis' denotes to the doctor's statements such as: "arthritis only afflicts the joints." Therefore, it is much more plausible that Alice is using *our* word 'arthritis', which lends support to Kaplan's notion of words.

3.4 Does Kaplan solve word-performance ambiguity?

Putnam observes that if we were to encounter somebody, let's call him 'Oscar', pointing at a snowball uttering: "This is a tiger!"[167] – there wouldn't be much use in talking to Oscar about tigers using *our* words 'tiger' or 'tigers'.[168, 169] Why not? Does this mean there are limits to the degree to which a speaker may be under- or misinformed about a given topic? If so, can Kaplan accommodate such limits with his theory?

I have earlier shown Cappellen's criticisms to be unfounded. Nevertheless, I agree with Cappellen that the speaker intending to repeat a word doesn't thereby *guarantee* her success. Unlike Cappellen, I am not concerned with radically deviant marks in terms of physical features. Instead, I am interested in cases where the speaker is fully capable to produce standard marks, such as "angina", but the speaker still fails in her attempt to repeat the respective public word. How can these phenomena be explained? What interferes with the linguistic commitment to a given public word? I raise doubts as to whether the intention to repeat a token *t* originally

---

[167] We are assuming that this snowball doesn't even have the shape of a tiger.
[168] Putnam (1975): p. 248.
[169] Some of you may think that Putnam's example is too far fetched. People simply don't really talk like that. During one of the primaries I came across an interesting response by one of the republican candidates, Governor Mitt Romney, to the following question: "Knowing everything you know right now, was it a mistake for us to invade Iraq?" He said:

> Well, the question is, kind of, a *non sequitur*, if you will. What I mean by that – or a *null set* – that is that if you're saying let's turn back the clock and Saddam Hussein had…[emphasis added].

This is from a republican presidential candidate debate aired on television June 5th 2007: http://myclob.pbwiki.com/Was+it+a+mistake+for+us+to+invade+Iraq

produced by *another* individual, which results in the production of *t'*, *suffices* for tokens *t* and *t'* to represent performances of the same word (that is for *t* and *t'* to be nodes on the same repetition tree) – as argued by Kaplan.

His proposal *seems* to work particularly well in contexts where one speaker enters the conversation with a "blank slate", or where one speaker is under- or misinformed about a topic and wants to learn from the more knowledgeable interlocutors.[170] It is remarkable how in many instances we appear to pick up a word *instantly* yet without knowing anything at all about its referent.[171]

> [I]n a tiny fraction of cases the [epistemic] connection is absent – semantics (or metasemantics) does not require it – and in these cases we have direct reference, and expressibility, but no apprehension.[172]
> A name may later take on the required epistemic connection when the referent appears upon the scene and is recognized as the named object.[173]

What Kaplan describes seems to be particularly true for "what-is-x-questions". The inquirer is only put *en rapport* with the referent *after* her question has been answered.

Capoeira is a Brazilian martial art that is always accompanied by songs. One of them begins 'Não é dinheiro, não é ouro…'. Although I don't speak Portuguese, except for a few words, I can ask "what does '*dinheiro*' mean?"[174] This would be easy to explain if the lyrics had been written on a board and I had pointed at the word in question and asked "what does this word mean?" But I didn't. I produced my own token of the Portuguese word 'dinheiro'. I repeated more than just some word-"form" that sounds like the Portuguese 'dinheiro'. I want to

---

[170] Ebbs (2002).
[171] We can do this even for words of a language we are not competent in. This is why I do *not* agree with Alward (2005) that Kaplan's proposal could be improved by requiring that the speaker must be overall competent in the language that the word *w* belongs to – in order to be in a position to repeat *w*. This is not wherein the mistake lies.
[172] Kaplan (1989).
[173] Kaplan (1989): footnote 98.
[174] I happen to know how to pronounce that word correctly, since I have heard the song many times. Even if I had mispronounced it, I would still have repeated *that* word.

know what the Portuguese 'dinheiro' means. I don't want to know about some word 'dinheiro' or other, I am asking about *that* word – the one used in the song in question. Suppose the song were in French, where the word denoting money can also mean silver. 'Ce n'est pas de l'argent, ce n'est pas de l'or…'. This time, it's a language I can speak. So I ask: 'Qu'est-ce que c'est du argent' (overlooking that it would be rather strange for a competent French speaker not to know what 'argent' means)[175]. Several aspects of this example seem to favor my being able to *use* the name 'argent' versus merely being in a position to mention it:[176]

a) My question is grammatically complete. So I must be using *some* word (albeit not necessarily the public word 'argent').
b) Other speakers would and ought to interpret me as asking a question about *argent*.
c) In the song the word 'argent' denotes money, so the proper answer to my question would be that money is a means of exchange that can take on a variety of forms (expressed in French of course).

I will later go on to use the French word 'argent' on many occasions successfully in conversations with other competent French speakers. I intra-subjectively individuate the various 'argent'-marks as utterances of one and the same word 'argent_N'. I view myself as having repeated the same word 'argent' (in the thick sense) – since I first produced the corresponding mark. This speaks in favor of Kaplan's notion of words that regards the individual conceptions a

[175] From the grammatical construction of the song excerpt I can tell that 'argent' is a mass noun, enabling me to ask a grammatical question without knowing at all what '*argent*' denotes.
[176] Most ordinary speakers don't pay much attention to the use-mention distinction, so how can we be sure that I am using the word 'argent', rather than mentioning it? Most ordinary speakers don't indicate syntactically whether the word is being mentioned, rather than used. In this case, it just happens that 'argent' is a mass-term while words are the kinds of things we can count. If I were asking about the word 'argent', then I should be using the count-noun construction: 'Qu'est-ce que c'est le (mot) 'argent''. I am not even sure how we would answer such a question.

speaker associates with her own words as irrelevant in interpreting her words (with the omission of the exception clause that I will discuss later).[177,178,179,180]

In the above example I had no substantial beliefs associated with '*argent*', but I was nevertheless able to mention and use it – or so it seems. Sometimes, though, the indexical intention – repeating *that* word (the third word of the original song, or the seventh word in the French translation) of *that* particular song – can be overridden.

I now want to call into doubt whether we really pick up the word with its pre-packaged semantic content as Kaplan has us believe. The following example speaks against immediate word-acquisition. If my analysis is correct, then Kaplan owes us an explanation with regard to

---

[177] Maybe you care to also inform me that there is another word 'argent' in French that denotes silver. Accounting for ambiguity is somewhat difficult for Kaplan, since he regards words such as 'bank' denoting financial institutions and 'bank' denoting the land mass adjacent to a river as distinct words (not as one and the same word with distinct uses or meanings). For Kaplan ambiguity can only be a relation between marks (or perhaps a *non*-resemblance notion of word *forms*).

[178] Evans makes a similar observation: "An example which might favourably dispose one towards the [causal] theory [of names] is this. A group of people are having a conversation in a pub, about a certain Louis of whom *S* has never heard before. *S* become interested and asks: 'What did Louis do then?' There seems to be no question but that *S* denotes a particular man and asks about him. Or on some subsequent occasion *S* may use the name to offer some new thought to one of the participants: 'Louis was quite right to do that.' Again he clearly denotes whoever was the subject of conversation in the pub" (Evans (1973, 1985) pp.6-7).

[179] Bach declares boldly that we pick up the referents of proper names via the name's *form*: "Since the hearer's mental token of the name 'inherits' the same object as the speaker's, the object of the hearer's thought is determined relationally, not satisfactionally. […] A token of a name can function as a *de re* mode of presentation because its reference is determined not by its meaning but by its ancestry. It plays this role by being of a certain form (sound or shape), generally the same as the one to which it is linked. When the meaning of the token does not matter, its referent cannot be determined satisfactionally. Indeed, that only its form matters is what constitutes its being used as a [proper] name. And that is what enables one to form *de re* thoughts about an unfamiliar object referred to by that name. Since the token of a [proper] name represents in virtue of its form, not its meaning, representational features can be perceived by the hearer, who can then and thereafter use mental tokens of the same name to think of (or refer to) the same object" (Bach (1987) pp. 32-33). Of course Kaplan would probably quarrel with the idea of a name *form*. Perhaps Kaplan could agree to a notion of two tokens displaying the same form, if they didn't do so based on a resemblance relation – perhaps we can say that the hearer intending to repeat token *t* and thereby produces *t'* results in *t* and *t'* constituting the same word or name *form*. But as I will explain towards the end of this chapter, this is not what Kaplan set himself up to do. He believes the speaker is thus repeating the *word* not merely the word *form.*

[180] I will discuss cases of this sort in greater detail in chapter 4.

what constitutes an appropriate intention to repeat a word-token *t* such that the resulting performance *t'* is a node of the same repetition tree as *t*.

Suppose a person, Jane, finds a piece of paper with only the mark 'Pompey' written on it. Let's stipulate further that the original producer of that mark meant to write the name denoting the famous Roman statesman known as 'Pompey$_P$'.[181] Jane has never before encountered this word. She subsequently wonders: 'What is Pompey$_J$?'.[182] The last token occurring in her question is intended as a repetition of the name she found on the piece of paper. Kaplan would have to say that Jane succeeded in repeating the public name 'Pompey$_P$' – denoting the Roman statesman Pompey.[183] So, is Jane's question inquiring about Pompey? Jane intends to find out about the name or word the original producer of the token wrote on the piece of paper – which would entail finding out about what it denotes. She would therefore regard information about that very token, as used by the person who produced it, as relevant. These details accord with Kaplan's theory and similar considerations may very well have inspired it.

Still, my inclination is to say that Jane is not *at this point* capable to use the name 'Pompey$_J$' to ask a question about Pompey. For one thing, why should she even assume that the name denotes an individual?[184] At this point Jane's word is a blank slate; she doesn't associate any substantial conceptions with it. Although she knows that the name is of some interest to her fellow speakers. Even though we allow speakers to be under-informed about a given topic, we

[181] I am using the subscript 'P' to designate that it is the public name that was produced on the piece of paper. The mark is a node on the public repetition tree: 'Pompey'.
[182] I am using the subscript 'J' to designate that the mark is one produced by Jane and that it is a word of her own idiolect. Whether 'Pompey$_P$' and 'Pompey$_J$' are one and the same word – in the thick sense – remains to be seen.
[183] The exception clause that I will discuss later doesn't help Kaplan here, since Jane doesn't have any conceptions regarding the name 'Pompey' and thus cannot have any conceptions that conflict with how the producer of the original mark uses 'Pompey'.
[184] Thanks to Ebbs for bringing this to my attention.

expect speakers to know *some* essential truths about what a given public word denotes in order to be entitled to participate in the given practice.[185]

Consider what would happen if Jane were to receive the following response by a person other than the producer of the original mark: 'Pompey was a historic Roman town-city'. Something isn't quite right, but I am sure that I don't want to argue that Jane was given false information about Pompey. Pompeii was a historic Roman town-city (destroyed by an eruption of Mount Vesuvius); Pompey was a prominent military and political figure in Rome.[186] While the standard pronunciations of 'Pompeii_P' and 'Pompey_P' are distinct but similar, one can easily see how Jane may attach the (standard) spelling or pronunciation of 'Pompeii_P' to 'Pompey_P' and the same goes for the person answering her question. Remember that even incorrectly spelled nodes can be causally linked to the repetition tree of a given word (i.e. 'Pompeii_P').

Interpreting Jane's question as one pertaining to Pompey, leads to unacceptable inferences. If we leave her intra-subjective word-individuations intact, she comes to believe that Pompey (the man) was a city, which is *absurd* (and renders anti-individualism vulnerable to serious self-knowledge problems). Imagine Jane were to tell us: "I wonder what life in Pompey was like prior to the eruption. I wish I could have experience life there." If we were to interpret her name 'Pompey_J' as denoting the same individual as our name 'Pompey_P' – Jane would be wishing to have lived in a person! We simply cannot make sense of such a desire. What states of affairs would have to be actual for her wish to come true? We find ourselves at a loss of how to begin answering this question.

[185] Recall Putnam's example, where a person points at a snowball and exclaims: "This is a tiger!"
[186] I intentionally picked two words that are not considered homonymic in anticipation of the objection that the case I consider is merely a problem of ambiguity.

Taking for granted that a particular token cannot constitute a node on two distinct repetition trees, we could, nevertheless, argue as follows: even though Jane believes to be using the same word '$Pompey_J$' (in the thick sense) over time, her various '$Pompey_J$' performances stand for different words – without her noticing that they do. As a result, it would be possible to interpret some of her '$Pompey_J$' tokens as repetitions of the word '$Pompey_P$' (i.e the question "What is Pompey?"), while other '$Pompey_J$' tokens constitute (misspelled) repetitions of the word '$Pompeii_P$' (i.e. her assertion "Pompey is a historic Roman town-city"). This would *clash* with Jane's intentions to keep repeating the *same* word with the tokens in question. Can we make sense of Jane's inability to keep track of her own word-repetitions, as described?

So far, I relied heavily on speakers being right about individuating their own words across time. Otherwise, it becomes *difficult* to explain why I am able to state satisfaction conditions for my own words, such as when I utter:

"*x* satisfies '<u>dog</u>' if and only if *x* is a <u>dog</u>."

The same ability is ascribed to other speakers. I do not want to give up this assumption of *infallible* individuation of one's own words. Kaplan's proposal assumes that individuals can keep track of their own words – otherwise, what use is the ability to repeat other speaker's words, if I cannot store testimony thus conveyed and use it as an inference in an argument at a later time? Given that a speaker has privileged access to her own conceptions associated with a given token, it seems puzzling at best why we should generally be capable to repeat other people's words (in the thick sense) if we cannot do so for our own words. For those reasons I argue that we must ascribe to Jane, as well as other speakers, the ability to successfully repeat

their own words, even if this means that Jane deviated from the accepted performance standards of how to pronounce 'Pompeii_P'.[187]

We interpret Jane (in retrospect) as having inquired about Pompeii, because we accept that her later beliefs as pertaining to Pompeii. Assuming this to be correct and given that Jane is assumed *infallible* at individuating her own words it follows that Jane simply used a deviant (but still recognizable) spelling or pronunciation for our word 'Pompeii_P': Jane's word 'Pompey_J' is our word 'Pompeii_P'. That means the content of Jane's question *q* was determined by how the person responding to *q* interpreted *q* and by the resulting beliefs Jane comes to acquire as a result about the referent of 'Pompey_J'.

If at some future point in time somebody alerts her that there are two similar words: one 'Pompeii_P' denoting a Roman town-city, the other one 'Pompey_P' denoting a prominent Roman statesman, Jane would revise her spelling of the word 'Pompeii_P' (from 'Pompey_J' to 'Pompeii_J'). She would not take herself to have confused Pompeii the town with Pompey the man. In this case Jane would change the spelling of her word 'Pompeii_J' from 'Pompey_J' to 'Pompeii_J' such as to adhere better with the performance standards of her linguistic community – these distinctly spelled tokens would nevertheless belong to the same repetition tree.

Recall, however, how I also remarked on Jane's inquiry being guided by the desire to know the meaning of the word written on the piece of paper. She would therefore regard information about that very token, as used by the person who produced it, as relevant. If Jane never gets the chance to find out what name or word the person intended to write on that piece of paper, then her name 'Pompey_J' clearly denotes the city Pompeii, not the man Pompey. What would

[187] I will explain in a later chapter that when Jane picked up the token 'Pompey' she created a mental file or register labeled 'Pompey'. At this point the file is empty and Jane cannot use this token t_1 to convey beliefs about the referent connected tot his file (yet). Gradually she will acquire information that gets stored in this file (i.e. Pompey_J ("our" Pompeii) is a Roman town-city…). Eventually we can view Jane as having acquired a *new* word 'Pompey_J' (*not* to be translated as out 'Pompey').

happen, if after years of using 'Pompey_J' as the name of the city, she were to discover that the name written on that original piece of paper is the name of a person and was written to denote Pompey? Jane would judge that her own word 'Pompey_J' wasn't an appropriate repetition of the original. I don't view her discovery that the original mark denotes Pompey, a Roman politician, as sufficient evidence that her word 'Pompey_J' thus (also) denotes Pompey.[188]

In this example Jane merely seems to have acquired the *mark* 'Pompey' but not the word 'Pompey_P' that denotes Pompey the person. Even though this very mark was produced with the intention to stand for the word 'Pompey_P' (denoting Pompey) and Jane meant to repeat that very word, she *failed* in her attempt. Whether Jane's "what-is-question" is about Pompey, when she asks: "What is Pompey_J?" depends on how she goes on to fill her conceptions regarding that mark. If those conceptions turn out to *radically* clash with the word that the mark she originally encountered stood for (a person is not a city or *vice versa*), then the causal link to the repetition tree may *fail* to be established.[189]

---

[188] In chapter six I propose to solve slow-switching Twin Earth puzzles, where the speaker appears to draw invalid inferences, with merged registers. The speaker stored information about two natural kinds for which two distinct public words exist in one register. I do not view this a plausible option in Jane's case. The objects thus merged as one kind in a single register must at least appear similar to the individual. This is not the case here. She is not confusing the man Pompey and the city Pompeii as one kind or individual. I therefore do not ascribe to Jane a register that merges Pompey and Pompeii.

[189] Kaplan conceives of the possibility of a *new* name being created: "One might consider two kinds of polar cases: In one case you intend to use (to *repeat*) a given common currency name with whatever referent it may have. ("What is Hesperus?" you ask, overhearing a conversation in which the name is used.) In the antipodal case, you intend to dub a particular thing using an apt generic name. In the formed case there is continuity, in the latter, creativity, a new name is created". He even admits of there being a possibility for one intending to *repeat* and already existing name, while one also attaches descriptive beliefs with the token produced with the intention to repeat that name which *conflict* with the success of repetition. "But there are those troubling cases (first thrust upon our consciousness by Keith Donnellan, and then Gricefully reconceptualized by Saul Kripke) that seem to lie in between: the man with the Martini, the false introduction, and their ilk" (Kaplan (1990): p.117) . Just imagine a case where your neighbor points at a woman across the street carrying several books and tells you that her name is 'Mrs. Wagner'. As it happens this woman, let's call her 'Virginia', is Mrs. Husting. Mrs. Wagner is a different woman, who also goes by 'Ursula' and lives several blocks down the street. Your neighbor knows that Mrs. Husting goes by the name 'Mrs. Husting', but he has poor vision and thought he saw Ursula when in fact he saw Virginia. You recall the event: "Mrs. Wagner was carrying an awful lot of

This means I must qualify my earlier conclusions about the '*dinheiro*' and '*argent*' example. Although I appeared instantly capable of using those words, mere exposure to the mark doesn't enable me to use the respective word that the mark was produced to stand for. It matters how I "fill" in my beliefs about what the word that the mark instantiates denotes. Those conceptions in turn are formed on the basis of information passed on to me by other speakers and their interpretation of the mark in question. Individual conceptions regarding a word cannot be entirely set aside when we individuate words.

Perhaps we can account for Jane's mistake as one trying to repeat two distinct words, but she is under the impression that they are one and the same. When Jane forms the belief: "Pompey is a Roman town-city" her mark 'Pompey', she thereby intends to repeat the word originally produced on the piece of paper, as well as, the word her informant was using. Mercier explains how this fails to be a well-formed intention on Jane's part: "Now, in order for my intention to use a word with *its* meaning to be a well-formed intention, it has to be the case that *the* word about which I have that intention *is* indeed *one* word."[190]

But when Jane first asks the question "What is Pompey?" she has not yet encountered her informant's word 'Pompey' (which is to be identified with our public word 'Pompeii_P') and can thus *not yet* have the intention to repeat his word(s). So, only some of her later 'Pompey_J' tokens could be linked to two distinct repetition trees ('Pompey_P' and 'Pompeii_P'). I don't think we can make sense of Jane's word 'Pompey_J' denoting a person *and* a city – especially if she (by then) intends to repeat a proper name.

---

books" (Kaplan (1990): p. 117). Since you have never met Ursula, it is natural to assume that your memory pertains to Virginia. Thus your name 'Mrs. Wagner' denotes Mrs. Husting, even though your neighbor's name 'Mrs. Wagner' denotes Mrs. Wagner (Ursula). Your intention to denote the woman across the street (Virginia) conflicts with your intention to repeat the name of your neighbor 'Mrs. Wagner' that denotes Ursula.

[190] Mercier, Adèle (1993): p. 79.

I have shown Jane to have failed in her attempt to repeat another person's word. Jane is still committed to repeating *some* public practice (she still remains a language *consumer*); she happened to *switch* the standard "labels". Her mark "Pompey" is linked to the public name 'Pompeii' (not the public name 'Pompey' – at least until other speakers correct her spelling and pronunciation). Should we really worry about such confusions if they only constitute a problem for ambiguous or relevantly similar word pairs? This is not an excuse Kaplan can help himself to, since he believes that "the difference in phonographology, the difference in sound or shape or spelling, can be just about as great as you would like it to be."[191]

As I clarified in my discussion of Cappellen's objections against Kaplan *not everything goes*. We would have a hard time to make the case that 'alligator' stands for 'onomatopoeia'. It is hard to imagine a person in good faith intending to repeat 'alligator' and produce 'onomatopoeia' instead. It is, on the other hand, possible to conceive of someone intending to repeat 'Pompeii' and produce 'Pompey' instead – the same goes for the word-pair 'alligator' and 'allegory'.[192] Anybody who has learned a foreign language will attest to this, even when we are already overall competent in that language we may still attach very non-standard marks to words of that language.

I have demonstrated how Kaplan's intention-based account doesn't yet provide us conditions for when two *inter-subjective* word tokens *t* and *t'* are of the same word type. Given that speakers may *fail* in their attempt to repeat a token (produced by another performer) that belongs to a public repetition tree, we need an account of what distinguishes successful from failed repetition attempts.

[191] Kaplan (1990): p. 101.
[192] And this *not* being a slip of the tongue or pen!

Perhaps Kaplan could be interpreted as offering an alternative criterion for the first step (i) of the word individuation process (U).[193] However, this is not what Kaplan describes himself as doing. Kaplan clarifies:

> There are two phonographic words "base" (meaning '*low*' and *bottom*), not, as the orthographic conception would have it, a single word with two meanings.[194]
> [W]e are for the most part, language *consumers*. Words come to us prepacked with a semantic value. If we are to use *those* words, the words we have received, the words of our linguistic community, then we must defer to *their* meaning.[195]

Kaplan attempts to solve word-ambiguity. There are two words '$bank_1$' (the financial institution) and '$bank_2$' (the slope at a river). They happen to be spelled and pronounced the *same* way and both represent nouns. Ambiguity thus means that one and the same mark can, depending on context, either constitute a *standard* replica of '$bank_1$' or '$bank_2$'. It is not the case that there is one word with distinct meanings or uses. Therefore, Kaplan distinguishes between the intention of repeating '$bank_1$' versus '$bank_2$', thereby rendering word-repetition *semantically* loaded. Eventually, we are interested in more than labels, we want to know what beliefs our interlocutor is trying to convey to us!

3.5 Thought examples demonstrating repetition failures

Receiving repeated significant misinformation about how a mark *m* is used in a given linguistic community *L* may inhibit a speaker's ability to communicate with members of *L* by using *m*. In the following two thought-examples I will show that there are also other sources for

---

[193] Ebbs's rejection of the two-step word individuation process:

(U) Two word tokens t and t' are of the same word type if and only if
  (i) t and t' are each tokens of the same orthographic or phonetic type, and
  (ii) facts about the ex-use of t and t' determine that they each have semantic values and that their semantic values are the same.

Ebbs (2009): p.112; Ebbs rejects the token and ex-use model of words (U).

[194] Kaplan (1990): p. 100.

[195] Kaplan (1989): p. 602.

miscommunication. Sometimes a given mark *m* stands for different words *w* and *w'* in different languages, such that *w* and *w'* ought *not* to be translated for one another – this causes the speaker's failure to repeat *w*. In the second example, either through radical underinformation or some misunderstanding, the speaker fails to establish appropriate linguistic connections to a public word practice.

Allow me to recall an experience I once had at the dentist. I was asked to pick from a long list of medical conditions those that I am suffering from at present or had been diagnosed with in the past. I affirmatively placed a checkmark next to 'angina'. In German – my mother tongue – '*Angina*' is ordinarily used to refer to strep throat also known as '*angina tonsillaris*'. English speakers, on the other hand, tend to ordinarily refer to *angina pectoris* with 'angina', whereas Germans would commonly name this condition '*Brustenge*' (chest constriction) or '*Stenokardie*'. Since chest constriction and angina (*tonsillaris*) are clearly distinct medical conditions, the ordinary German word '*Angina*' is *not* to be translated as the ordinary English word 'angina'. Even after many years of immersion in English speaking communities, I hadn't learned the (ordinary) English word 'angina' yet. Perhaps you can already see where this is going. I clearly meant to convey that I have had strep throat in the past. Admittedly, at the time, I was wondering why this would be of any concern to a dentist. Then again, most of those medical forms seem overly elaborate, so I didn't continue puzzling over the matter. The dentist, on the other hand, was interested in whether I am suffering from *angina pectoris*. When going through my papers the technician asked further questions about my suffering from "angina". Since she showed a significant degree of concern I suspected some miscommunication and clarified my response. In the end, I was able to convey what I had intended to by paraphrasing.

The initial misunderstanding could very well have gone unnoticed. Time constraints often render us less inquisitive than we should be.[196] I thought that that word 'angina' was to be translated as 'Angina' into German – which is *false*. I was using the word 'angina' to mean *angina tonsillaris*, while the technician used 'angina' to mean *angina pectoris*.

What is so difficult to explain in this example is that I certainly meant to use the same word 'angina' as the one printed on the medical survey;[197] however, I had specific expectations as to what kind of word it is, unlike in the case of '*dinheiro*'. Whether I ought to be considered having had a false belief about angina (pectoris) or having used a *different* word 'angina (*tonsillaris*)' depends on how *I* go on to use repetitions of my token 'angina'.

There are several ways of how I could interpret my own mistake. Do I say: "Oh I thought angina causes pain in the throat, not in the chest." Or do I instead go on to reflect: "The technician is using a different word 'angina' from mine. We are talking past one another." Given my translation into ordinary German the latter is the more plausible interpretation of my word. It matters very much which beliefs are "dear" to me. What I cannot do is view my attempt to repeat the technician's word 'angina$_T$' as successful *and* hold on to the translation that my "English" word 'angina' is to be translated as my German word 'Angina'.[198] If I want to maintain that the belief I meant to convey was that at some point in my life I have suffered from strep throat, then I must – if I am committed to successfully communicate with the technician – add a *new* word 'angina' to my idiolect.

---

[196] Verdonik (2010).

[197] I actually re-used an already existing mark – originally intended to stand for the English word 'angina'. When an individual re-uses an already existing mark *m*, *m* does not thereby necessarily retain its previous denotation (and its membership in a specific repetition tree). This case is similar to the one discussed earlier with the found mark 'Can you spare a quarter' that acquired word-hood or semantic content through the communicative intentions the speaker (Cappellen) associated with that mark.

[198] A legitimate worry is that we generally don't reflect on our words this carefully and thus may miss such inconsistencies.

Perhaps we can brush this example aside as one where the speaker simply lacks a well-formed intention, since the utterances she intends to repeat and views as performances of the same word are actually performances of two distinct words. I intended to repeat the word 'angina' written on the questionnaire, but at the same time also thereby thought I was repeating a word to be translated into German as 'Angina'. In anticipation of such examples Kaplan writes:

> The second kind of error is a short circuit – two different circuits got wired together – zap, the whole thing goes up in smoke. I am inclined to think that when two different common currency words are wired together in this way in a given black box, which then pulls from that common source and transmits, nothing whatsoever is being said. Is it transmitting the first word? Is it transmitting the second word? I think there is just no answer to that question.[199]

Even though I saw myself as committed to two distinct public repetition trees (that I *falsely* interpreted as translations of one another) I view this (in retrospect) as a case where my linguistic commitment to one was so overwhelming that it *inhibited* my commitment to the other repetition tree. Kaplan correctly identifies the phenomenon, but I disagree with his conclusion, one can clearly make sense of my thought that I would express in English thus: "I have had angina as a child." It states that I have had strep throat as a child. I had attempted to repeat an English word, but *failed* to do so. Nevertheless, my utterance was a repetition of *some* public word.

In the next thought example, I demonstrate how sometimes speakers intend to repeat some public word *w*, but the resulting token marks the beginning of an entirely *new* repetition tree (one that doesn't exist yet in any language). Evans presents such an example, which we now know to be historically incorrect, but no less philosophically illuminating. He writes:

> We learn from Issac Taylor's *Names and their History* (1898):
> In the case of 'Madagascar' a hearsay report of Malay or Arab sailors misunderstood by Marco Polo... has had the effect of transferring a corrupt form of the name of a portion of the African mainland to the great African Island.[200,201]

[199] Kaplan (1990): p.109.
[200] Evans (1973, 1985): p. 11.

Although Marco Polo meant to repeat the word of his Malay or Arab encounters, we must conclude that he failed to do so. Otherwise, unbeknownst to us we are referring to some part of mainland Africa and not to the island we believe to be talking about, when uttering 'Madagascar'. Since the referential chain reaching back from our tokens to those of Marco Polo are assumed intact, he must have somehow failed to pick up the word of the foreign sailors. This again leads us to the conclusion that the intention to repeat a word doesn't guarantee our success in doing so. Whereas in my previous 'angina' thought example my word 'angina' was able to connect to some public word, the German word 'Angina', Marco Polo's word 'Madagascar' had no practice to "latch on" to – except for the one that would have made his name denote some part of mainland Africa.

What allowed Marco Polo's token 'Madagascar' to break off the Malay and Arab public repetition tree 'Madagascar'? How was it possible for a new word-chain to be initiated? Mercier points out how this development is most often *not* driven by the more competent speakers or experts[202] (i.e. the eventual creation of a new word chain for 'livid' denoting the

---

[201] From an interview with Lepore about a paper in progress: "Here's an example that John [Hawthorne] and I [Lepore] put in the paper ["On Words,"]. Most people would agree that performances of the word 'moan' and the word 'mean' are totally separate words. I can't imagine anyone denying it. However, if you take the cluster of performances of the one word and the cluster of performances of the other word, you could trace them back to a single cluster of performances in Saxon England. Now, from that, I don't think anybody would feel compelled to say, 'Oh, isn't that interesting: in Saxon England they had a homonym that was two words.' I think most people would agree that whatever they were doing in Saxon England, those performances were performances of one word. So that raises the interesting question, since those performances of that single word somehow or other provoked and inspired performances of distinct words, contemporaneously, when did these two words come into existence, exactly?" (Baggini, Julian (2009) "Words don't come easy" in *The Philosopher's Magazine*, 43; http://www.philosophypress.co.uk/?p=67).

[202] Mercier (1993): pp. 85-86. I cannot agree with Mercier that those new word creations are based on mistakes. We don't want to say that people are misusing the (original) word 'livid' to denote the facial color red or pink. This would not get the point across, since then we would admit that those people indeed are repeating the original word 'livid'. Mercier and I want to argue that they are using a *different* word 'livid'. I also don't want to say that the use they attach with the mark 'livid' ceases to be a mistake once dictionaries condone it. This would entail that prior to the new dictionary entry '$livid_R$' (meaning the color of a reddish face) some people were *mis*-using the original word '$livid_{BG}$' and post dictionary

opposite color of a pale face; a new expression 'begging the question' recently entered the lexicon, stating that a question ought to be raised[203]). Sometimes, speakers view themselves as language consumers, but due to some misunderstanding, they unintentionally break off the existing linguistic practice and give birth to a *new* word. Marco Polo didn't mean to impose his geographic expertise or his social status on the Arab (or Malay) sailors when he named some island 'Madagascar'. Clearly, it isn't up to just anybody to name an island and have that name enter public usage. Still, Marco Polo merely intended to use the *same* name as his encounter, but failed in his attempt.

3.6 An exception clause threatens to undermine anti-individualism altogether

Kaplan is not oblivious to the kinds of examples I discussed above and he offers a solution of his own. I will show that the exception clause Kaplan proposes is in need of being contained, or it could threaten to undermine standard (social) anti-individualist thought examples.

Kaplan suggests that a word may undergo a change in reference.[204]

---

amendment they are *suddenly* employing a *new* word 'livid$_R$' *correctly*. Marco Polo's mistake can only consist of his attempt to repeat the word of the locals, when he is really giving birth to a *new* word.

[203] There is a (similar) older word 'livid' meaning bluish gray; there is also a (similar) older expression 'begging the question' denoting a circular argument. I suppose if somebody begs the question there is indeed a question that begs to be asked, namely: how does she support her conclusion?

[204] Perhaps the word 'livid' is such an example. Apparently many people are under the misconception that it denotes the color of a reddish face, whereas it actually denotes a bluish gray color. George imagines a person checking what 'livid' means in the dictionary and thereby *revising the meaning* of her word 'livid' (George (1990): p. 289).
As I will lay out later, I won't accept that a person can be so radically wrong in interpreting her own word. George's analysis of the case is along the lines of Kaplan's notion of word. The individual intended to repeat the *public* word 'livid' and only later found out what it means. Prior to consulting the dictionary, she said of Peter's face: "Peter's face is livid". Suppose Peter's face was red. I simply cannot make sense of the claim that her utterance states a false belief (especially since she believes that 'a livid face' describes a red face).
For reasons I will describe later, I would say that when she looks up the word 'livid' in the dictionary, she acquires a *new* word 'livid'.

> [Can] a common currency *name* undergo a change in referent? […] It is a matter for further analysis to say whether such an entity could change meaning (or reference). It is certainly no part of my conception that it cannot. [205]

We can see that he hesitates to endorse this possibility. And so he should! Much of what attracted us to the common currency notion of words depends on the interlocutors talking about the same topic (which pre-supposes that the reference remains constant[206]) – the existence of trans-theoretic terms is evidence for this![207]

He proceeds by suggesting the following strategy:

> Try to show that something like Donnellan's *referential use* is involved in cases of apparent change of referent of a given common currency name. Then try to show that this sort of referential use is sufficiently like creating a new common currency name from the genus of a given name so that by the time a new 'semantic referent' appears, a new name does also. I don't know that such a strategy would be successful.[208]

Kripke suggests in a similar vein the following solution to Evans's Madagascar thought-example:

> In all cases [where a reference shift occurs], a present intention to refer to a given entity (or to refer fictionally) overrides the original intention to preserve reference in the historical chain of transmission.[209]

According to Donnellan,

> [a] speaker who uses a definite description referentially in an assertion [...] uses the description to enable his audience to pick out whom or what he is talking about and states something about that person or thing.[210]

It is not quite clear how we should take this to apply to words. In Donnellan's example a person *P* is being mistaken for drinking a Martini (perhaps he is only holding a Martini glass which

---

What about the fact that linguists consider '$\text{livid}_{\text{RED}}$' and '$\text{livid}_{\text{BLUISH-GRAY}}$' to be the same word from an etymological point of view? Again, this is not a problem for my view, since I am after an epistemic notion of word-sameness. It would be nice if all our pre-theoretic uses of the term 'word' could be captured in one theory, but this is unlikely.

[205] Kaplan (1990): p.118.

[206] I am thinking of the reference of our word 'gold' remaining constant throughout changes in theories pertaining to the nature of gold over time.

[207] Putnam (1973, 1975); Putnam (1974, 1975); Putnam (1975, 1975).

[208] Kaplan (1990): p.118.

[209] Kripke (1972): p.163.

[210] Donnellan (1966): p.267.

happens to be filled with water).[211] The assertion "The man who is drinking a Martini is handsome", according to Donnellan, may still be accurately interpreted to pertain to *P* even if *P* is *not* (at the time the conversation takes place) drinking a Martini. Now suppose that the speaker instead uttered: "James Bond is handsome." She meant for her audience to connect that James Bond's trademark drink is the Martini and that she thus meant to denote *P* with the fictional name 'James Bond' (the man who appears to drink a Martini but actually is drinking water, unbeknownst to the people talking about him).

It is not entirely clear what Kaplan had in mind. When a description is asserted referentially, according to Donnellan, the speaker intends for his expression 'the Martini drinker' to keep referring to *P* (within a given context) even if the description is *false* of *P* (at time *t*). The intention to refer to *that person P* overrides the attempt to make a true assertion (namely that *P* is drinking a Martini).

I am not sure how this is supposed to work when we are trying to explain why a speaker *violates* public practice. Marco Polo's intention to refer to *that island* (the one he presently stands on) *overrides* his intention to repeat his Arab encounter's word. Donnellan contrasts the referential use with the attributive use and argues that a "speaker who uses a definite description attributively in an assertion states something about whoever or whatever is the so-and-so."[212]

Perhaps the parallel to the attributive use would be "Whatever place my interlocutor denotes with 'Madagascar' is what I will call 'Madagascar'." A more appropriate name here would be 'deferential use of expressions'.

[211] Even though the definite description may not uniquely pick out anybody, it could still referentially pick out one salient person or object as the one being spoken of.
[212] Donnellan (1966): p.267.

According to Donnellan it is the speaker who has a sense of whether she wants her description to pick out the object she is attributing a certain quality to referentially or attributively. I think this is an insight we can exploit for our present explanatory goals: when does a speaker fail in her attempt to repeat a given word? When Marco Polo failed to repeat his Arab or Malay interlocutor's word 'Madagascar' he didn't intend to *go against* his status as a language *consumer*. Just like the speaker whom Donnellan calls to our attention didn't mean to assert a description that failed to denote. It is once those speakers get faced with their failure that they hold on to their referential intentions – "regardless of whether *this man* is drinking a Martini, *he* is the one I am talking about." Marco Polo may say (after he has used 'Madagascar' in a many conversations) "regardless of what my Arab or Malay encounter meant with 'Madagascar' I mean *this island* when I say 'Madagascar'." What seems to be clearly different between the distinctions attributive versus referential and deferential versus referential is that there is no problem with using a description *D* in some contexts attributively, but referentially at others. It seems to me that Marco Polo must commit to one use *once and for all* or else his proper name 'Madagascar' would sometimes denote a geographic area on mainland Africa and at other times it denotes an island – which goes against our epistemic goals.

Wouldn't the individualist here intercept and argue that this is what she has been arguing all along? Words are being used referentially and not deferentially. We need a way of containing the referential use. A convincing story that tells us how the speaker commits herself to a public language practice in such a way that she is willing to violate her referential intentions is needed.

Couldn't we say that Alice had the referential intention to have 'arthritis' denote the condition that causes pain in some thighs? I certainly had the referential intention to have 'angina' denote the condition that caused the pain in my throat (many years ago). While we want Kaplan and

Kripke's exception clause to operate in the angina thought example, we do *not* want it to come into effect in the arthritis thought example – or else anti-individualism collapses into individualism!

The idea that my referential intentions were *stronger* than Alice's does not create a robust enough barrier between the two antagonistic theories. Sometimes strongly held beliefs are challenged and get revised, for instance, the belief that space is Euclidean, or the belief that all adult swans are white. [213]

My belief that all adult swans are white may be so strong that I am unwilling to believe the Australian ornithologist that this black bird in front of me is a swan (even under the assumption that she is sincere). At that moment I may even decide that – unlikely as it is – she simply is wrong about that bird. Ludicrous, I think to myself a *black* swan? I view myself in *disagreement* with the relevant expert. But then I take the time to do some research on the matter and sure enough, it turns out there are indeed black swans in Australia. Only after consultation of a second source am I willing to abandon the belief in the truth of which I formerly was so certain. It seems that over time even powerful referential intentions are open to revision.

Getting back to Alice the arthritis patient, what seems to play a crucial role is that Alice didn't mean to have the belief "arthritis is the kind of condition that occurs in (thigh) bone tissue" *override* her intention to repeat her doctor's word 'arthritis'. She is prepared to defer to the expert even if that means that some of her referential intentions may turn out to be false.[214]

[213] There are black adult swans in Australia.
[214] This may change if the doctor were to point at a snowball and exclaim: "This is arthritis!". Even Alice's deferential intentions have their limits.

3.7 Refining amendments

Burge and Putnam acknowledge that speakers can fail in their attempts to repeat a public word – mere exposure to a mark is not sufficient to acquire the word the respective mark stands for. In this section I will discuss under what conditions Putnam and Burge view the identity between public and idiolectal words to be missing. Given that they are, like Kaplan, social anti-individualists, I hope to lean on Putnam and Burge's respective analyses to refine Kaplan's notion of words such as to account for the repetition failures I mentioned above. Something that puts the speaker adequately *en rapport*[215] with the referent of the word seems to be lacking. How can we establish this epistemic link between speaker and referent without thereby regressing back to individualism?

In order to (roughly) determine whether a speaker can be considered to partake in a public word practice, Putnam suggests a "passing muster" clause:

> Extension may be determined socially, in many cases, but we don't assign the standard extension to the token of a word W uttered by Jones *no matter how* Jones uses W. Jones has to have some particular ideas, and skills in connection with W in order to play his part in the linguistic division of labor.
>
> […] We shall speak of someone as having *acquired* the word 'tiger' if he is able to use it in such a way that (1) his use passes muster (i.e. people don't say of him such things as 'he doesn't know what a tiger *is*', 'he doesn't know the meaning of the word "tiger", etc.) […].[216,217]

[215] I am using the phrase '*S* is not en rapport with *F*s' as a way of saying that she is not using the public word *w* to denote *F*s.

[216] Putnam, Hilary (1975, 1975): pp. 246-247.

[217] Putnam goes on to clarify: "and (2) his total way of being situated in the world and in his linguistic community is such that the socially determined extension of the word 'tiger' in his idiolect is the set of tigers". The second part of Putnam's condition for participating in the public 'tiger' talking practice is *question begging*. When we wonder whether Alice is using our word 'arthritis' we are asking whether with her word 'arthritis' she denotes all and only arthritic conditions. When we are asking whether Jane is using our word 'Pompey' we want to know whether she is denoting Pompey with her name 'Pompey' or whether she is denoting a different object Pompeii.

We don't want to ascribe outrageous or as Putnam states "mind-boggling" statements to rational agents. From a select group of key truths, we expect competent speakers to *not* fail to maintain a critical number of those truths.[218,219]

> What I contend is that speakers are *required* to know something about (stereotypical) tigers in order to count as having acquired the word 'tiger'; something about elm trees (or anyhow, about the stereotype thereof) to count as having acquired the word 'elm'; etc. [...] The linguistic community [...] has its minimum standards, with respect both to syntax and 'semantics'.[220,221,222,223]

---

[218] Neely considers the objection that a speaker may never use some esoteric word in conversations. She proposes to reformulate the requirement of "passing muster" counterfactually by claiming that "a person is competent if her use would pass muster with her community (under the appropriate circumstances)" (Neely (2005): p.45). She ultimately favors an account that assesses a speaker's competence with regard to a word based on beliefs the speaker has, rather than beliefs that were attributed to her. She is worried that speakers could deceive others by bluffing even minimal competence (ibid.: pp.46-47) and shies away from an account that offers two clauses with distinct functions for the determination of word competence (ibid. p.39). I believe that Putnam consults the beliefs a speaker has, as well as those we ascribe to her regarding a given topic, due to the reciprocal roles resulting from participating in a practice. The participant must (minimally) master certain rules or tasks and be perceived as (minimally) mastering them by other participants.

[219] Perhaps it is permissible to doubt whether orangutan are animals (perhaps Joe believes they are spy robots from Mars), but if Joe were to also say the following: "Orangutan are the size of an ant and they are transparent" – it seems impossible to understand Joe's word 'orangutan' as denoting orangutan.

[220] Putnam (1975, 1975): p. 248

[221] Putnam goes on to explain that how many true beliefs one must to associate with a word and its referent is context and culture dependent: "In our culture speakers are required to know what tigers look like (if they acquire the word 'tiger', and this is virtually obligatory); they are not required to know the fine details (such as leaf shape) of what an elm tree looks like. English speakers are *required by their linguistic community* to be able to tell tigers from leopards; they are not required to be able to tell elm trees from beech trees" (Putnam (1975, 1975): pp. 226-227). King considers that different categories of words may have different standards of competence associated with them (King (1998): p.169).

[222] Neely points out that treatment of competence with regard to a public word is an idealization. Sometimes we must acquire the competence of a cluster of expressions or concepts. She writes: "Terms such as 'set,' 'set membership,' and 'element of a set' are very closely related; just as we cannot understand the notion of set membership without having some idea of what a set is, so too we may wonder whether it is truly possible to be competent in 'set' without knowing what it is to be a member of a set. Competence in one notion seems linked to competence in the other two; it may well be that we must consider them together" (Neely (2005): p. 29).

[223] Neely raises several important criticisms against Putnam's account of word acquisition. She objects to the idea that all speakers minimally competent with regard to tiger must share the *same* stereotype (it is perhaps better to speak of paradigms here). "It is possible to imagine a case where we simply must have sufficient overlap with other speakers; I may know that tigers are large orange and black felines, while Jones knows that tigers are orange and black felines from Asia, and Smith knows that tigers are large carnivorous animals from Asia" (ibid. p. 42-43). I would even go so far as to claim that no such overlap is necessary. We could imagine one person defining 'money' such: money is a medium of exchange, while a second person (the diligent child who puts her pocket money in the piggy bank) may think of it as

This fits with our previous observations that we simply cannot make sense of a person confusing tigers with snowballs or somebody confusing Pompey with Pompeii – if you believe that Pompeii is a person, then you don't understand (at all) what Pompeii was. We can, however, imagine a patient wondering whether her arthritis has moved from her knees to her thighs. Putnam's passing muster clause is merely a rule of thumb and doesn't provide us with specific advice of how to proceed with our word-individuation in such tricky cases.

Putnam's passing muster clause is not sufficiently fine-grained, as the following case will show. I will provide an example where a speaker S_1 intends to repeat a word *w* with S_1's production of token *t* and another speaker S_2 interprets *t* as standing for *w* – that is S_1's use of *t* passes muster as *w* – but *t* turns out to be a repetition of a different word *w'*. Suppose that Charles is from England and familiar with the British English word 'robin$_{BE}$' that denotes the European Robin. British speakers have described the bird to Charles and he remembers that it is easy to recognize by its distinctive red chest. He has not yet had the pleasure to see one with his own eyes. Let's further assume that Charles has not yet spoken to Americans about (American) robins. I take for granted that Charles's word 'robin' cannot be a node on two distinct repetition trees; it will either have to constitute a performance of the British 'robin$_{BE}$' or 'American

---

the stuff to accumulate in her piggy bank, a means to savings. A third person may think of that very same (coin and paper) money as a useless relic. Still, we would want to argue that those people can all talk about money – even though they will clearly have some disagreements about how one ought to treat money. Neely is worried that people with physical restrictions (i.e. blindness) may not be able to share some key beliefs generally present in other people's paradigms i.e. that tigers are yellow and black striped (ibid.: p. 42).
More importantly, she points out Putnam's *dual* account of word acquisition: "This ambiguity in Putnam's talk highlights the distinction between belief and attributed belief. The beliefs a person has and the beliefs others attribute to him do not always coincide" (ibid.: p. 43-44). As an approximation Neely suggests: "I thus suggested pursuing a belief-based account of competence in which the community sets standards for competence but does not judge whether someone meets those standards" (ibid.: p. 58). I will ultimately argue that defining which beliefs a speaker must have or be ascribed to by her linguistic community is the *wrong* approach. Which beliefs a speaker must have in order to be considered competent in a public word *w* or having acquired *w* will depend on the retrospective judgments the speaker makes.

'robin$_{AE}$'.[224] Given the description of the case I find it obvious that Charles uses the British word 'robin$_{BE}$', not the American word 'robin$_{AE}$'.

During a phone conversation with an American friend Joe, Charles is asked: "Have you ever seen a robin$_{AE}$ before?" Charles responds: "No, I have not yet seen a robin, but I am keeping my eyes open for one."[225] Can Putnam explain how Joe should interpret Charles's word 'robin'? Charles's mark 'robin' passes muster – Joe would not say of him that he doesn't know the meaning of the word 'robin$_{AE}$' or he doesn't know what a robin$_{AE}$ is. Interpreting his mark to stand for 'robin$_{AE}$', Charles's performance passes the "minimum standards" of Joe's linguistic community "with respect to both syntax and 'semantics'."[226]

It is not clear how stringent Putnam wishes his passing muster clause to be interpreted. The more cases of "talking past" one another one attempts to account for through the passing muster clause, such that in the conversation imagined above we should not individuate Joe's word 'robin' to denote the American Robin – the more we threaten the standard anti-individualist thought examples. Since Charles expects to be able to view this bird in Europe (which is *not* the

[224] In later chapters I discuss the case of Oscar being switched back and forth between Earth and Twin Earth. I don't allow for Oscar's word 'water' to constitute a node on two different repetition trees. Please consult chapter six for further discussion.

[225] Now you may argue that Charles also meant for 'robin' to be a repetition of the word his fellow British speakers are using. I don't deny that. This is part of the problem; Charles attempts to repeat the word he heard from Joe and the word he has head from his British fellow speakers – it turns out that this not a well-formed intention. I just don't think that this is sufficient evidence to argue that Charles thus doesn't say anything at all. I think that if Charles were told that there is the European Robin and there is a different bird-species: the American Robin and Joe was asking him about the latter, then Charles would view himself as having misunderstood Joe. It somewhat depends on how committed Charles is to the use of 'robin$_{BE}$' of his British English linguistic community. If he has had many conversations where he used and heard the word 'robin$_{BE}$', he would take himself to have *consistently* used that same word (all the way up to his conversation with Joe). Hence, he would take himself to have mis-communicated with Joe.

[226] I don't think that the division of linguistic labor could do much work here. We would need an explanation as to why Charles defers to the British experts and not the American experts. The intention to repeat Joe's word 'robin' could be presented as evidence to have Charles potentially defer also to American experts.

natural habitat of the American Robin) – we could view this as a *false* belief about American Robins potentially threatening his performance to pass muster.

This couldn't be the correct interpretation of Putnam's passing muster clause, since it would get us into trouble with Burge's arthritis thought example. Alice has the false belief that arthritis can sometimes spread from the knee to the thigh. If false beliefs mean that the speaker fails to pass muster, then Alice is not even minimally competent with regard to the public word 'arthritis' – which is a charge we want to avoid (as anti-individualists). Not every false conception or lacking true belief should be cause for exclusion from the respective public linguistic practice – or else we threaten the entire social anti-individualist project. Therefore, Putnam's passing muster clause is either going to be too strong such as to undermine some standard (social) anti-individualist thought examples, or too weak such as to fail to account for Charles' lack of competence with regard to the American word 'robin$_{AE}$'. I don't see how any set of conceptions that render a speaker competent with regard to 'robin$_{AE}$' would ever make one fail to pass muster with regard to 'robin$_{BE}$' (and *vice versa*).

Some entities are very similar in the sense that many characteristics attributed to entity *a* are also true of entity *b* and *vice versa*. There are only a few significant characteristics $c_1 \ldots c_n$ that aren't true of both entities and thus can potentially be referred to as criteria that distinguish *a*'s from *b*'s. At the heart of social anti-individualism lays the assumption that a speaker *S* can be under- or misinformed about what a given word *w* denotes and *S* can still use *w* to effectively talk about *w*'s referent. It can therefore *appear* to the linguistic community that *S* is making a true claim about *a*'s, even though *S* is making a false claim about *b*'s. The passing muster clause cannot tell us whether a speaker is using the American word 'robin$_{AE}$' or the British word

'robin$_{BE}$' – because the conceptions any speaker associates with either term should pass muster in both linguistic communities.

Burge, like Putnam, acknowledges the appropriateness of deviant interpretations of other people's words in *some* cases:

> There are, of course, numerous situations in which we normally reinterpret or discount a person's words in deciding what he thinks. [... We do so in situations where] a foreigner (or anyone) utters something without comprehension.[227]
> Misuses or failure of understanding exemplified by malapropism, tongue slips, extreme 'category' misuses, the first uses of words by very young children, and the fumblings of foreigners, all normally and rightly occasion reinterpretation. Most other cases are more complex. […]
> It is not part of my view that just because a person is using the same word forms as others in a given social network, the person's words express the same concepts that his fellows' words do.[228,229]

He stresses how difficult it is to decide, at times, whether a speaker merely succeeds in repeating a mark or whether she also succeeds in repeating the word (in the thick sense):

> Of course, it is often difficult to find evidential grounds for attributing an object-level attitude *as opposed* to its metalinguistic counterpart. […] We do want a general account of these cases [of reinterpretation].[230]

Determining whether a speaker is competent with regard to a given word is important to Burge. His description of what competence entails reminds us of Putnam's "passing muster" clause:

> The expression, '*understanding (mastering) a notion*' is to be construed more or less intuitively. Understanding the notion of contract comes roughly to knowing what a contract is. […] Talk of notions is roughly similar to talk of concepts in an informal sense.[231]

According to Burge, being part of a linguistic community comes with responsibilities or commitments.

> People are frequently held, and hold themselves, to the standards of their community when misuse or misunderstanding are at issue. One should distinguish these cases, which seem to

[227] Burge (1979, 2007): p. 118.
[228] Burge (2007): p. 176.
[229] Such re-interpretations are very much a last resort for Burge. He suggests that we can even make sense of a person confusing an orangutan with a fruit drink. Contrary to Burge, I don't find this believable. See Burge, (1979, 2007): p.120.
[230] Burge (1979, 2007): p. 126.
[231] Ibid.: p. 102.

> depend on a certain *responsibility* to communal practice, from cases of automatic reinterpretation.[232]

When "reinterpretation" of a speaker's word or words is appropriate, she never did or has not yet committed herself to a given public linguistic practice or word repetition tree. Burge attempts to clarify when it is appropriate to regard another speaker as simply not using a mark in the same way that her community does – that is, she is *not* competent in the respective public word:

> A person born and bred in the parent community might simply decide (unilaterally) […] to fashion his own usage with regard to particular words, self-consciously opting out of the parent community's conventions in these particulars. In such a case, members of the parent community would not, and should not, attribute mental contents to him on the basis of homophonic construals of words.[233]

Clearly, when I was visiting the dentist, I didn't intend to opt out of the 'angina' linguistic practice. I was convinced that I was using the mark on the questionnaire like the technician and the dentist would. The same goes for Marco Polo who sincerely attempted to participate in the 'Madagascar' linguistic practice as a *consumer*.

Burge suggests other possible conditions that may engender reinterpretation:

> Individuals can fashion idiosyncratic uses of communal words. If their usage corresponds to their own understanding, and they do not rely in unconscious ways on others for fixing the applications of their words or concepts, individuals can cut themselves off from the communal usage. […] Any dependence on others for linguistic or psychological content derives from reliance on others through certain types of causal relations to them.[234]

In Burge's arthritis thought example Alice, in contrast, does rely on the doctor for "filling in" her missing knowledge of arthritis and what her word 'arthritis' exactly denotes. She adapts her conceptions in accordance with the knowledge the doctor offers her. Alice allows the doctor to *persuade* her on this issue and *defers* to his knowledge about arthritis. Recalling my visit to the dentist, I was not willing to defer to the technician as to whether "angina" is the kind of condition that causes pain in the chest, rather than pain in the throat – as I thought. I was

[232] Ibid.: p. 119.
[233] Ibid.: p. 147.
[234] Burge (2007): p. 176.

unwilling to let her persuade me as to where the condition that my word 'angina' refers to causes pain.[235] I also didn't think that she was offering me false or mistaken information. It is not as if I thought to be more informed about "angina" than she; I quickly came to suspect that we were talking past one another and thus not disagreeing. In turn, her mark 'angina' and my mark 'angina' do not belong to the same repetition tree.

Burge further describes the mechanism of deference thus:

> The language does not present a standard of competence independent of individuals' activity. Minimal competence consists in conformity to the practice of others. 'Greatest competence' consists in abilities to draw distinctions, to produce precisifications, to use numerous linguistic resources, to offer counterexamples to proposed equivalences – that elicit the reflective agreement of other competent speakers. We may imagine a vast, ragged network of interdependence, established by patterns of deference which lead back to people who would elicit the assent of others. […] To put it crudely, a person counts as among the *most competent* if he or she would be *persuasive* to other competent speakers in the use and explication of the language. The point about persuasion is fundamental.[236]

So, some speakers are more active in shaping the practice of how a given word is used, while others *follow* their lead. This pre-supposes that the minimally competent speakers *M* (who maintain a significant amount of false conceptions and may also be under-informed) allow the more competent users *C* to *correct their* conceptions regarding the topic accordingly and, if needed, members of group *M* will generally defer to the knowledge of members of group *C*.[237,238,239] That doesn't mean that a language "consumer" must always bend her beliefs to those of experts. There are people who doubt that global warming is man-made and they still consider

---

[235] I don't think that my intention to make an assertion in English is sufficient evidence for my having been using my word '$angina_N$' to denote angina (pectoris). The person who is pointing at the snowball and exclaims: "This is a tiger!" presumably also believes that she is making an intelligible assertion in English.

[236] Burge (1986, 2007): pp. 258-259.

[237] Putnam alludes to a "division of linguistic labor": "[w]henever a term is subject to the division of linguistic labor, the 'average' speaker who acquires it does not acquire anything that fixes its extension. […] it is only the sociolinguistic state of the collective linguistic body to which the speaker belongs that fixes the extension" Putnam (1975, 1975): p.229.

[238] Kaplan is also a proponent of dividing speakers into consumers and producers (Kaplan (1989)).

[239] Evans proposes a similar division. He labels the less competent 'consumers' and the experts 'producers' Evans (1982): pp. 376-377.

themselves to be using the mark 'global warming' with the same meaning – describing the same condition – as the experts. They view themselves in *disagreement* with the experts. When the defining characteristics are concerned, the language consumer must show *some* willingness to defer to the relevant experts. Was I willing to defer to the technician with regard to "angina"? Up to a point I was. I trusted the dentist's assessment of needing to know whether I had suffered from angina$_N$ in the past in order to offer me better dental treatment (even though I had my doubts I diligently took the time to fill out the laborious questionnaire). If she had told me "angina can cause serious complications during a variety of dental treatments" – I would have believed her and would have made the appropriate changes in my conceptions regarding angina$_N$. I was not willing to defer to her when it came to determining whether angina$_N$ causes pain in the chest or the throat. I considered myself to *know* that it causes pain in the throat. This, judgment went hand in hand with my belief that we were *not* talking about the same condition. To what degree a speaker's tolerance for corrections is *actually* being tested is arbitrary. Had the technician been less inquisitive, I would have appeared to be the language consumer. A proven attitude of deference can be deceiving, unless it has been considerably put to the test.

Angina pectoris and *angina tonsillaris* happen to be quite different medical conditions. If, on the other hand, the two terms denote rather similar kinds of things such as the European Robin or the American Robin, the speaker may not notice that she *ought not to* defer to the expert. Recall the earlier robin thought example, where I assumed myself familiar only with the British word 'robin$_{BE}$'. Imagine my American friend to be particularly fond of American Robins. He reads several books devoted to the species and spends hours studying their behavior in his backyard. When he tells me "the Robin$_{AE}$ is the state bird of Michigan and can frequently be sighted in that state" I consider this to be something I didn't know about robins$_{BE}$. I thus defer to my friend,

even thought I shouldn't have, since his statement is false of the European Robin. A speaker's willingness to defer to more competent speakers of a word *w* thus cannot be a sufficient condition for minimal competence with regard to *w* – or else I would have automatically become competent with regard to American Robins. What Burge, however, gets right, is that if I were to defer frequently to my American friend (or other competent speakers of the word 'robin$_{AE}$'), then my word 'robin' would also denote the American Robin.[240]

Sometimes people are particularly stubborn and are unwilling to defer to the expert community, which leads me to conclude that an attitude of deference on significant characteristics of a given phenomenon does not represent a necessary condition. Admittedly there are only a handful of such examples. Some Americans doubt whether global warming is man made, even though there exists an overwhelming consensus among experts that such a correlation exists. Global warming skeptics still view themselves as talking about the *same topic*. Perhaps deference is secondary and what matters is whether after prolonged linguistic exchanges the interlocutors take each others words *at face value* (whether they extend practical judgments of sameness of extension to each others words).[241]

The Madagascar thought experiment is tricky. Marco Polo viewed himself as a *language consumer*, but he failed to defer to his interlocutor who introduced him to the same 'Madagascar'. Insufficient information about what their mark 'Madagascar' denotes was conveyed to Marco Polo such that a misunderstanding had him believe that 'Madagascar' denotes the island he was standing on. He was not under the impression of having a firm "grip" (that is, significant understanding) of the topic and was thus very much open to being corrected by others. We don't really know why (perhaps nobody *dared* correcting Marco Polo), but the

[240] I will further discuss cases of this sort in chapter six.
[241] Ebbs (2009): chapter 4.

corrections didn't happen (or they were misunderstood). In Macro Polo's case, deference would have played a *constituting role*. Had he been appropriately set right by the relevant competent speakers, Marco Polo would have become a competent speaker of their word 'Madagascar'. It seems to me that the more the speaker views herself as competent with regard to a given word *w* and believes herself to be learned about what *w* denotes, the less open she is to corrections regarding significant characteristics of what *w* denotes.

In the arthritis thought example, the speaker was willing to defer and minimally competent; the angina thought example showed the speaker to be unwilling to defer and not minimally competent – those cases led us to suspect a close connection between a speaker's willingness to defer and minimal linguistic competence. In contrast, the global warming skeptics are unwilling to defer and minimally competent; Marco Polo was willing to defer, but lacked minimal linguistic competence – those cases had us doubt the determining role of deference. Still, Marco Polo's interaction with the Malay (or Arab) sailors was probably too insignificant and the attitude of global warming skeptics strikes many as borderline rational or at least ignorant. It is probably no accident that those skeptics are *generally* hostile towards contributions by academics. The degree to which experts are generally trusted appears to be partially dependent on culture.

It is partly up to the speaker whether she is acquiring a new *word* or whether she is acquiring a new *belief*. We cannot argue that if a speaker intends to repeat an expert's word, she *must* commit herself (in lure of being the layman) to the expert's word. Only after already having individuated the speaker's word can we assess whether her interlocutor can be viewed as having superior expertise. I know more about the referent of the German word '*Angina*' than the English-speaking dentist. Of course, we want to shield ourselves from counterexamples where

speakers are excessively stubborn. If Marco Polo were to have said "I don't care what *they* call 'Madagascar'. I am Marco Polo! I thereby *declare* this island to be Madagascar…" we would take him to have deliberately *opted out* of a public practice. The powerful counterexamples against Kaplan all involve a certain naïveté on the part of the speaker, since they "humbly" attempted to repeat the words of their interlocutors.

I am sympathetic with Burge's strategy of taking the speaker's attitude of deference to more competent speakers seriously when interpreting her words. Alice is willing to be corrected by her doctor, while I am *not* willing to be so corrected with regard to what my (original) word 'angina' is true of (due to my additional commitments to the German word '*Angina*'. Sometimes it takes more than one expert source for the information that conflicts with our deeply engrained beliefs to "sink in". Even though I am quite certain that there are no black swans, I am willing to revise my beliefs in the face of overwhelming evidence from respectable sources to the contrary.

3.8 Conclusion

We have seen that merely intending to repeat another person's word is not sufficient to do so. Kaplan doesn't achieve what he sets himself up to deliver at the inter-subjective level. We need to explain how it is possible that the false beliefs Burge's arthritis patient has regarding the doctor's word 'arthritis' (that she means to repeat) do *not* prevent her from repeating that word, whereas in the above angina thought example my *false* beliefs regarding the technician's word 'angina' do interfere with my repetition of that word. What "makes" those cases so different?

I will elaborate upon Burge's idea that the speaker must be willing to defer to other competent speakers. I argue that it is the speakers' own *informed* word-individuations that render the two cases different at the semantic level. Even if the patient were to be encouraged to make further

inquiries about the doctor's word 'arthritis' she wouldn't cease to identify her word 'arthritis$_P$' with his word 'arthritis$_D$'. My dental visit progressed differently. The more I learned about the technician's word 'angina$_T$', the more I was convinced that I was using a *different* (albeit homophonic and homographic) word. The judgments the speaker makes about her own words and how they relate to the words of her interlocutors must be taken seriously. For now I want to observe that Kaplan doesn't seem to have the resources to distinguish Burge's arthritis thought example from cases where speakers intend to repeat a word but *fail* to do so. The intentions to defer exist in *both* examples, but the speaker only succeeds in one of the two.

I discussed Kaplan's suggestion that something like Donnellan's referential use may be at play in situations where the speaker fails in her intention to repeat a given word. Since Kaplan doesn't exactly tell us how this is supposed to work, it remains unclear whether the strategy is successful. What I would like to carry into the next chapter is the observation that the speaker generally has a sense of which conceptions (or referential intentions) she would be willing to *revise* in order to do justice to her intention to repeat a given word.

CHAPTER 4

UNDER WHAT CONDITIONS ARE TWO UTTERANCES PERFORMANCES OF THE SAME WORD? PROPOSAL: RESTROSPECTIVE JUDGMENTS

Kommunikationsgestörete interessieren mich am allermeisten. Alles was ich als komisch empfinde, entsteht aus der zerbröselten Kommunikation, aus dem Aneinander-vorbei-reden.
-- Victor von Bülow (alias 'Loriot')

Communication-defective people interest me the most. Everything that I find comical emerges out of shattered communication, out of talking-past-each-other.
-- Victor von Bülow (alias 'Loriot'

4.0 Introduction

I am fascinated by the observation how, at times, we realize in retrospect that during some conversation we were entirely talking *past* our interlocutors. A multitude of conditions can bring about such miscommunication. I will focus on those resulting from semantic misunderstandings, where the utterance of one speaker is similar to some utterance of her interlocutor without those constituting performances of the same word.[242] Such linguistic exchanges *prima facie* comprise a threat for Kaplan's notion of words, which relies on the speaker *succeeding* in her attempt to repeat her interlocutor's word *w* such that her own word *w'* (thus produced) "inherits" its reference from *w*.

In the last chapter I recalled my visit to the dentist, where the technician was using 'angina' to mean *angina pectoris*, while I was using the mark 'angina' to stand for *angina tonsillaris*. I was intending to successfully communicate with my audience; in fact, I "re-used" one of the technician's marks by placing a checkmark next to 'angina' on the *pre-printed* questionnaire. I observed that in spite of my intention to repeat my interlocutor's word 'angina',[243] I *failed*. My word 'angina' was already bound by prior commitments to a different public repetition tree: the

[242] The respective word is here assumed *not* to be indexical.
[243] That is, I fail to repeat the word in the semantically loaded thick sense of using a word.

German word '*Angina*'. If there are two distinct public words sharing a relevantly similar form,[244] that compete having tokens of the speaker's word *w* as instances, we have to figure out which corresponding repetition tree *w* is related to, such as to individuate *w* unambiguously.[245]

Evans's counterexample reveals how a person can unknowingly create an entirely *new* repetition tree, while intending to repeat some already existing public word. As the historically inaccurate story goes, Marco Polo attempted to repeat the word 'Madagascar', which he encountered for the first time during an exchange with his Arab (or Malay) encounter. He developed referential intentions that 'Madagascar' denotes the island both speakers were standing on, while in fact, his interlocutor was using 'Madagascar' to denote some area on mainland Africa. Marco Polo viewed himself as a language *consumer* intending to repeat an *already existing* foreign name; nevertheless, he *failed* in his attempt. Kaplan (and Kripke's) explanations of how reference inheritance can fail, exploit the circumstance that Marco Polo's referential intentions *conflict* with the reference of the Arab (or Malay) word 'Madagascar'. The same goes for my referential intentions to use 'angina' to denote the condition that caused severe pain in my throat many years back. So far, Kaplan and Kripke's exception clause functions as intended.

Since Kaplan and Kripke's suggestions towards a solution are vague, it isn't entirely clear what the notion of *referential intentions* entails. In Burge's arthritis thought-example – one of the standard thought examples and a crucial source of support for anti-individualism the patient

[244] That is, we recognize them as displaying a relevantly similar form (see chapters one and two for further details).

[245] Evans raises a similar concern pertaining to ambiguous marks within a given language: "The precise statement of the way in which individual utterances exploit the general practices of the community must be a little complicated, because of the possibility of there being several distinct practices in the community involving the same name [*form*]. If a speaker is to refer to something by using a name, then it is necessary that he manifest *which* name-using practice he intends to be, and to be taken to be, participating in" (Evans (1982): p. 384).

intends to refer to a specific condition that may cause pain in the thigh when she utters 'arthritis', but that condition couldn't be arthritis;[246] hence, the conflict with her intention to repeat her doctor's word 'arthritis'. With many false beliefs potentially being at variance with the intention to preserve reference in the historical chain of word transmission the recommended exception clause potentially threatens to *undermine* anti-individualism altogether. Kripke and Kaplan need it to be the case that Alice's false beliefs do not constitute the kind of referential intention that *override* her deferential intentions – her intention to repeat her doctor's word 'arthritis'. In this chapter I will be exploring ways to contain this referential use such as to render it compatible with the standard anti-individualist thought experiments.

I will propose an approach that occupies the middle ground between (social) anti-individualism and the kind of individualism that individuates a speaker's words without input from the speaker's linguistic community. According to my proposal, the content of the mental register doesn't by itself individuate the individual's word *w*, although it cannot be entirely ignored either when it comes to individuating *w*. Mental registers are the individual version of

[246] Burge suggests "If their usage corresponds to their own understanding, and they do not rely in unconscious ways on others for fixing the applications of their words or concepts, individuals can cut themselves off from the communal usage" (Burge (2007): p. 176). Suppose there are several conditions a, b, and c that could potentially cause such pain in the thigh. The patient thus still relies on medical experts to determine which one of those conditions, whether it is a, b or c, that afflicts her thigh. Her definition for the new word 'arthritis' is vague. Does it denote arthritis plus a, or arthritis plus b…? The heart of the matter cannot be the element of vagueness, since the experts' definition of arthritis is similarly vague. We know so little about arthritis that we don't even know whether it is a conglomerate of distinct conditions or one single condition. Nevertheless the patient wants to have her pain alleviated. She therefore relies on experts identifying the true cause of her pain. A referential intention of 'arthritis' refers to whatever condition causes the pain in my thigh (and my knees) seems hesitant to begin with. At the same time it isn't as if Marco Polo may not want to rely on some experts to determine the exact physical extension of Madagascar. How far into the sea does an island reach? This is why I think Burge's other suggestion of a willingness of being corrected offers a deeper understanding. After years of filling out maps and official documents with his word 'Madagascar' Marco Polo finds out that the original word 'Madagascar' he heard refers to a place on mainland Africa. It is unlikely that Marco Polo would now decide: I had all those maps drawn *wrong*. He would probably judge: well, their word 'Madagascar' is not our word 'Madagascar'; we will have to keep the two separate. The arthritis patient, in contrast, even if she believed for years that arthritis is the kind of condition that afflicts the muscles, she should be open to correction from medical experts.

Kaplan's public repetition trees. The speaker *S*'s own inter-personal judgments of word-sameness resulting from her intentions to repeat a token *t* previously performed by *S* (herself) and the resulting performance *t'* make it the case that *t* and *t'* belong to the same (mental) register or word-repetition tree. I argue that the speaker has linguistic control over which tokens belong to the same word type. We consult the mental register's content, when it comes to deciding whether it is linked to any public repetition tree and if so, which one. I will argue that in the end we must leave it up to the speaker's own judgments whether she interprets her word *w* as repeating a particular public word *w'*. Of course, we often don't know enough about the potential public word *w'* to reliably assess whether *w'* and *w* constitute performances of the same word. This is why the speaker must make such judgments from an *informed perspective* – one not every speaker will actually take the time to acquire. I thus propose a hypothetical rule: if *from an informed perspective* the speaker were to interpret *w'* and *w* as performances of the same word, then *w* is linked to the public repetition tree that *w'* belongs to. That is, *w* and *w'* share the same semantic ancestry. Conflicting referential intentions only interfere with the word-identity of *w* and *w'*, if – *from an informed perspective* – the individual decides that her beliefs pertaining to what *w* denotes cannot be interpreted as beliefs pertaining to what *w'* denotes.[247]

At the heart of my account lies the observation that the speaker cannot be wrong about individuating her *own* words in terms of *intra*-subjective practical judgments of sameness of satisfaction.[248,249] This does *not* entail that she cannot be wrong about how she translates her

[247] This is only a sufficient condition. My proposal leaves out cases where the speaker may never acquire an *informed perspective*. Seeking out information about the public word *w'* may effect the content of *w*. If Mr. Anchovy, whom I will introduce later, would have consulted a dictionary about the mark "lion" early on, then he would have been more successful with his attempt to repeat the public word 'lion'.
[248] Ebbs describes commitments of this sort as 'practical judgments of sameness of satisfaction' ('PJSS' for short). He explains: "I call such judgments *practical* because they are exercises of a learned yet typically non-deliberative ability that is acquired over a period of years. I call an exercise of this ability a *judgment because* we will reject it if we come to think it is wrong" (Ebbs (2009): p. 5).

words into the public language. Quite the contrary, in order for the speaker to always be right about her own word identities, she must sometimes deem *wrong* translations into the public language which she previously judged appropriate. Remember that I used to think that some tokens of my word 'angina$_1$' were correctly to be translated into ordinary English as 'angina'. I also maintained that some tokens of my word 'angina$_2$' were to be translated into ordinary German as '*Angina*'. The ordinary English word 'angina' and the ordinary German word '*Angina*' denote entirely different medical conditions. As a result, some identities have to be abandoned: either 'angina$_1$' is a different word from the ordinary English public word 'angina', or 'angina$_2$' is a different word from the ordinary German word '*Angina*', or 'angina$_1$' and 'angina$_2$' are not tokens of the same word. Since the speaker judges her own utterances 'angina$_1$' and 'angina$_2$' to be instances of the same word, I argue that this is an identity we cannot revise (or else we threaten the *a priority* of logical reasoning).[250] If we ascribe to the speaker two distinct words that merely share the same form 'angina' – even though the speaker herself is oblivious to that fact – this may force us to ascribe to her irrational or invalid inferences which she is unable to correct (without input from others). We usually don't consider the false beliefs of regarding one entity to be two, or the other way around, confusing two distinct entities to be one, as grounds for accusing a person of irrationality. This is why I maintain that the speaker is infallible about *intra-subjective* practical judgments of sameness of satisfaction – an individual is never wrong about thus individuating her *own* words.[251]

---

[249] See also Fiengo & May (2006): p.53.
[250] See my chapters five and six for further analysis.
[251] Unless there are serious neurological defects disrupting the proper function of the individual's cognitive apparatus.

Kaplan therefore was exactly right: I cannot fail in my intentions to repeat my own words![252] This is why in my own account I ascribe *intra*-subjective practical judgments of sameness of satisfaction priority over *inter*-subjective practical judgments of sameness of satisfaction. In order to do justice to the inferences an individual is in a position to draw from her own (limited) first person perspective, I introduce mental registers to allow us to keep track of how the individual individuates her own words.

4.1 Lion taming: a thought example

From Monty Python's "Vocational Guidance Counselor"[253]

> [Background: A chartered accountant, Mr. Anchovy, consults a vocational guidance counselor about a suitable new career. After Mr. Anchovy rejects the counselor's advice of keeping his current job, the following dialogue unfolds:]
>
> Anchovy: But don't you see, I came here to find a new job, a new life, a new meaning to my existence. Can't you help me?
> Counselor: Well, do you have any idea of what you want to do?
> Anchovy: Yes, yes I have.
> Counselor: What?
> Anchovy: (boldly) **Lion taming**!
> Counselor: Well yes. Yes. Of course, it's a bit of a jump isn't it? I mean, er, chartered accountancy to lion taming in one go. You don't think it might be better if you worked your way towards lion taming, say, via banking...
> Anchovy: No, no, no, no. No. I don't want to wait. At nine o'clock tomorrow I want to be in there, taming.
> Counselor: Fine, fine. But do you, do you have any qualifications?
> Anchovy: Yes, I've got a hat.
> Counselor: A hat?
> Anchovy: Yes, a hat. A lion taming hat. **A hat with 'LION TAMER' on it.** I got it at Harrods. And it lights up saying 'lion tamer' in great big neon letters, so that **you can tame them after dark when they're less stroppy.**
> Counselor: I see, I see.
> Anchovy: And you can switch it off during the day time, and claim reasonable wear and tear as allowable professional expenses under paragraph 335C...

[252] This doesn't mean that the individual cannot have delusional memories. When I utter u_1: "I suffered from angina$_1$ [strep throat] at the age of seventeen" – I may very well be wrong about that. If I believe that my own words 'angina$_1$' and 'Angina$_2$' keep track of the same medical condition, then I *know* that my utterance u_2: "In meinem siebzehnten Lebensjahr litt ich unter Angina" [I suffered from strep throat at the age of seventeen] and u_1 express the *same* state of affairs.

[253] This sketch appeared in the Flying Circus TV Show, episode 10; it also appeared in the movie "And Now for Something Completely Different,". I would like to thank Steve Wagner for bringing the relevance of this sketch to my attention.

Counselor: Yes, yes, yes, I do follow, Mr. Anchovy, but you see the snag is... if I now call Mr. Chipperfield and say to him, "look here, I've got a forty-five-year-old chartered accountant with me who wants to become a lion tamer", his first question is not going to be "does he have his own hat?" He's going to ask what sort of experience you've had with lions.

Anchovy: Well I ... **I've seen them at the zoo.**

Counselor: Good, good, good.

Anchovy: **Lively brown furry things with short stumpy legs and great long noses. I don't know what all the fuss is about, I could tame one of those. They look pretty tame to start with.**

Counselor: And these, er, these lions ... how high are they?

Anchovy: (indicating a height of one foot) Well they're about so high, you know. They don't frighten me at all.

Counselor: Really. **And do these lions eat ants?**

Anchovy: **Yes, that's right.**

Counselor: Er, well, Mr. Anchovy ... **I'm afraid what you've got hold of there is an anteater.**

Anchovy: A what?

Counselor: **An anteater. Not a lion. You see a lion is a huge savage beast, about five feet high, ten feet long, weighing about four hundred pounds, running forty miles per hour, with masses of sharp pointed teeth and nasty long razor-sharp claws that can rip your belly open before you can say 'Eric Robinson', and they look like this.**

(The counselor pulls out a large picture of a lion and shows it to Mr. Anchovy who screams and passes out.) […]

Assuming that the hat-maker was using the public word 'lion' we know that Mr. Anchovy has been exposed to our word 'lion'. He clearly means to repeat our word 'lion', or else he wouldn't think that the neon letters 'LION TAMER' intelligibly describe to others the profession he would like to pursue. At first, the counselor takes Mr. Anchovy to express his desire to switch his profession from chartered accountancy to lion taming, but he quickly develops suspicions as to whether Mr. Anchovy really knows what he is talking about. The counselor soon suspects that a better interpretation for Mr. Anchovy's word 'lion_A' is the public word 'anteater'.

The counselor is clearly not taking Mr. Anchovy's word 'lion_A' to denote lions. When he inquires: "And do *these* lions eat ants?", his choice of words shows that he is *distancing* himself from a face-value interpretation of Mr. Anchovy's word 'lion_A': he considers Mr. Anchovy's "lions" to be *unlike* actual lions (different from what the public word 'lions' denotes).[254]

[254] The consultant's surprise is an indication that he *at first* took Mr. Anchovy's word 'lion_A' to denote lions. As the conversation goes on, however, he retracts this interpretation and takes 'lion_A' to denote some other animal (perhaps anteaters).

At some point Mr. Anchovy must have been exposed to our public word 'lion'. This is what caused the creation of a mental register (or file) with the label 'lion'.[255] Registers are theoretical devices reflecting how Mr. Anchovy can "keep track" of his various tokens of his word '$lion_A$' throughout conceptual changes. He acquires richer conceptions over time without changing what his word '$lion_A$' denotes.[256, 257] Since I am asserting the possibility that Mr. Anchovy's word '$lion_A$' may not be adequately connected to *any* public word, public concepts do not provide me with a reference point to "keep track" of the various tokens of '$lion_A$'. Even if Mr. Anchovy's word '$lion_A$' ends up being fictitious, he had in mind the same topic when uttering '$lion_A$' from whence he purchased the lion tamer hat until his conversation with the counselor.

From the above conversation we can infer some of the conceptions Mr. Anchovy associates with his $lion_A$-register over time:

$Lions_A$ are animals.

$Lions_A$ can be tamed.

$Lions_A$ are less stroppy in the dark.

$Lions_A$ can be seen at the zoo.

$Lions_A$ are lively.

---

[255] Some understand these registers to constitute cognitive mechanisms. See for example Schroeter (2008): "Mental Files = $_{df}$ the stable cognitive mechanisms that bind together a changing body of attitudes and cognitive dispositions in such a way that, from the subject's perspective, it seems obvious […] that they pertain to the very same subject matter". I am not committed to this interpretation of 'register'. Especially when we are first exposed to a word – recall my 'Pompey' example in chapter three – we do not yet associate *any* conceptions with it. There are (not yet) any cognitive mechanisms that the subject intends to interpret as pertaining to the same subject matter. How would consulting the cognitive content of what the speaker associates with 'Pompey' at this point be of any help?

[256] "Or suppose that I learn more about how to discriminate water from other (possible or actual) colorless, tasteless, potable liquids. In such cases, I learn something about chairs or wate that I did not know before. In these cases it is simply not true that the reference of my words 'chair' and 'water' must change. Although it is trye that my conceptions – my explication – changes […]" (Burge (1989): p.284).

[257] "Usually, all participants begin the discussion without being able to give a precisely correct normative characterization; […] Participations commonly regard their object-level thoughts (thoughts about, say, chairs) as undergoing correction in the course of the inquiry" (Burge (1986): p. 261).

Lions$_A$ have brown fur.

Lions$_A$ have short stumpy legs.

Lions$_A$ have great long noses.

Lions$_A$ eat ants.

Lions$_A$ aren't frightening…

We don't know how Mr. Anchovy acquired these beliefs, perhaps people deliberately misinformed him or he repeatedly misunderstood what people intended to sincerely convey to him about lions. We do know that we can't interpret his mental register with the label 'lion$_A$' to pertain to or be *en rapport* with lions. Mr. Anchovy clearly doesn't interpret his own word 'lion-$_A$' to denote lions either, or he wouldn't be so shocked when faced with the picture of an actual lion. We can reasonably assume that Mr. Anchovy would concur with the counselor's interpretation that his word 'lion$_A$' does not denote lions but perhaps anteaters. Upon finding out what the expression displayed on his hat really denotes he will very likely retract his application for lion tamer – Mr. Anchovy is not that reckless!

Unlike Marco Polo, Mr. Anchovy is not in a position of power, nor is he an animal expert. Other speakers are unlikely to catch on to his idiosyncratic use of the mark 'lion'. This means Mr. Anchovy ought to make some changes. If he thinks that the public word 'anteater' is a good interpretation for his lion$_A$-register, then he can simply change the label of his register from 'lion' to 'anteater'. From then on he will be able to successfully communicate his beliefs about the original lion$_A$-register by talking to others about anteaters. If, on the other hand, Mr. Anchovy has several significant beliefs associated with his lion$_A$-register which are entirely incompatible with our public word 'anteater' then Mr. Anchovy has to look for a better match or perhaps admit that his word 'lion$_A$' simply doesn't refer and that it is a fictitious term.

Of course Mr. Anchovy could doubt whether the counselor knows what he is talking about. After all, what does a vocational guidance counselor know about animals? If Mr. Anchovy wants to be epistemically responsible he ought to consult some reliable source of reference, such as an encyclopedia or dictionary. 'Lion' and 'anteater' are the kind of words one is likely to find referenced in such sources. He will eventually find out that the counselor was speaking truthfully.

The informed perspective is achieved by consulting a *wide selection* of reliable resources, such as *experts, respected dictionaries or encyclopedias*. This should allow the speaker, here Mr. Anchovy, to adequately "grasp" what the public word in question expresses. For example, he looks up 'lion' and reads: "A large tawny-colored cat that lives in prides, found in Africa and northwestern India." My suggested approach pre-supposes that the speaker understands most of the words contained in the description of what 'lion' denotes.[258]

My proposal accords with Putnam's observation that if we were to encounter somebody, let's call him 'Oscar', pointing at a snowball uttering: "This is a tiger!"[259] – then there wouldn't be much use in talking to Oscar about tigers using *our* words 'tiger' or 'tigers'. I cannot imagine Oscar consulting a dictionary and looking up what 'tiger' stands for and then declare as a result that his word 'tiger_O' is the same as the public word 'tiger' (in English). It seems to me that from an informed perspective Oscar would come to realize that his use of the mark 'tiger' is confusing. If he wants to communicate successfully with his linguistic community he must adopt a *new* 'tiger_N' register and use the label of that new register for communicative purposes.

[258] This is something I am entitled to assume, since my proposal is not meant to constitute an analysis of linguistic competence with regard to a language, but instead explores competence with regard to a particular (public) word, thus taking for granted the speaker's overall competence in the respective language.
[259] We are assuming that this snowball doesn't even have the shape of a tiger.

When Mr. Anchovy decides that the public word 'anteater' "fits" his $lion_A$-register, I thereby don't mean to argue that *all* the characteristics Mr. Anchovy finds listed under 'anteater' must also be contained in his $lion_A$-register. It is likely that some characteristics are *missing* in his $lion_A$-register and that furthermore, some of his conceptions will *conflict* with what is being claimed of anteaters in the reference book – Mr. Anchovy is not an anteater expert! What is crucial is whether Mr. Anchovy can interpret himself as having some *false beliefs* about anteaters and being *partly mis-* or *under-informed* about them (i.e. there is no profession dedicated to taming anteaters). It is not plausible that Mr. Anchovy would interpret his $lion_A$-register to be *en rapport* with lions once he consults a reputable encyclopedia about the entry 'lion'. His own potential judgment decides the case!

I can't require that a speaker must *actually* consult expert sources in order for her word-token *t* to constitute a performance of a public repetition tree. If I merely require that we have her *hypothetically* consult a dictionary and judge her token *t* to constitute a repetition of the given entry, then I will get results conflicting with how the speaker keeps track of her words. Imagine we had Mr. Anchovy consult a dictionary shortly after his initial exposure to the public word 'lion' (perhaps after he purchased the lion-tamer hat). This would change how Mr. Anchovy would come to use his word '$lion_A$' – he would not come to use it to denote anteaters. So we must ask what Mr. Anchovy would decide if he were to consult a dictionary at the end of his performance history of '$lion_A$'.[260]

Once Mr. Anchovy realizes that the counselor described an existing type of animal, which other members of their linguistic community refer to with 'lion', a *new* mental register will be

[260] As I will discuss in the last chapter, it is not sufficient to hypothetically ask the speaker to go through this exercise, because sometimes a person's use may seem mature, but then undergo a significant change in use. Oscar associates new stuff with his word 'water' after he has been switched to Twin Earth.

created with the title 'lion$_N$', one to be kept *separate* from his original 'lion$_A$' register. This entails that when Mr. Anchovy acquires new information about one entity, he will not associate it with the register of the other entity. I may have two friends with the name 'Jessica Ramirez'. If I find out that one of them is about to get married, I don't thereby infer that the other Jessica Ramirez is also about to get married. I must keep information about those two individuals separate, except for characteristics that happen to be true of both of them (i.e. they are both women, they both have brown hair, they are both of Mexican decent…).[261] I cannot infer characteristics of one Jessica Ramirez on the basis of characteristics of the other Jessica Ramirez.[262]

4.2 The labels of non-mature registers

I don't deny that in most cases the hearer's word or mental register will inherit reference from the public word – Mr. Anchovy's 'lion$_A$' is an exception! Usually people are epistemically responsible: we attempt to pass on true and justified beliefs. When a person encounters a word for the first time she will inquire about it: "what is a lion?" or "what does 'lion' mean?". The audience *automatically* interprets such questions as inquiries about lions.[263] This is why typically the newly formed lion-register will *predominantly* be filled with characteristics that are true of lions.[264] It is this mechanism of interpreting questions such as "what are lions?" as pertaining to lions – or the reference of other public terms we are familiar with – which will

[261] I will discuss a case where two individuals are mistakenly linked to one proper name register later.
[262] Unless they are identical twins, in which case one may infer some genetically determined characteristics to hold true of both women that go by that name, even though one only experienced those characters to be true of one Jessica Ramirez.
[263] See also Ebbs (2009): p. 5.
[264] I don't mean for other people to use considerations of charity here. I am arguing that it is their duty to interpret Mr. Anchovy's word 'lion$_A$' as denoting lions such as to facilitate his ability to participate in this public practice – until they are presented with evidence to the contrary.

eventually *create a lion register* (a register that is linked to the public 'lion'-repetition-tree within the cognitive apparatus of the novice. Nevertheless, the mere existence of a register with the label 'lion' is not sufficient for it to be *en rapport* with lions (contain beliefs about lions), even if that label is an intended copy of a token of the public word 'lion'. It matters what gets associated with the register over time, since this will effect how the individual will relate her word '$lion_I$' to the public word 'lion' in the future (or any other words for that matter).

If, however, Mr. Anchovy considers *from an informed perspective* that our public word 'anteater' is a good translation for '$lion_A$', then '$lion_A$' is (and always was) individuated through the public word 'anteater'. This means that Mr. Anchovy can be under- or misinformed about what his word '$lion_A$' denotes. Facing this kind of "switching behind the scenes" should not raise doubts as to whether we can generally communicate successfully. As we have seen in Mr. Anchovy's case, it takes some considerable accumulation of unfortunate exchanges for one to be so confused about what other speakers talk about. Even if in that process of collecting content for one's lion-register the individual adds some false information, she won't thereby be prevented to interpret her word '$lion_I$' as a performance of the public word 'lion' in the future. Only a very unusual *accumulation* of unfortunate circumstances must have led Mr. Anchovy to predominantly receive information that either doesn't pertain to lions to begin with, or is simply false.

How can anyone then ask a question about *lions* without having acquired the word 'lion' (in the *thick* sense), without the ability to describe them and without the ability to directly point at them (unless the inquirer already knows the answer to her question)? When Mr. Anchovy was first exposed to the public word 'lion' let's imagine him intending to repeat that very same word in order to find out about the things it denotes: "What is a lion?" We would generally interpret

this to be a question about lions. But how can this be a sensible and grammatically complete question inquiring about lions, if the speaker lacks the word 'lion'?[265]

For Kaplan this question is easily answered. The novice attempts to repeat the word that was produced by the competent speaker (in the thick sense). According to Kaplan such attempts (generally) succeed, the novice thereby repeats the *word* 'lion' of the competent speaker denoting lions and thus asks a question that pertains to lions. It thus seems that a new word is acquired abruptly – Kaplan's account doesn't suggest a gradual acquisition process. How could Mr. Anchovy possibly, at this point, that is, *prior* to having received an answer, be considered as having acquired the public word 'lion' (in the thick sense) – if he doesn't even know anything about lions. His use couldn't even pass muster (except for in inquiries as imagined above).

If the novice is merely using a word of the same *form*, then anything goes – that is, any answer would be potentially apt. A mere label of a register, devoid of conceptions could *turn out* to denote anything whatsoever. Perhaps this is a consequence of the acquisition of competence in a word developing gradually, but it surely would be odd if all our what-is-x questions[266] were grammatically incomplete (unless the inquirer already knew the answer to her question). Furthermore, why do we so strongly believe that if somebody were to answer: "lions are robots operated by aliens from Mars" – she thereby fails to answer the original question? This neglects how we – as competent speakers – should *facilitate* membership in our linguistic community, not deceive people about what our word 'lion' denotes. Therefore, teaching novices our public

[265] Of course I could indexically refer to another person's word 'lion'. I could ask: "What kind of entity does the third word of your last sentence denote?". But this is not what is happening here.
[266] Usually when people have a register devoid of content or almost devoid of content their questions take on a very general form. The analysis I am offering below would also apply to specific questions that are linked to an almost empty register such as: "Are lions dogs?" or "Do lions eat meat?" – we can't really imagine competent speakers raising such questions.

words and explain what they denote – to the best of our knowledge – is our *responsibility* as members of the linguistic community.

The competent speaker ought to provide an answer to the novice (Mr. Anchovy) that will *facilitate membership in the linguistic community* with regard to the word 'lion'. This is why the competent speaker ought to interpret Mr. Anchovy's question as an inquiry about lions. When the competent speaker answers she has accepted *linguistic control*. In such situations people simply don't respond: "Lions aren't anything because you don't know what you are talking about."

At the same time, for a speaker's token *t* to inherit reference form the competent speaker's use of the corresponding public word *w* (i.e. 'lion') the speaker *S* must be receptive to the linguistic control. *S* must be paying attention to their utterances, trust their sincerity and knowledge. *S* must also be willing to *be corrected* about what *S* thinks is true of what *t* denotes by the competent speakers of *w*.

The speaker's token *t* inherits reference first and foremost from future tokens (i.e. *t'*) that constitute repetitions of *t* – from those the speaker associates rich conceptions with. Whether *t* inherits reference from a given public word *w*, depends essentially on whether *S* would interpret *t'* as a repetition of *w* (from an informed perspective).

I cannot argue that reference is always *inherited* from the interlocutor's word 'lion', that is, the competent user of the word 'lion', since this could potentially violate my previous claim that intra-subjective word-individuations must be *correct*. Suppose Mr. Anchovy at some point t_1 asked a competent user of the word 'lion': "What are lions_{t1}?". Perhaps the interlocutors get interrupted and proceed to forget Mr. Anchovy's question. If his word 'lion_{t1}' is *at this point* denoting lions, then this will conflict with his later retrospective judgment that his word 'lion_{A}' is

not to be translated as the public word 'lion' (and it doesn't and never did denote lions) – assuming 'lion_{t1}' and 'lion_{A}' are nodes on the same intra-personal repetition tree.[267] But then how can the question have been an inquiry about lions?

I don't see how I can avoid driving a wedge between how the question is and ought to be interpreted and what the question actually inquires about *from Mr. Anchovy's perspective* – unless I am willing to allow speakers to be wrong about their intra-subjective word individuations. The 'lion_{A}'-register that is thus being created won't *ever* connect to the public 'lion' repetition tree. This is only a problem for cases where the speaker's word-form is not appropriately connected to the public language. For other individuals who pick up the word 'lion' and ask questions such as: "what is a lion?", but later go on to fill their 'lion_{I}'-registers with conceptions compatible with the matured register to denote lions, there is no such discrepancy between the two perspectives. The competent speaker who answers the question *ought to* interpret the question as pertaining to lions and the quasi-word 'lion_{I}' contained in the question turns out to be connected to a lion-register, which in turn is linked to the public-'lion'-repetition tree. *Being a language user comes with responsibilities*.

I propose to make sense of the above observations as follows: The word *w* (whether public or idiosyncratic) is acquired *instantly*, by speaker *S* even though she only *gradually* learns how to use it as a tool in communication or reasoning processes; competence with regard to *w* is acquired gradually! If the word *w* is idiosyncratic (such as Marco Polo's word 'Madagascar'

[267] But he will create a *new* register with the label 'lion_{N}' that will denote lions, since he wants the ability to talk to others about lions. Let me also clarify that the associated register with 'lion_{t1}' is assumed empty, which means that at t_1 Mr. Anchovy's 'lion_{t1}' is not yet a full-blown word with the possible denotation of all and only anteaters. As I will explain later his mark 'lion_{t1}' is merely a quasi-word. Still, it is linked to the same register that his fully developed word 'lion_{A}' will also be linked to. This is why strictly speaking we cannot say that 'lion_{t1}' and 'lion_{A}' are instances of the same *word* since the former is merely a quasi word. Nevertheless, 'lion_{t1}' is linked to the same *topic*.

was at first), then the content of the mature register[268] constitutes the target for the conceptions *S* ought to come to associate with *w* over time. If the word *w* is public, then the content of the public word constitutes the target for the conceptions *S* ought to come to associate with *w* over time.[269]

4.3 The informed perspective

I stipulate that the topic of a mental register doesn't alter over time[270] (such as to avoid potential self-knowledge problems that also cause epistemic problems[271]). I furthermore consider the speaker's register as connected to a *public* register (with rare exceptions). I am also saying that it is *up to the speaker* to judge whether she views herself committed to a given public word-practice (that is, a public repetition tree) when she uses her word *w*. We can clearly conceive of speakers changing their mind on whether their word *w* is to be identified with a given word *w'*. For example, I used to think that my word '$angina_I$' is the same word as the ordinary English word 'angina' and later came to realize that it is better translated as 'strep throat'. I can't permit that in such cases the topic of my mental register changes from 'angina' to 'strep throat'.[272] As a result, I am defining a point of reflection from which the speaker can view herself as committed *once and for all*. Post identification, she interprets discrepancies between the content of her own register and popular or expert opinion as one of substantial disagreement

[268] The content of a mature register is what a token *t* of that repetition tree expresses.
[269] I allowed for Mr. Anchovy's word '$lion_A$' to possibly constitute a node on the public repetition tree 'anteater'. In that case Mr. Anchovy will have to change the spelling of 'lion' to 'anteater'.
[270] I admit that splitting one register into two, or merging two registers into one is a change of topic, but it is one the individual is aware of. This is why those structural alterations do *not* lead to her affirming invalid or irrational inferences she simply cannot recognize as invalid or irrational (without corrections from the outside).
[271] Please consult chapter six for further analysis.
[272] This is guaranteed by how I define 'register'.

on a given topic. In order to attain such a stable point of reflection, I propose that the speaker judges her commitment to 'arthritis', 'tiger', 'angina', 'Madagascar' etc. *from an informed perspective*. This position is achieved by consulting a *wide selection* of reliable resources, such as experts, respected dictionaries or encyclopedias. This should allow the speaker to adequately "grasp" what the public word in question expresses. For example, she looks up 'tiger' and reads: "A very large cat with a yellow-brown coat striped with black; native to the forests of Asia but becoming increasingly rare." My suggested approach pre-supposes that the speaker understands most of the words contained in the description of what 'tiger' denotes.[273]

When a speaker commits to a public word *w* with a given register, then the tokens t_1 through t_n connected by that register express what *w* expresses. The speaker must display an *openness* to allow more knowledgeable speakers to *correct* her about what her tokens t_1 through t_n denote. In cases where she decides *not* to be persuaded against her own conceptions, she thus views herself in genuine *disagreement* about what *w* denotes. These are all signs of the attitudes of a *language consumer*.

I do rely on speakers within a linguistic community having been trained to adhere to the performance standards of how to pronounce, gesture or write words. As discussed in chapter two this can sometimes lead to productions other speakers may not instantly recognize (recall my mistaken intonation as 'reciprocal'). Still, with a little "help" of how to look at the token, they should be able to recognize the connection and how perhaps a novice performer could have come to produce *t*. Due to this comportment as a consumer, the speaker will only want to violate his status as a consumer if there are radical discrepancies between what (a sensible number and

[273] This is something I am entitled to assume, since my proposal is not meant as an analysis of linguistic competence with regard to a language, but as an analysis of competence with regard to a particular (public) word, thus taking for granted the speaker's overall competence in the respective language.

unbiased selection of the more) competent speakers believe about the referent of *w*. So, when Mr. Anchovy looks up 'lion' he won't be able to extend a practical judgment of sameness of satisfaction to the public word 'lion' (even though he may now come to acquire it as a *new* word which will require its own register distinct from the old 'lion_A'). If he now comes to realize that other members of his linguistic community call 'anteaters' what he used to call 'lions_A' and he wants to remain committed to effectively communicating with his linguistic community, then he ought to change the label of his register to 'anteater'. This is interesting development suggesting that an *already existing* register can come to be identified with a public word – this constitutes a slight deviation from the scenario Kaplan had in mind.

Some may wonder how Mr. Anchovy knows which ones are the appropriate entries to check. How would he know that 'lion' is the relevant word to look up? Mr. Anchovy will begin with those words that have a similar *form* as his 'lion_A'. Given that Mr. Anchovy is already a competent English speaker and thus meets the performance standards with respect to most other words he uses, I can assume that what he considers similar word *forms* (based on the physical characteristics) is going to largely *agree* with the judgments members of his linguistic community will make. So, he should check the entries 'lion', 'ion', 'lamb', 'lemon' and move on to increasingly more deviant word-forms… until he finds 'anteater'. Our vocabulary is finite, so this task can *in principle* be mastered.

When other public words (sometimes from other linguistic communities) compete with each other for word-sameness with a speaker's idiolectal word, it can become more difficult for the speaker to decide which word she is committed to. Examples of these were my word 'angiana_N' and Charles's word 'robin_C'. I think that in my case it wasn't reasonable to take myself as having meant *angina pectoris*, especially, since I was aware of my commitment to the German

word ‘Angina’. In Charles’s case I stressed that he had *many* interactions with British English speakers using his ‘robin$_C$’ and regarding it as a repetition of his British interlocutor’s word ‘robin$_{BE}$’, but that with the exception of a few cases (i.e. his conversation with Joe over the phone), he lacks these linguistic interactions with American speakers. If, on the other hand, Charles has a critical amount of substantive conversations involving the tokens ‘robin’ in conversations with both linguistic communities, then this may constitute a really complicated case of a merged register, which I will discuss in chapter six.

The question remains, how could Charles in principle find out which word he is using, since he could potentially look at either entry ‘robin$_{BE}$’ and ‘robin$_{AE}$’. What I am trying to say is that Charles should look at *both* entries but take into account that he has at many times intended to repeat ‘robin$_{BE}$’ and only a few times intended to repeat ‘robin$_{AE}$’. Since his register cannot be individuated by both words, he will have to choose one of the two public words in question. The way I described the case, ‘robin$_{BE}$’ is the better candidate.

Neely has serious initial reservations against speaker-centered approaches to word-competence:

> Presumably the speaker alone does not determine what beliefs are relevant; I cannot merely declare myself competent with respect to ‘tiger’ on the grounds that *I* think I have enough beliefs about tigers to qualify. […] if this were our standard for competence the notion would be useless; competence would be idiosyncratic to each speaker. In such a case, calling someone competent would be devoid of general content.[274]

What are Neely’s main concerns?

a) How does the speaker know which beliefs are relevant when it comes to competence with regard to ‘tiger’?
b) How does the speaker know that she has enough beliefs about tigers to qualify as competent in the use of the word ‘tiger’?

[274] Neely (2005): p.48.

c) The standards for competence with respect to any word (i.e. 'tiger') cannot be idiosyncratic.

Demanding that speakers (at least hypothetically) compare their own word against the corresponding public word from an informed perspective should put to rest some of Neely's concerns, since it introduces (somewhat) *uniform* standards of competence. This approach doesn't allow individuals to decide which beliefs concerning 'tiger' are relevant. When consulting reliable sources she thereby *finds out* what the appropriate experts of her community believe are the salient beliefs a competent speaker ought to have regarding what 'tiger' denotes. This responds to a) and c). It is true though that (once she has adequately consulted those resources) I leave it *up to the speaker* to decide whether her corresponding 'tiger_I'-register contains *those* community wide accepted descriptions listed and whether her register can be sincerely interpreted as containing a sufficient amount of those descriptions. When judging whether her word can be thus translated into the public language, the speaker needs to consider:

a) whether any of her own beliefs *conflict* with the encountered descriptions or publicly accepted beliefs; and
b) whether she interprets the conflict as a *substantial disagreement*, and if so, should she
    i) drop her respective beliefs because they are mistaken, or instead,
    ii) hold on to them because she regards the received view as partly mistaken on those aspects.[275]

Since speakers are taken to be competent in the language overall, which pre-supposes competence in how to participate in the practice of speaking some common language in general, I assume that speakers have been sufficiently homogenously *trained* by their linguistic

[275] This latter option must be used with great caution.

community such that from this informed perspective and with identical content in the register at hand, speakers are *not* going to make radically divergent decisions on whether they are competent in a given word. In other words, there is no Mr. Anchovy twin in some twin world (which is *identical* to our own world) who has the *same* conceptions about what his word 'lion' denotes as Mr. Anchovy, but the twin sincerely comes to the conclusion during the conversation with the vocational guidance counselor that he simply had very *strange* beliefs about lions and that he still intends to apply for a lion taming position. I simply don't find this scenario very believable. There could, however, be gray cases, where – even from an informed perspective – the conceptions contained in a register could reasonably be viewed compatible with the corresponding public concept or *not.*[276] This is where I want to leave it *up to the speaker* and would allow for inter-subjective discrepancies in judgment.

Unlike Neely, I don't think we commit ourselves to a concept or word lightly. If I decide that my word 'angina$_I$' is not the ordinary English word 'angina', then I thereby also admit that all my past conversations with English conversation partners containing the word 'angina' non-vacuously were unsuccessful – in the sense that I didn't successfully convey my beliefs concerning angina$_I$ (strep throat, that is). This judgment comes with significant costs!

Question b) is particularly interesting – what happens if the individual's tiger$_I$-register is still relatively empty and she then compares its content to the 'tiger' entry of a reliable dictionary?

[276] In cases where "experts" radically disagree about how to even *frame* a given research project – a case where one could say what the topic is still developing and hasn't found a stable equilibrium yet – it may indeed be difficult for the layman to make up her mind on whether her idiolectal word is adequately related to the corresponding public word (words such as 'ghost', 'monster', 'happiness' perhaps). I venture to argue that in such cases we all have such high variations in understanding that there is no well-defined common topic: at best the contents of our registers sufficiently overlap and we *roughly* have the same topic in mind. Even if a well-defined framework of how to talk about the phenomena in question were to be developed, it would be hard to link the layman to the corresponding public word, since she had very little basis from which to ever judge her possible commitment. In some such cases, experts may simply not exist (yet).

How does the individual know enough about her own developing register to judge whether it is compatible with the corresponding public concept? Actually, this is getting things entirely confused! This case is not at all a problematic case, but the way our word acquisitions *ought* to work. While the register is still developing, the speaker is exposed to appropriate information[277] regarding the relevant public concept *tiger* and her register thus correctly develops into one pertaining to tigers. Admittedly, it is hard to say at what point we ought to consider the speaker's *w*-register sufficiently matured (regardless of whether it is related to a public concept) such that she can be considered expressing something with *w*. When it comes to words that are *not* linked to public repetition trees, it is up to the speaker: when does she believe she is expressing something substantial with *w* – as a result of careful and sincere reflection (it shouldn't be merely as a result of a "gut" feeling)? If *w* is appropriately linked to a public repetition tree – I propose the same test: does the speaker believe to be saying something *substantial* with *w*? I simply don't find it believable that a speaker would be under the illusion of uttering something (of substance) when her register is still meagerly filled with content.[278]

Of course, speakers sometimes wrongly view themselves committed to public repetition tree *a*, but then later realize that they were really committed to repetition tree *b*. This was demonstrated by my angina thought example. I thought that my word 'angina_I' and the ordinary English word 'angina' were performances of the same word. As I later came to realize, my word 'angina_I' was better translated as 'strep throat' – it was thus *not* a node of the 'angina' public repetition tree (it was a node on the 'strep throat' public expression tree). This is why I suggest this commitment is decided from an *idealized* standpoint of evaluation at the *end* of the

[277] 'Appropriate' here is always relative to the current state of knowledge
[278] Some technical terms such as 'muon' that rarely make it into conversations of the larger public may be exceptions; it seems permissible for laymen to have very meagerly filled registers.

performance history of a given register. Nevertheless, it is a standpoint *accessible* to the speaker in principle.

Imagine King Louis II of Bavaria having traveled to Australia (which of course he didn't). His favorite animal was the swan, which he pictured white, once fully grown (since all adult swans in Bavaria, or Europe in general for that matter are white). During his visit he encounters a bird looking very much like an adult swan: a rather large "duck-ish" looking water-bird with a very long neck – but covered with *black* feathers all over. His guide informs him (through a translator) that this animal is a swan. Assuming his translator correctly transmitting the words of the guide, King Louis may trust the guide and alter his image of swans accordingly (inferring that there are white and black species). If we want to interpret King Louis's word '*Schwan*' to be the same as our word 'swan' (or '*Schwan*') today (as opposed to merely denoting a subspecies of swans) and assume the same for the Australian guide, then King Louis *ought* to maintain one single register that keeps track of his own word '*Schwan*' and the word 'swan' of the guide jointly – thereby coming to the conclusion that swans can be white or black.

It is quite possible that King Louis may have made an (initial) judgment to the contrary by holding on to the distinguishing feature that *all* (adult) swans are *white*, such that he is unwilling to extend a practical judgment of sameness of satisfaction to the word 'swan' the guide is using. So far, it seems King Louis could have decided either way. But if he had consulted an ornithologist on the matter to make a decision from an informed perspective, he would have come to realize that there are indeed black and white species of swans, thus concluding that the Australian 'swan' is the same as his word '*Schwan*'.

We can easily describe the converse situation, where an American traveler to England at first extends a practical judgment of sameness of satisfaction to the British word 'robin'. As a result,

she *falsely* keeps track of the British word 'robin' and the American word 'robin' in *one single register*. Again, I rely on her coming to see her mistake and revising her registers accordingly, if she were to consult the appropriate expert sources.

Of course, speakers sometimes change their mind on how their own words relate to those of others. I suggest that from an informed perspective the speaker commits herself to a given public word *once and for all*. Even if at some point she *disagrees* with experts on a given topic; for instance, whether global warming is caused by human activity, or whether it is simply the result of natural cycles of warming and cooling which the planet undergoes (and would undergo to the *same extent* without the existence of human beings). Even though this constitutes a *radical* disagreement, we don't view it as one that leads to a change of topic between laymen and the majority of scientists – we can imagine a person holding such beliefs and still *minimally* engaging with the public debate on global warming. This disagreement is mostly about the causes of global warming, not *what* constitutes global warming (although too many substantial disagreements on causes and effects can add up to a disagreement on what one is talking about). I rely on it *never* being the case that from an informed perspective a speaker commits herself to engaging in the public practice of global-warming talk, or what have you, but then at a later point in time radically disagrees with how to define 'global-warming' (or the respective word). Could a person who sincerely points at a snowball and utters: "This is a tiger! [279]" have at some point in the past interpreted (from an informed perspective) her own word 'tiger_{I}' and the public word 'tiger' as interchangeable? Speakers, who view themselves committed to a public linguistic practice after having consulted the appropriate sources, will rarely, if *ever*, take back this

[279] This is not meant to state that the snowball looks like a tiger, but that *literally* the snowball *is* a tiger.

commitment. Such in-decisiveness may threaten inter-subjective communication of beliefs about the external world.

Why are public words so important? Sure, they explain why Putnam can be under- or mis-informed about elms – given that Putnam deferred to speakers who know more about elms (i.e. Quine). But then how can we account for the possibility of experts being wrong or misinformed about a given topic? The arthritis patient can defer to her doctor, but whom can her doctor defer to – given how little medical experts know about arthritis? The idea of public words or concepts is supposed to *keep track* of a single *research topic* or project (i.e. arthritis or elms) such that our present experts can defer to the knowledge of future experts.[280] There is nothing objective that *makes* it the case that the arthritis patient *ought* to defer to her doctor – it has to make sense to her from her own (informed) perspective. Today's experts cannot actually obtain an *informed* perspective from which the content of their own registers looks incomplete and partly mistaken, since the information is not *yet* "out there" to be acquired. How can they view themselves committed to the way future experts are pursuing the research in question (i.e. research on arthritis)? I can propose that future medical experts view themselves as continuing the *same* research project, but that would still amount to an *external* perspective. I could, however, rely on an overlap of experts between those living today and future experts. Some of today's experts (group a) may still be alive when the *new* discoveries about arthritis are being made by future experts (group c) and will be able to interpret the new resulting definitions of what constitutes arthritis as apt. The experts that are not alive anymore (group b) are then linked to the use of the

[280] This still leaves some questions unanswered. In the counterfactual world in which experts never discovered that there are two types of jade – perhaps their chemical theories are too unsophisticated to distinguish nephrite from jadeite – is it still true that there are two kinds of jade in *that* world? In such a world there wouldn't be any future experts to defer to who have the concept of the two kinds of jade – even thought jadeite and nephrite exist both *in that world*. Then again, I don't think this is a question that a public word theory must answer.

word ‘arthritis’ of group c, via group a: since members of group b viewed themselves as investigating the same topic as group a (and vice versa) and members of group a view themselves as investigating the same topic as group c (and vice versa), group b can therefore be assumed to be causally linked to the investigation of group c such that the word ‘arthritis’ denotes the same medical condition(s) in all three groups a, b and c.

How can the novice speaker judge who is competent? How can she assess which resources are reliable, given that she is not herself an expert on the topic? That is why we have social processes that select experts. Even though we may at times pick people to fill that role who have *some* false beliefs, or individuals who may even be motivated to deliberately misinform us[281] – we nonetheless selected people who are particularly knowledgeable on the topic and who are well positioned to frame the research project by *defining* the topic and manage to keep track of it over time through progressive changes in attitudes and beliefs. I do assume that we wouldn’t trust experts that collectively lead us radically astray – such that they “feed us” beliefs that are *entirely* false of the topic to a degree that the aggregate of those deceptions render us too confused to even talk about it. It appears to me that this would entirely defeat the purpose: if they prevent us from thinking about the topic intelligibly in the first place, experts can’t deceive us anymore either.[282] Isolated cases of deliberate misinformation by single experts or a profession as a whole shouldn’t undermine my proposal (i.e. false information on whether a given artificial sweetener is carcinogenic may threaten our health, but ought not threaten our linguistic competence with regard to the name of that sweetener). Widespread deception across

[281] Neely (2005): chapter 3.
[282] Unless, of course we live in the distopia conceived by George Orwell described in *Nineteen-Eighty-Four*. In a society where history is always revised according to the whims of the governing party, without the *slightest* attempt at accuracy, but instead blatantly aiming at manufacturing public consent to whatever political agenda is being pursued – it seems the academic practice, formerly known as ‘history’ ceases to exist entirely.

a vast array of aspects regarding what a name refers to would render the novice *incapable* of judging whether her name *n* expresses the same as the public name *n'*. Although she did her best to attain the necessary information by consulting a variety of expert sources she cannot even begin to make sense of what *n'* denotes. I doubt that in such a case any layman, except for perhaps those who are in on the fraud, is still able to meaningfully talk about the topic in question. Perhaps each individual privately keeps track of their mental register without those being connected in any significant way to public discourse.

Since we know that scientific progress sometimes alters definitions gradually in the short run, but sometimes radically in the long run, it may very well be the case that a speaker if she were able to consult some of our dictionary entries today, she would be entirely at a loss. Does that entail that she is thus *not* linked to the content of that entry? This could clearly cause problems. Therefore, a speaker's informed perspective is always relative to the received knowledge of her linguistic community *at a given time* (*her* present time). For instance, a proponent of the phlogiston theory may be utterly confused if she were presented with our dictionary entry on 'combustion', such as to suspect that the word 'combustion' ought not to be confused with her own word 'combustion'. Since I want to allow for trans-theoretic terms, I want the above linguistic judgments of the phlogiston proponent to be *wrong*. All the informed perspective requires is that the laymen would identify her word 'combustion_{L}' with the word 'combustion_{EP}' used by experts (of her time). The commitment to the *present* word 'combustion_{ET}' is established throughout history by experts attempting to repeat the words of previous experts; their word 'combustion_{ET}' is a node on the same repetition tree that 'combustion_{EP}' belongs to.[283]

[283] Grice and Strawson insist on there being statements we simply cannot make (literal) sense of and therefore necessitate re-interpretation (Grice & Strawson (1956)):

(1') "My neighbor's three-year old child is an adult".

Since the speaker commits herself to repeating 'combustion$_{EP}$', she thereby commits herself to repeating 'combustion$_{ET}$' – even though she could perhaps not have actually made sense of consulting today's experts on the topic.

4.4 Comparisons

Individualists can use the conceptual tool of registers to their advantage. By allowing individual speakers to keep track of a particular topic through registers, individualists are in a position to explain why a mark *m* uttered while the respective register is relatively empty and a mark *m'* uttered while that same register is rich in content constitute performances of the same word *w*: if *m* and *m'* are appropriately connected to the same register, they constitute performances of the same word *w*. In the previous chapter I proposed that when the individual speaker intends to repeat the same word with both marks *m* and *m'*, then those marks stand for the same word (in the thick semantic sense of word-sameness); the marks in question are thus appropriately linked to the same register.

Registers also allow us to explain why Putnam may utter different words when producing 'elm$_P$' and 'beech$_P$' respectively, although both registers were stipulated to contain the same

---

Imagine the following utterance:

(2) "My neighbor's son is turning into a woman".

First sex-"change" surgeries were attempted in the 1920's (those first surgeries didn't yet change an individual's sexual organs to those of the opposite sex, they merely removed sexual organs). We can imagine how in the 1800's we would not have been able to interpret an utterance such as (2) literally. Even if one day with the right scientific advancements we were able to administer pills that speed up the cognitive development of children, so that a three-year old child has cognitive capacities comparable to a young adult – that doesn't mean that we can understand utterance (1') literally *today* such that we can conceive of somebody uttering (1') sincerely and have meant it literally today. From the perspective of our present scientific abilities (1') seems absurd – even if it turns out to be true in the future. These scientific developments shouldn't entail that the word 'adult' is a different word then from what it is today – in the sense of what it denotes.

conceptions.[284] Recall that the challenge was Putnam's confession that he attaches identical descriptions to both words.[285] If *m* and *m'* are connected to distinct registers, they constitute performances of different words. The individualist will in the end still individuate the word in terms of the *content* of the mental register. In order for words to persist throughout conceptual changes, individualists could propose that words are to be individuated in terms of the content of *mature* registers. This is admittedly merely a sketch, but since I don't intend to endorse this view, I will leave it at that.

Can individualists explain why a speaker can be under- or mis-informed about a given topic? They could claim that we are psychologically wired to adopt information from people whom we judge more knowledgeable on a given topic. This information would appear in the individual's matured register and can then be accounted for within an individualist framework. In contrast, I don't require that the individual *actually* view herself as committed to a public word. I regard it as a *hypothetical* condition: if the individual were to judge herself committed to the public word *from an informed perspective*, then she is so committed regardless of whether she actually ever invests the required time and effort into acquiring this informed perspective and then reflects on the use of her own words. The word 'arthritis' of Burge's patient denotes arthritis *regardless* of whether she goes through with the visit to her doctor during which she is told that arthritis *only* afflicts joints. What I don't want to ignore are cases where a speaker (in retrospect) *actually* judges herself *not* committed to a public word – just like I had come to realize that my word 'angina' (uttered during some visit to the dentist) wasn't committed to the ordinary English 'angina'-repetition-tree. Individualists can only account for this transmission of reference of the

[284] Please consult chapter two for the full analysis of this case.
[285] Putnam (1975, 1975): pp. 226-227.

doctor's word 'angina' to the patient's word 'angina' – if the patient *actually* makes the respective changes to the content of her mental register.

I propose to take seriously what the individual associates with her own words. At the same time, I do not suggest we individuate words in terms of contents of mature registers (unless the speaker's word fails to be connected to any public repetition tree or collections thereof). Doing so would wrongly ignore the commitment to *public words*. I demonstrated in chapter two how individualism fails to explain why Putnam's word 'elm_P' is the same word as Quine's word 'elm_Q', but that Putnam's word '$beech_P$' cannot be translated as Quine's word 'elm_Q'. Given that Putnam associates the *same conceptions* with both words, they bear *equal claim* to individuation with Quine's word 'elm_Q' (within the individualist framework). Kaplan offers an elegant solution to the problem: when Putnam utters 'elm_P' he thereby means to repeat Quine's 'elm_Q' or some other token which is a member of the same repetition tree that Quine's 'elm_Q' is a node of.

What Kaplan's as well as the individualist's account fail to properly acknowledge is that Putnam has gotten *substantial* information right! The word 'elm' and 'beech' are each respectively true of distinct entities. Putnam believes that his words 'elm_P' and '$beech_P$' denote different kinds of trees (since Putnam indeed is competent in the public words 'elm' and 'beech' we can also say that he believes that elms and beeches are different kinds of trees) – he just cannot further substantiate how they are different (without deferring to experts or consulting an expert source). For Kaplan this trivially follows from Putnam being competent in 'elm' and 'beech'. Kaplan cannot describe a case where Putnam still effectively talks about elms and beeches, but mistakes the words 'elm' and 'beech' for one and the same.

The German words '*Semmel*' and '*Brötchen*' denote the same kind of bread bun (the former word being favored in Southern-German communities, the latter being favored in Northern-German communities).[286] A competent German speaker (of either linguistic community) usually knows that:

x satisfies '*Semmel*' if and only if x satisfies '*Brötchen*'[287]

Competent speakers would not be tempted to assert something of *Semmeln* that they wouldn't also automatically assert of *Brötchen.*[288] This is something significant to have gotten right! Why couldn't Putnam have *falsely* believed that 'elm' and 'beech' behave like '*Semmel*' and '*Brötchen*'?

Neither Kaplan nor individualists can explain the mistake that occurs if Putnam were to think that 'elm_P' and '$beech_P$' constitute performances of *the same word* (that is, if he had two alternative labels for one and the same mental register – where he ought to have one register for each label). Individualism rejects the notion of public words, so they entirely lack the resources to relate registers to public words.

According to Kaplan, speakers are (almost) automatically competent in the use of a word (in the thick sense) once they are exposed to one of its tokens and attempt to repeat that token. Bach explains in detail how this transfer of reference works:

> Now if the speaker is thinking of something by name, he is entertaining a mental token of the name; when he refers to it by name, he produces a physical token of that name; and the audience, upon hearing that token, forms a mental token of the same name, which he can then retain in memory. Since the hearer's mental token of the name 'inherits' the same object as the speaker's, the object of the hearer's thought is determined relationally, not satisfactionally. […]

[286] Those words behave somewhat like 'soda' and 'pop' in American English.
[287] This is of course not how a competent speaker would put it, unless she had the relevant philosophical training. We could, however, ask her the right kind of questions from which we can infer that she holds this bi-conditional to be true.
[288] That doesn't mean it is advisable to go into a Bavarian bakery ordering "*Brötchen*" – it seriously violates linguistic conventions to do so.

> [The] reference of a [proper name] is determined not by its meaning but by its ancestry. It plays this role by being of a certain form (sound or shape), generally the same as the one to which it is linked. […] Since the token of a name represents in virtue of its form, not its meaning, its representational features can be perceived by the hearer, who can then and thereafter use mental tokens of the same name to think of (or refer to) the same object.[289,290]

Kaplan asserts that with the appropriate intentions reference of a word can be inherited from the word of the competent speaker – as just described by Bach. Even if Putnam were to use 'elm_P' and '$beech_P$' as labels of *one and the same register*, Kaplan has difficulties explaining how Putnam could ever fail to be competent in 'elm' and 'beech' (since Putnam intends to repeat the public word 'elm' with his 'elm_P' and the public word 'beech' with his own word '$beech_P$'), unless Putnam had seriously faulty referential intentions[291] (which he does not[292]). If Kaplan remains faithful to his own proposal, he can't deny Putnam membership in the respective public repetition trees – even though something has gone *seriously* wrong. He therefore simply

---

[289] Bach (1987): pp. 32-33. It is somewhat ironic that this immediate inheritance of reference is popular when it comes to proper names. Come to think of it, conflicting referential intentions can occur much more immediately when they pertain to proper names. The proper extension of kind terms is always somewhat in question. Not only can the laymen speakers be mistaken about that, even experts can be wrong about the extension of a word. When it comes to proper names (with the exception of possibly split and merged registers) the speaker knows at least that her name is true of one individual object only. Imagine my neighbor points across the street and tells me: "Oh look, Patrick Espinoza is shoveling snow in front of his driveway". As it turns out my neighbor is shortsighted and the person shoveling snow is somebody else. Who does my memory – "Patrick Espinoza was shoveling snow in front of driveway x" – pertain to? I clearly have the referential intention to refer to the man I saw *that day* (whoever he was), even though at the same time, I intend to repeat the name my neighbor uttered (but his name 'Patrick Espinoza' denotes Patrick Espinoza). It seems to me that when it comes to proper names it is much easier to fill the corresponding register with conceptions violating the transmission of reference – all it takes is one direct perception of the object one relates with the name in question that conflicts with the reference of one's interlocutors name.

[290] Although Kaplan believes that with the appropriate intentions, reference of kind terms and proper names transfer thus, Bach only agrees this to be true for proper names.

[291] This is Kaplan's exception clause, which I showed to be so strong that it threatens to undermine anti-individualism altogether.

[292] Although Putnam intends to refer to what 'beech' denotes with 'elm' (and vice versa) this confusion doesn't strike me as grounds to deny somebody membership in the 'elm'-speaking practice (or 'beech' speaking practice for that matter). We can make sense of somebody wondering whether there is any difference between beech and elm trees in a way that we simply cannot make sense of somebody wondering whether there is a difference between snowballs and tigers.

stipulates that if somebody who believes one of her own words *w* to be linked to two distinct public repetition trees, then "nothing whatsoever is being said" by *w*:

> I am inclined to think that when two different common currency words are wired together in this way in a given black box, which then pulls from that common source and transmits, nothing whatsoever is being said. Is it transmitting the first word? Is it transmitting the second word? I think there is just no answer to that question.[293]

According to Kaplan, the word of the individual must either be linked to the first public word or the second – he cannot conceive of it being its own muddled entity (somehow individuated by both public concepts). Kaplan is correct about us not being able to straightforwardly translate this muddled word into the public language.

This is how I understand the case: Imagine the individual has a register labeled 'jade'. The public language contains the word 'jadeite' and 'nephrite', but counterfactually nothing equivalent to our actual word 'jade'. The individual's word behaves like our actual word 'jade' that I stipulated non-existent in the public language of the counterfactual world – it denotes nephrite and jadeite (without distinguishing between the two distinct chemical compounds). In the counterfactual public language 'jade' cannot be translated without some paraphrasing reservations.

> x satisfies 'jade' if and only if, x satisfies 'nephrite' or x satisfies 'jadeite' (without the ability to distinguish nephrite from jadeite).

I don't see why our actual word 'jade' should fare any better than those words that have muddled or merged public concepts. Since we view assertions such as: "You are wearing a beautiful jade ring" as expressing something that makes sense, we should accept that muddled words – muddled in the sense that they merge two separate public repetition trees) – say something

[293] Kaplan (1990): p. 107.

meaningful as well. I will discuss the implications of anti-individualism on discursive justification further in chapter six.

When it comes to the flip-case, a speaker mistakes one name for two, Kaplan surprisingly judges differently. Here, the speaker's way of storing the respective words doesn't interfere with competence in the respective word:

> A transmission of a word may come into the black box and then another transmission of the same word may come into the black box, and inside the black box this very same word might get stored in two different locations, *stored*, so to speak, as the transmissions of two different words. […]
>
> Somehow, in the black box, the different branches of the same word (i.e. the different input utterances of the same word) weren't all properly linked together. And that means that when the black box emits, it thinks it has two words, and it will make a choice as to *which* 'word' it is going to emit.[294]

Kaplan seems to give precedence to the success of repeating a public word over successfully judging whether the speaker's own utterances constitute performances of the same word. The speaker, Peter, is exposed to an utterance of the public name 'Paderewski$_{P1}$', he then *repeats* that utterance and uses it as a label of one of his registers: 'Paderewski$_{I1}$'. At a different occasion Peter is exposed to another utterance 'Paderewski$_{P2}$'. Again, he *repeats* that utterance and uses it as a label of one of his registers: 'Paderewski$_{I2}$'. Since the individual believes the utterances 'Paderewski$_{P1}$' and 'Paderewski$_{P2}$' to constitute performances of different (common currency) names – that is, he understands them to each denote a different individual – he keeps track of what they denote in *distinct* registers. This is where the individual is mistaken, the utterances 'Paderewski$_{P1}$' and 'Paderewski$_{P2}$' constitute performances of the *same* common currency names and denote one and the same individual.

Kaplan wants to keep the identities between the words 'Paderewski$_{I1}$' and 'Paderewski$_{P1}$' intact, that is, maintain that they constitute performances of the same name. The same goes for the pair 'Paderewski$_{I2}$' and 'Paderewski$_{P2}$' – given his criterion of repetition, Kaplan does not

[294] Kaplan (1990): pp. 106-107.

conceive of an individual ever failing to repeat a public word or name (except for those cases where referential intentions seriously conflict with the reference of the interlocutor's word and cases where public concepts are merged). As the Paderewski puzzle goes it is assumed that 'Paderewski$_{P1}$' and 'Paderewski$_{P2}$' are utterances of the *same* name. This means that the individual is *wrong* about thinking that 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' are utterances of different names. It seems rather odd, to say the least, that we always succeed in repeating words produced by other speakers, but shall be able to *fail repeating our own words*. Kaplan may object that when Peter intends to repeat 'Paderewski$_{I1}$' he doesn't fail to do so, but rather, doesn't notice that he is *also* repeating 'Paderewski$_{I2}$'. Since he believes that 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' constitute different names (in the sense that they refer to different individuals), he explicitly means *not* to repeat 'Paderewski$_{I2}$' when he repeats 'Paderewski$_{I1}$'. In other words, according to Kaplan, the individual *cannot* keep track of the identities (and non-identities) of her own words.

This is a dangerous failure to make space for in any semantic theory. The ability to repeat a public word, such as 'tiger' is only meaningful, if the speaker knows how to keep her corresponding 'tiger$_I$'-register *separate* from registers of other words (i.e. 'anteater'). We already know that she cannot do that in terms of similarity of content of the registers or else, for Putnam 'elm$_P$' and 'beech$_P$' are both alternative spellings of one and the same word. I think allowing an individual to fail in her intentions to repeat her own words, creates a theoretic "mess". I think Kaplan shouldn't give precedence to inter-subjective word-individuation over intra-subjective word-individuation. By introducing this possibility of doubt towards a speaker's ability to keep track of her own word-identities, he implodes his own proposal. If individuals cannot keep track of their own words in problematic cases, why should they do any better in the

"normal" cases (where they do *not* confuse two words to be one, or one word to be two)? The individual cannot distinguish the "normal" from the problematic cases *from her own perspective*! It seems to me that Kaplan's proposal actually relies on our ability to at the very least keep track of our own words. Wasn't this an unquestioned assumption we made all along? Since the supposed false individuation of words depends on factors that she is *not aware of* (unless she is being corrected from the outside), she is incapable of correcting such "mistakes".[295]

I will argue, instead, that Peter correctly judges his own words 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' to represent *different* words. He split one public name into two separate registers. Fiengo and May ascribe the individual the same linguistic control:

> We will find that the *beliefs* of individuals will be critical in the statement of those identity conditions on words: two occurrences of words are occurrences of the same word for an individual if and only if that individual believes they are the same.[296]

This constitutes a failure to repeat the public word 'Paderewski$_{P1}$' when he utters: 'Paderewski$_{I1}$'. Nevertheless, *both* registers are *linked* to one and the same public word 'Paderewski$_{P}$' (of which 'Paderewski$_{P1}$' and 'Paderewski$_{P2}$' are instances. This link is not merely assumed in principle, but can be tested: if the speaker finds out that there is only one Paderewski and that 'Paderewski$_{P1}$' and 'Paderewski$_{P2}$' denote that Paderewski – will he as a result of that information *merge* his registers 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' into one register? If so, then her words 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' were all along both linked to 'Paderewski$_{P}$' – without her having been *adequately* competent in 'Paderewski$_{P}$' (since this pre-

[295] I am puzzled by how deflationary accounts of meaning are supposed to work, unless we assume that speakers can express the meanings of their own words as in the following example: "'robin' is true of x if and only if x is robin". If external factors (i.e. switches between linguistic communities) could bring it about that between the first token "water" and the second token 'water' the speaker looses track of how to individuate her own tokens, then the utterance could be *false*. Unless I can tell whether I switched linguistic communities (which in the Twin Earth thought experiments one can't) I wouldn't be able to distinguish the unproblematic or true utterances from the false ones.

[296] Fiengo & May (2006): p. 53.

supposes that he doesn't "store" *it* under distinct registers). Can a word *w* be linked to a public repetition tree without *w* constituting a repetition of the corresponding public word? Remember how I argued earlier that somebody who does not yet know anything about tigers may utter the word *form* 'tiger' in a question such as "what is a tiger?" without yet being competent in the use of that word. In that case the quasi-word 'tiger' is linked to the public word tiger without it constituting a performance of that public word.

When Peter realizes that his words 'Paderewski_{I1}' and 'Paderewski_{I2}' denote the same person, he will have to merge his two registers into one 'Paderewski_{I}' – or else he may in the future be tempted to assert contradictions. It is this *new* merged register that he then identifies with the public word 'Paderewski_{P}'. Since the conceptions associated with both registers 'Paderewski_{I1}' and 'Paderewski_{I2}' get transferred to the *new* 'Paderewski_{I}' register, there is a semantic *continuity* between those registers. This semantic continuity does not pre-suppose semantic identity. It is this relationship of *continuity* that establishes a *semantic link* between 'Paderewski_{I1}' and 'Paderewski_{I2}' and 'Paderewski_{P}' (but again, one that is weaker than identity). It is not merely because 'Paderewski_{P}' and 'Paderewski_{I1}' *overlap* in content (that is, Peter reads and hears many things about whom other people call 'Paderewski_{P}' that he also believes to be true of 'Paderewski_{I1}') that establishes this connection. This overlap is a *result*, not a cause of his semantic commitment. It is the identity between his later *merged* register 'Paderewski_{I}' (that arose from registers 'Paderewski_{I1}' and 'Paderewski_{I2}') and 'Paderewski_{P}' that establishes that link. The realization that 'Paderewski_{I1}' and 'Paderewski_{I2}' ought to be merged represents a *substantial change* and will affect his inferential behavior. I will discuss Kripke's Paderewski puzzle in greater detail in chapter five.

The abrupt acquisition of competence, endorsed by Kaplan, brings with it significant explanatory shortcomings:

(i) An inability to account for deference failures: a speaker attempts to repeat a public word but fails to do so and thus fails to defer to other competent users of the word in question.

(ii) We cannot make sense of utterances that merge two separate public concepts[297] unless a corresponding public word exists (i.e. 'jade' is true of nephrite and jadeite).

(iii) We cannot retrace inferences that involve words linked to registers that *split* two separate public concepts[298] without ascribing irrational beliefs to the individual (which she, however, is incapable of recognizing as such without input from other speakers).

Does the speaker believe that her own marks *m* and *m'* constitute performances of the same or different words (in the thick sense of 'word')? The content of the registers are the basis on which the speaker judges whether with uttering her word *w* she is committed to a given public word *w'*. The speaker will also take into account memory of past conversations she had during which she used *w* and how this possible individuation via *w'* will affect the truth-value of her past assertions. If this commitment to some public repetition tree can be established (in actuality or hypothetically), then the content of her *w*-register doesn't (and never did) individuate *w*: *w* and all past productions of *w* are individuated via the public word *w'*.

This is how my register-based version of anti-individualism compares to Kaplan's anti-individualism and register-based individualism in its explanatory power:

[297] This is what causes self-knowledge problems and will be analyzed in chapter six.
[298] This is what happens in the Paderewski puzzle, which I will discuss in the next chapter.

Corresponding thought example	Anti-individualism without registers	Anti-individualism with registers	Individualism with registers
Arthritis patient	+ Under- or misinformation about what one's own word denotes possible	+ Under- or misinformation about what one's own word denotes possible	- Under- or misinformation about what one's own word denotes *not* possible
Mr. Anchovy	- Immediate acquisition of competence with regard to a word	+ Acquisition of competence with regard to a word gradual	+ Acquisition of competence with regard to a word gradual
Marco Polo & dental visit	- Involuntary deference failure not possible	+ Involuntary deference failure possible	[Not applicable]
Paderewski puzzle	- Ascriptions of irrational inferences to an otherwise rational agent possible[299]	+ Ascriptions of irrational inferences to an otherwise rational agent impossible	+ Ascriptions of irrational inferences to an otherwise rational agent impossible
Self-knowledge puzzles[300]	- Utterances where the individual believes two distinct public words to be one are nonsensical.	+ Utterances where the individual believes two distinct public words to be one are meaningful, but are difficult to translate into the public language	[Not applicable]
	Individual variations in understanding due to different beliefs	Individual variations in understanding due to different beliefs and register structure	Individual variations in understanding due to different content associated with registers

[299] I have in mind here the kind of irrational inferences that the individual cannot recognize as such from her first-person perspective without external input – she is unable to make the apparently necessary logical corrections on her own.
[300] Those will be discussed in the following chapters.

4.5 Conclusion

I aimed at developing a position that occupies the middle ground between social anti-individualism and the kind of individualism that individuates a speaker's words without input from the speaker's linguistic community. According to my proposal, the content of the mental register doesn't by itself individuate the speaker's word '$lion_I$', although it cannot be entirely ignored either when it comes to individuating '$lion_I$'. We consult the mental register's content when it comes to deciding whether it is linked to any public repetition tree and if so, which one. I don't regard it as impossible that a person's register may simply lack an appropriate corresponding public repetition tree it can "latch on" to (i.e. I may have a register that denotes all and only philosophy books that some recluse spider does or used to hide in *and* be under the impression that the label of my register is a word used in my linguistic community). Nevertheless, I view most words as allowing for a translation into the public language that doesn't require paraphrasing. The main advantage of individualism is that it can easily explain away invalid and contradictory beliefs that under certain conditions anti-individualism cannot help but ascribe to speakers.[301] In my own proposal, I want to hold on to the social anti-individualist idea that speakers can be under- or mis-informed about a given topic – but do so without threatening the *a priority* of their logical reasoning capacities.

[301] The speaker is unable to recognize those beliefs as contradictory or invalid even after careful reflection.

CHAPTER 5

ARE INDIVIDUAL SPEAKERS RELIABLE WHEN IT COMES TO INDIVIDUATING THEIR OWN WORDS? TEST-CASE: KRIPKE'S PADEREWSKI PUZZLE

5.0 Introduction

Social anti-individualism *à la* Burge and Kaplan is driven essentially by epistemic concerns. The individual's acquisition of knowledge is in part structured by other members of her linguistic community. We can't be experts in all fields relevant to our research endeavors. Kripke's Paderewski puzzle exploits the resulting epistemic gap between how the speaker *perceives* the reference of the name 'Paderewski' and its *actual* reference to ascribe (simultaneously) a set of contradictory beliefs to a rational speaker. In order to circumvent such threats to the individual's logical abilities, I attribute maximum linguistic control to the individual over her own word individuations. I will reveal the fine-grained individuation of belief ascriptions on the basis of individual conceptions as misguided, since they hinder the proper transfer of justification from testimony T to the belief τ formed through T. If the topics of T and τ only ever partially overlap – unless all conceptions are shared between testifier and audience – then T can never be fully understood by the hearer.[302] I will apply the notion of mental registers[303] in a new way such that

[302] My own solution to the Paderewski case doesn't lead to this result. More is required between 'Paderewski_{I1}' and 'Paderewski_{P}' for communication to succeed (somewhat). This communication is limited because it doesn't allow for a transfer of *knowledge*. It is Peter's merger of his two 'Paderewski_{I1}' and 'Paderewski_{I2}' registers that demonstrates that he was all along committed to the public 'Paderewski_{P}'. This is why I stress that the two 'Paderewski_{I1}' and 'Paderewski_{I2}' constitute a split *public* concept – not merely a split idiosyncratic concept. This kind of commitment is lacking towards the British 'robin_{BE}' for speakers who employ the American 'robin_{AE}' – even though there is an overlap of register content this does not suffice for the kind of communicative success that I have in mind. Thanks to Ebbs for pressing the concern I address here.

[303] In the relevant literature, mental registers are sometimes referred to as 'dossiers' or 'files' instead.

our belief ascriptions of the described case are rendered compatible with epistemic concerns that originally drew us to (social) anti-individualism.

5.1 Kripke's Paderewski puzzle applied to social anti-individualism

Peter believes (falsely) that there are two famous individuals known as 'Paderewski' – one a celebrated pianist, the other a prominent Polish statesman. As it turns out, the pianist named 'Paderewski' and the politician called 'Paderewski' are actually one and the same person. If prompted, Peter would assert:

(i) Paderewski has musical talent.

(ii) It is not the case that Paderewski has musical talent.[304]

(i) is uttered as a statement about Paderewski, *the pianist*, while (ii) is uttered with Paderewski *the politician* in mind. Given that the two occurrences of 'Paderewski' are instances of the *same* (public) proper name, we ascribe contradictory beliefs to Peter, merely by connecting (i) and (ii) through a conjunction:[305, 306]

(iii) Peter believes that Paderewski has musical talent and that it is *not* the case that Paderewski has musical talent.

Furthermore, we as *observers* are maintaining a contradictory thought:[307]

---

[304] The same standards for musical talent are assumed throughout the paper.

[305] Assuming that the substitution of instances of one name with instances of a different but co-referential name is truth-preserving represents a *substantial* assumption. It is trivial that the substitution among instances of the *same* name preserve truth-values.

[306] Usually the Paderewski puzzle is interpreted as an attack on direct reference accounts of meaning. Owens (1995) agrees with me that anti-individualistic intutitions provide all the motivation we need for the puzzle to arise.

[307] (iv) follows from (iii) if a strong disquotational principle is presupposed.

There is also a strengthened 'biconditional' form of the disquotation principle, where once again any appropriate English sentence may replace 'p' throughout: *A normal English speaker who is not reticent will be disposed to sincere reflective assent to 'p' if and only if he believes that p.*

(iv) Peter believes that Paderewski has musical talent and it is not the case that Peter believes that Paderewski has musical talent.

Peter is oblivious that the content stated in (iii) asserts a contradiction – he is incapable to *see* that he must choose between (i) and (ii). From Peter's first-person "perspective" (i) and (ii) neither agree nor disagree in content. Unless some serious neurological condition is inhibiting the proper function of an individual's cognitive apparatus, we assume that individual to think rationally. We cannot comprehend how somebody who understands some statement of the form $\lceil \alpha \neq \alpha \rceil$ – that is, understands what it asserts, could sincerely believe it to be true. Taschek explains:

> [T]here is no making sense of persons, in the relevant sense, without the possibility of assessing their rationality. This means at the very least, being able to assess their beliefs for logical consistency, to evaluate their inferences for validity, and to identify intersubjective agreements and disagreements between their beliefs and those of others.[308]

In Peter's case, social anti-individualism *à la* Burge and Kaplan *prima facie* conflicts with the *a priority* of logical reasoning. According to Kaplan, given that Peter intends to repeat the public word 'Paderewski$_P$' in (i) as well as (ii), both assertions contain the *same* name without Peter noticing that he thereby produced two performances of the same name. Thus, there is no getting around ascribing to Peter a contradictory belief.

Owens argues that ascribing "permissible" contradictions to a person is compatible with her rationality. He explains that a

> fully rational subject may subscribe to contradictory beliefs for the simple reason that *belief contents* are, in part, a function of contextual factors, and so recognition of sameness and difference among such contents may require much more than mere rationality.[309]

Kripke (1976): p. 249.

[308] Taschek (1998): p. 323.

[309] Owens (1989): p. 311; see also Owens (1995): p. 251.

Owens simply denies that Kripke's thought-example constitutes a puzzle. Unless we can present a way to contrast permissible from non-permissible contradictions, it seems to me that Owens merely stipulates that rational individuals simply don't assert non-permissible contradictions. I think that we can learn more about language from Kripke's puzzle, if we attempt to avoid ascribing contradictory beliefs to Peter.

To render the original puzzle convincing to the reader, allow me to flesh out the circumstances. When Peter listens to a CD he received as a gift, he is pleasantly surprised by the merit of the performance. Curious about the person playing, he studies the label and reads 'Paderewski' – listed under 'pianist'. This is how he comes to hold (i). While listening to the recording he is also reading Paderewski's biography, forwarded to him by the friend who gave him the very CD capturing his admiration. Since the recording turns out to be so enjoyable, however, Peter's attention oscillates between music and text. He reads the beginning section, where Paderewski's early interest in music is mentioned, as well as his teacher's assessment of Paderewski's insufficient talent for a professional career in music. Peter then *skips* the sections, where Paderewski's persistence towards his aspirations and his resulting fame in the musical world are recorded. Peter's attention returns to the text for the parts describing his political role in Poland's reunification after WWI as its elected prime minister. Peter thus concludes (ii).

There are strong reasons to believe that Peter learned the *public* name 'Paderewski$_P$'. He acquired the *public* word 'Paderewski' – from *separate* public sources: (a) the one printed on a CD; (b) the one used in a biography. Those two occurrences are instances of one and the *same* name. Peter intends to use the same name that is printed on the CD when uttering (i) and he intends to repeat the name that is used in the biography when he declares (ii) – those names just happen to be identical. He certainly doesn't deliberately opt out of the 'Paderewski'-linguistic

practice. He wants to communicate his thoughts *effectively* when he utters 'Paderewski' in both instances (i) and (ii). Peter holds substantial beliefs, that would be *true* if we were entitled to interpret his word 'Paderewski' as our word 'Paderewski'.

Without explicitly stating it, Peter's friend clearly attempts to convey to him that Paderewski the famous pianist and Paderewski the famous Polish statesman are one and the same individual – thereby drawing attention to Paderewski's extraordinary life. Did Peter learn anything from the CD or the biography? It seems natural to say he has *learned* from the biography that, amongst other things, Paderewski was a politician. Information imparted to Peter via the *public* name 'Paderewski'. Does this imply that Peter understands the public name 'Paderewski' (and thus is entitled to use it)? We need an account of names that allows Peter to "*understand*" our name 'Paderewski' *without* thereby entitling his use of it. That is, without identifying his utterances of "Paderewski" with the public name 'Paderewski' such that (iii) must be viewed as an instance of the schema $\lceil \alpha \neq \alpha \rceil$.

If Peter were to claim:

(v) As a young man Paderewski didn't have red hair.

we would consider ourselves to *disagree* with him, since Paderewski was known for his vibrant red hair. We would naturally be tempted to take Peter's words in (v) to express what we would express with (v) – a proposition which we consider to conflict with our belief that as a young man Paderewski did have red hair. These observations speak for Peter's name 'Paderewski' in (i) and (ii) to be identical with the public name 'Paderewski$_P$'. Does the successful transfer of information go hand-in-hand with word-sameness of the words in testimony T and the belief τ acquired as a result of T?

Some considerations, however, clash with our assumption that Peter uses the public name 'Paderewski'. If we were to tell Peter:

(vi) Paderewski had red hair.

he would probably ask us "which one?". In comparison, it would never appear to *us* to raise such a question. Isn't it blatantly obvious? Peter believes that Paderewski, the pianist, is not Paderewski the politician – and that's that! However, if we translate that into our public language, this is a longer version for:

(vii) Paderewski is not Paderewski.

(vii) is an instance of the schema:

(vii') $\alpha \neq \alpha$

It is very difficult to comprehend how somebody who understands some statement of the form $\lceil \alpha \neq \alpha \rceil$ – that is, understands what it asserts, could sincerely believe it to be true.

Upon reflection, the *prima facie* evident explanation is not sensible in our own words. Imagining an attempt to correct Peter's beliefs: how would *our* utterance 'Paderewski is Paderewski' – an instance of the schema $\lceil \alpha = \alpha \rceil$ – be informative? [310]

In order to facilitate the discussion, let's distinguish the occurrence of 'Paderewski' in (i) and (ii) as 'Paderweski_{I1}' and 'Paderewski_{I2}' respectively. Recall, that it is part of the description of Kripke's case that Peter's names 'Paderweski_{I1}' and 'Paderewski_{I2}' denote the same individual, namely Paderewski (the politician and statesman). If Peter were to claim about some person

[310] In his opening paragraph to "On Sense and Nominatum," Frege writes:

> "a = a" and "a = b" are sentences of obviously different cognitive significance: "a = a" is valid a priori and according to Kant is to be called analytic, whereas sentences of the form "a = b" often contain very valuable extensions of our knowledge and cannot always be justified in an a priori manner".

Frege (1892).

called ‘Paderewski’ that he_1 has musical talent and of some *other* person called ‘Paderewski’ that he_2 has no musical talent and $he_1 \neq he_2$ – the puzzle wouldn’t even arise!

We are facing a dilemma: we must individuate Peter’s names – ‘$Paderweski_{I1}$’ and ‘$Paderewski_2$’ – such that they (a) co-refer with one another in order to avoid ascribing an overt contradictory belief to Peter and (b) co-refer with the public name ‘$Paderewski_P$’ such as to allow for the transfer of knowledge and successful inter-subjective communication. Since I wanted to be cautious about the substitution of co-referential terms, we didn’t thereby prove that (i) and (ii) contradict.

When we are attempting to solve Kripke’s puzzle, we are trying to do justice to two seemingly opposing explanatory goals:

a) Account for the successful communicative exchange between Peter and his linguistic community with regard to utterances containing the name ‘Paderewski’ (non-vacuously).

b) Avoid ascribing to Peter contradictory beliefs of the form $\lceil p \wedge \neg p \rceil$

While b) tempts us to conceptualize Peter’s contents of belief as entities that *change* in accordance to Peter’s ways of perceiving the referent of the name ‘Paderewski’, we would thereby drive a wedge between the content of his belief *c* and the semantic content of Peter’s utterances that are meant to convey *c*. Unless the audience *A* shares Peter’s ways of perceiving the referent of ‘Paderewski’, some element of *c* remains inaccessible to the audience. Of course we generally take for granted that a name doesn’t communicate all our attitudes towards a given referent to *A*. These attitudes are generally overlooked, since we don’t take them to have logical implications.[311]

[311] Even though these claims are reminiscent of a Fregean notion of *Sinn*, I want something much weaker in its effect than *Sinn*. Fregean approaches to this problem are struggling accounting for a speaker S_1 who can only think of Paderewski under the mode of pianist can successfully communicate with another

Perhaps we could argue that there are in fact two public words 'Paderewski', since the man in question is known for achievements in different fields. One could easily imagine how most people tend to keep up with happenings in one field but not both. The resulting one-sided and partial conceptions may have led to distinct public senses (some kind of non-referential contribution to the meaning of a name).[312] Consider an utterance made by a speaker informed of Paderewski's political career but not his musical fame:

(viii) Paderewski has red hair.

While the following one is made by a speaker who knows about Paderewski's contributions to music, but not his political ones:

(ix) Paderewski has red hair.

Even if we want to argue that (viii) and (ix) do *not* make the same assertion (thereby allowing us to argue that (i) and (ii) are not contradictory) – which I find highly dubious – we wouldn't want for *every* possible conception to create a *new* (public) name 'Paderewski'. Disregarding for a moment the original epistemic motivations of anti-individualism, this sketched solution won't dissolve the puzzle. In theory every imaginable conception regarding Paderewski can open up the possibility for Kripke's puzzle. Peter, for some reason or other, may falsely believe that there are two famous politicians who happen to have the same (generic) name 'Paderewski'. He believes that one of them is passionate about the piano in his spare time, while the other one doesn't care much for music. He thus attributes musical talent to one, but not the other – even though there is only one public name for Paderewski, the politician.

---

speaker S_2 who only thinks of Paderewski under the mode of politician. In my own account, nothing speaks against those individuals both sharing the same public word 'Paderewski$_P$' and hence communicating successfully.

[312] I have to thank Steve Wagner for encouraging me to consider this view in the present context. Perhaps this is why the Superman and Clark Kent thought example is so convincing. We attach different senses to Clark Kent in the persona of Clark Kent compared to the persona Superman.

5.2 Kaplan's analysis of the puzzle

In chapter three I discussed Kaplan's exception clause that explains how a speaker can fail in her attempt to repeat a public word. In this example, Peter doesn't seem to hold any significantly false beliefs that the referential use exception clause could exploit (such as to create a conflict between referential intentions and the intention to repeat the public word 'Paderewski$_P$'). Of course, Peter holds the false belief that Paderewski the pianist is not the same individual as Paderewski the politician. But as I explained above, this is not merely some false belief; it is one that cannot be true under any circumstances. Kaplan cannot get around ascribing Peter a contradictory belief of the form $\lceil p \wedge -p \rceil$ but he is perhaps in a position to contrast the kind of contradiction Peter asserts against those that we would deem *unacceptable*.

Kaplan's analysis of the puzzle is as follows:

> A transmission of a word may come into the black box and then another transmission of the same word may come into the black box, and inside the black box this very same word might get stored in two different locations, *stored*, so to speak, as the transmissions of two different words. […]
>
> Somehow, in the black box, the different branches of the same word (i.e. the different input utterances of the same word) weren't all properly linked together. And that means that when the black box emits, it thinks it has two words, and it will make a choice as to *which* 'word' it is going to emit.[313]

He seems to give precedence to the speaker successfully repeating a public word over her successfully repeating her own words (over time). He accuses Peter of not being able to keep track of his own utterances that constitute performances of the same word: "the different branches of the same word […] weren't properly linked together". Since identity between public and idiolectal word is assumed, this means that Peter repeats the name 'Paderewski$_{I1}$' when uttering 'Paderewski$_{I2}$' – even though he does *not* intend for these two performances to constitute repetitions of the same name (in the thick sense of *name*).

[313] Kaplan (1990): pp. 106-107.

Peter is exposed to an utterance of the public name 'Paderewski$_{P1}$', he then *repeats* that utterance and stores it as a label of one of his registers: 'Paderewski$_{I1}$'. At a different occasion Peter is exposed to another utterance 'Paderewski$_{P2}$'. Again, he *repeats* that utterance and stores it as a label of one of his registers: 'Paderewski$_{I2}$'. Since Peter believes the utterances 'Paderewski$_{P1}$' and 'Paderewski$_{P2}$' to constitute performances of different (common currency) names – that is, he understands them to each denote a different individual – he keeps track of what they denote in distinct registers. This is where Peter is mistaken, the utterances 'Paderewski$_{P1}$' and 'Paderewski$_{P2}$' constitute performances of the *same* common currency names and denote one and the same individual. Peter is "taking one name to be two".[314]

Mercier explains how sometimes speakers fail to have a well-formed intention when attempting to repeat *a* word, because the various utterances the speaker groups together as *one* word are in fact utterances of several words: "Now, in order for [a spealer's] intention to use a word with *its* meaning to be a well-formed intention, it has to be the case that *the* word about which [she has] that intention *is* indeed *one* word."[315]

Peter commits the converse of the mistake Mercier describes. He groups utterances into distinct word-categories, when on the contrary those utterances are all performances of the *same* word. Does this mean that Peter succeeds in his attempt to repeat the "names" in question?

Kaplan wants to keep the identities between the words 'Paderewski$_{I1}$' and 'Paderewski$_{P1}$' intact, that is, maintain that they constitute performances of the same name. The same goes for the pair 'Paderewski$_{I2}$' and 'Paderewski$_{P2}$' – which is partly a result of Kaplan not conceiving of an individual ever failing to repeat a public word or name (except for those cases where referential intentions seriously conflict with the reference of the interlocutor's word and cases

[314] Kaplan (1990): p.108.
[315] Mercier (1993): p. 79.

where public concepts are merged[316]). As the Paderewski puzzle goes, it is assumed that 'Paderewski$_{P1}$' and 'Paderewski$_{P2}$' are utterances of the same name – "this very same word might get stored in two different locations." This means that Peter is *wrong* about thinking that 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' are utterances of different names. It seems rather odd, to say the least, that we always succeed in repeating words produced by other speakers, but shall be able to *fail repeating our own words*. Kaplan may object that when Peter intends to repeat 'Paderewski$_{I1}$' he doesn't fail to do so, but rather, doesn't notice that he is *also* repeating 'Paderewski$_{I2}$'. Since he believes that 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' constitute different names (in the sense that they refer to different individuals), he explicitly means *not* to repeat 'Paderewski$_{I2}$' when he repeats 'Paderewski$_{I1}$'. In other words, according to Kaplan, Peter *cannot* keep track of the identities (and non-identities) of his own words.

This is a dangerous failure to make space for in any semantic theory. The ability to repeat a public word, such as 'lion' is only meaningful, if the speaker knows how to keep her corresponding 'lion$_I$'-register *separate* from registers of other words (i.e. 'anteater'). We already know that she cannot do that in terms of similarity of content of the registers or else, for Putnam 'elm$_P$' and 'beech$_P$' are both alternative spellings of one and the same word. The alternative is to have the public concepts keep track of the mental registers *for* the speaker. But then we are forced to ascribe contradictions to the speaker – even though he is assumed rational. Shouldn't a statement such as 'Sarah is Sarah' be known *a priori* by the speaker if she uses the same proper name (in the thick sense) twice to assert the identity statement? In the case where there are identical twins both called 'Sarah' and the speaker mistakes them for one and the same person, she would fail to repeat proper names (since each of her assertions 'Sarah' picks out both girls).

[316] See chapter three for a more detailed discussion.

According to Kaplan the public repetition trees '$Sarah_1$' and '$Sarah_2$' would somehow keep track of the performances for the speaker, let's call him 'Joe', such as to keep them separate. But then which conceptions does he associate with what '$Sarah_1$' denotes and which ones does he attribute to what '$Sarah_2$' denotes? Let's say $Sarah_1$ plays the piano and $Sarah_2$ plays the guitar. Imagine that at different occasions Joe asserts: "Sarah plays the piano" and "Sarah plays the guitar." If the public repetition trees do the work of individuating Joe's words, then he would *not* actually assert something false or at least confused (namely that "Sarah" plays *both* the piano and the guitar). According to Kaplan, Joe would draw an invalid inference from the two utterances in questions, which leads him to infer the false conclusion that "Sarah" plays the piano and the guitar. But this seems odd. It isn't as if Joe somehow doesn't understand the logical structure of the two utterances – it is that the premises are false or confused to begin with!

I think allowing an individual to *fail* in her intentions to repeat her own words, creates a theoretic "mess". Kaplan shouldn't give precedence to inter-subjective word-individuation over intra-subjective word-individuation. By introducing this possibility of doubt towards whether speakers are able to keep track of their own word-identities, he implodes his own proposal. If individuals cannot keep track of their own words in problematic cases, why should they fare any better in the "normal" cases (where they do *not* confuse two words to be one, or one word to be two)? The individual cannot distinguish the "normal" from the problematic cases *from her own perspective*! It seems to me that Kaplan's proposal actually relies on our ability to at the very least keep track of our own words. Wasn't this an unquestioned assumption we made all along? Since the supposed false individuation of words depends on factors that she is *not aware of*, she is incapable of correcting such "mistakes". In chapter one we observed that apart from the intra- and inter-subjective word-individuations we were left with nothing that could explain how the

various tokens are grouped into a word-type. Inter-subjective word-individuations clearly pre-suppose correct intra-subjective word-individuations.

Some readers may wonder whether I simply have forgotten what camp I belong to. As an anti-individualist it shouldn't come as much of a surprise to me that words are partly individuated by *external* factors that the speaker may not be cognizant of. But just consider how far this failure reaches. We take it as a given that we can state truth-conditions for our own words, such as:

x satisfies 'dog' if and only if x is a dog.

How can I be so certain that the word 'dog' used on the right hand side of the biconditional and the word 'dog' mentioned on the left hand side of the biconditional are instances of the same word? The biconditional holds true if and only if the used and the mentioned word are performances of one and the same word (in the thick sense). The speaker is supposed to *know* whether such biconditionals obtain, since she knows whether she is pulling the respective utterances from the same "black box". We are simply *repeating* the word mentioned on the left hand side of the biconditional at the appropriate place on the right hand side of the biconditional – but this time we use it. If there are exceptions where the speaker fails in her attempt to so repeat her own word, or she unknowingly repeats the same word, then this very basic assumption suddenly crumbles. Peter would wholeheartedly assert that:

If x is Paderewski$_{I1}$, then x does not satisfy 'Paderewski$_{I2}$'

and

If x is Paderewski$_{I2}$, then x does not satisfy 'Paderewski$_{I1}$'

If speakers cannot keep track of their own words in this manner, then in those cases they either don't know what they are thinking, or there is a discrepancy between thought and statement (the

one the speaker believes to express the thought accurately). Sure, we accept that individuals may sometimes struggle to express a rather complex belief, but 'Paderewski has musical talent' doesn't strike me as complicated enough to warrant such confusion.

Can we contain the above skepticism about whether speakers are reliably keeping track of their own words by arguing that this happens only when their word-individuations – such as 'Paderewski$_{I1}$' is a different name than 'Paderewski$_{I2}$' – are at odds with the judgments of the linguistic community as a whole (since other competent speakers judge 'Paderewski$_{P1}$' and 'Paderewski$_{P2}$' to constitute performances of the same name)? It would be an external state of affairs (something outside of the speaker's consciousness, but potentially accessible) that would distinguish permissible contradictions from unforgivable ones.

Allowing for intra-subjective word-individuations to depend on factors external to the speaker implies that speakers are unable to keep track of their own words. Such a solution to the puzzle would come at too high a price!

5.3 Modes of presentation to the rescue?

Most responses to the puzzle look towards *modes of presentation* to dissolve it. A straight-forward individualist approach is no solution: not *all* elements of Peter's *modes of presentation* related to 'Paderweski$_{I1}$', or 'Paderewski$_{I2}$' respectively, determine reference – unless we are willing to accept that those names do not refer at all, which is absurd (there is no person called 'Paderewski' who is a famous statesman, but not also a pianist).

Modes of presentation also play a role in the Neo Fregean response, the Hidden Indexical Theory, as well as the Guise Theory – to name the main ones. The disagreement is at what level these modes of presentation come into effect. Do they partly individuate propositions, like the

Neo-Fregeans would have it?[317] Do they instead leave the content of assertions unaffected, such as to individuate the belief ascriptions as a three-place relation between a subject, a proposition and a mode of presentation?[318] According to the guise theory modes of presentations don't affect the semantic content of assertions or belief ascriptions, they merely play a role in *how* the propositions are believed.[319] All these views have had to face ample criticisms, which I won't address here.

Whether we understand modes of presentation to amount to: uniqueness properties, general properties, percept tokens, "stereotypes", "characters", public language expressions, mentalese expressions, functional roles or causal chains[320] they individuate thoughts in a fine-grained manner such that the thoughts expressed by (viii) and (ix) are distinct.[321] Suppose Peter entertains the following while thinking of Paderewski the pianist:

(a) Paderewski has red hair.

He entertains this while thinking of Paderewski the politician:

(b) Paderewski has red hair.

Introducing modes of presentation (at any level) allows us to say that Peter's thoughts expressed by (a) and (b) are not the same thoughts (in content or at least in manner), which is key to the resolution of the puzzle.

---

317 Forbes (1990); Forbes (2006).

318 Prior to having rejected HIT in Schiffer (1987) he first discussed a version of it in Schiffer (1977). See also Crimmins & Perry (1989); Crimmins (1992); Recanati (1993).

319 Salmon (1986); Salmon (1989); Soames (1987); Soames (1988); Soames (2002).

320 Schiffer (1992): p. 511.

321 For the guise theory (17) and (18) would merely be distinct in manner – how this could have implications on the resulting schemata is not clear to me. At best this would allow us to explain why the contradiction in (7) is *hidden* from the agent – but the contradiction doesn't disappear. Why a rational agent, who fully understands his own thoughts cannot "see" that (5) and (6) entail a contradiction remains puzzling. Schiffer develops a similarly spirited criticism against the HIT. Since there are no explicit examples of the three-place form of belief ascriptions, they are at best adverbial qualifiers, in which case the puzzle is not dissolved. The claim that Ralph kissed Emily in the most exciting way, still presupposes that Ralph kissed Emily (Schiffer (1992): p. 518).

Introducing modes of presentation at any level drives a wedge between the cognitive content of the belief the speaker seeks to convey and the semantic content of the utterance that she produces to do so – thereby threatening the successful communication of any belief, unless speaker and audience share the same modes of presentation with regard to a given topic. Taschek explains

> A general difficulty for any such approach arises from the fact that while the demand to do justice to our *intra*subjective logical assessments tends, notoriously, to require that senses be individuated in quite fine-grained ways, the demand to do justice to our *inter*subjective logical assessments tends to preclude such fine-grained distinctions[322].

A solution to the puzzle must also be able to deal with the inverse of the Paderewski puzzle.[323] I will argue that modes of presentation approaches are *not* equipped to do so. I conclude that they therefore do not hit the heart of the problem that the Paderewski puzzle unveils.

In the original version there is a public name 'Paderewski_{P}' that denotes one particular person. Peter has two homophonic and homographic names 'Paderweski_{I1}' and 'Paderewski_{I2}'[324] that are "linked" to the public name 'Paderewski'. 'Paderweski_{I1}' denotes a pianist, while 'Paderewski_{2}' denotes a politician. Unbeknownst to Peter the pianist and the politician are one and the same individual: Ignacy Ian Paderewski.

Consider the inverse puzzle: Peter believes that there is *one* public name 'Paderewski_{I}' where in fact there are two distinct homonymous names 'Paderewski_{P3}' and 'Paderewski_{P4}' – such that non-vacuous occurrences of one name cannot be interchanged with occurrences of the other without affecting the truth-value of the sentence as a whole. Suppose furthermore we live in the counterfactual world in which there is an Ignacy Paderewski (named 'Paderewski_{P3}') and a Jan Paderewski (named 'Paderewski_{P4}'). Ignacy is a famous pianist (and not a politician); Jan is a

[322] Taschek (1998): footnote 11, p. 331.
[323] See also Fiengo and May (2006): chapter 4.
[324] I am using subscripts to distinguish the two homophonic and homographic names.

famous Polish politician (and not a pianist). The two men are identical twins and look exactly alike the actual Ignacy Jan Paderewski. In the counterfactual world, Peter *falsely* comes to believe that the pianist named 'Paderewski$_{P3}$' is the same person as the politician known as 'Paderewski$_{P4}$'. The resemblance of the two men on photographs has reinforced that belief. Usually people refer to the two men in question only by their respective last name. Even if that wasn't the case, Peter so firmly believes that Paderewski the politician and Paderewski the pianist are one and the same, he would conclude that the person referred to by 'Paderewski' is called 'Ignacy Paderewski' or 'Jan Paderewski' – depending on whether the speaker is using Paderewski's first or middle name. Peter can pick out the person known to others as 'Paderewski$_{P3}$' from a crowd. He can also recognize the person known to others as 'Paderewski$_{P4}$' and distinguish him from (most) other men (not from his identical twin). Peter will call both men 'Paderewski$_{I}$' and believe to have selected the same individual at both occasions.

Peter is being told on separate occasions:

(c) Paderewski$_{P3}$ has red hair.

(d) Paderewski$_{P4}$ has red hair.

(e) Paderewski$_{P3}$ has musical talent.

(f) Paderewski$_{P4}$ has a talent for public speaking.

Other speakers will usually discern from context whether Ignacy or Jan Paderewski was meant by either one of those utterances. Thus, speakers who are aware of both individuals (and their shared generic name) will know to distinguish (c) from (d), or will at least be aware of the need for disambiguation. Suppose Peter trusts the person who utters (c) – which belief will he form in response? Peter would express it thus:

(g) Paderewski$_I$ has red hair.

That's the key question: is (c) equivalent to (g)? There is no reason for not also viewing it equivalent to (d) since Peter doesn't distinguish between 'Paderewski$_{P3}$' and 'Paderewski$_{P4}$'. That is the problem: (c) and (d) are not equivalent in what they express. Most of us probably feel uneasy to call (g) true (let alone justified). Something must have gone wrong during the process of testimony although Peter had all intentions to use the "word-form" 'P-a-d-e-r-e-w-s-k-i' like other speakers in his linguistic community.

If presented with (e) and (f) Peter will conclude:

(h) Paderewski$_I$ has musical talent and a talent for public speaking.[325]

From which it follows that

(i) There is some person that has musical talent and a talent for public speaking.

Although quite possibly true, (i) clearly doesn't follow from (e) and (f), since 'Paderewski$_{P3}$' and 'Paderewski$_{P4}$' do not denote the same individual. It would be an illogical inference to draw. Just like in the original Paderewski puzzle we are unwilling to judge Peter irrational. While Peter in the original example realizes *too little* – he ought to choose between his two observations (i) and (ii), Peter in the reverse case goes beyond the evidence expressed by (e) and (f) – he draws a conjunction he ought not to and infers *too much*.

A possible objection to my demand that the inverted puzzle should be given equal consideration as the original puzzle may go as follows: Although Peter is assumed rational in both, he doesn't even appear to be using the public names 'Paderewski$_{P3}$' or 'Paderewski$_{P4}$' in the inverse case. In the original thought example we *successfully communicate* with Peter using

[325] Depending on what characteristics are true of the two men in question, Peter may be puzzled about testimonies that seem *contradictory* to him. For instance, he may (truthfully) be told that Paderewski$_3$ likes to eat green peas while at a different occasion he is being told (truthfully) that Paderewski$_4$ does not like to eat green peas.

the public name 'Paderewski$_P$', which is what leads us to believe that Peter is competent in the use of that name. The same cannot be said for the reverse case. The reference of Peter's name 'Paderewski$_I$' is neither equivalent to that of 'Paderewski$_{P3}$' nor 'Paderewski$_{P4}$'. We should not assume communicative success when his community talks to him about Ignacy or Ian with the names 'Paderewski$_{P3}$' and 'Paderewski$_{P4}$' respectively – since Peter's name 'Paderewski$_I$' doesn't distinguish between those individuals. I disagree with this analysis. Peter acquires information imparted through the name 'Paderewski$_{P3}$' from (c). While a competent speaker will be able to discern whether (c) concerns Ian Paderewski or Ignacy Paderewski – or will at least know that it could apply to either one of them – Peter will conclude that it is true of both of them (while lacking the awareness that there are two distinct individuals possibly being spoken of).

Somebody whose cognitive system is not at all "linked" to the public names 'Paderewski$_{P3}$' and 'Paderewski$_{P4}$' will merely conclude that somebody called 'Paderewski' has red hair – information of much inferior epistemic value. Suppose there are thousands of individuals whose names are spelled 'P-a-d-e-r-e-w-s-k-i', then Peter, in contrast, at least narrows the subject of the statement down to two individuals (although he believes he has narrowed it down to one). We wouldn't want to say that the person who hears "Paderewski has red hair" but doesn't know whether the speaker was talking about Ignacy, didn't learn anything beyond 'somebody called 'Paderewski' has red hair'. Information about Ian or Ignacy was conveyed via the name 'Paderewski$_{P3}$'. The speaker who doesn't know how to disambiguate the utterance still epistemically fares better than Peter, since at least he is aware of the need for disambiguation. I conclude that communication is successful – albeit not fully successful – in the original, as well as the reverse case and must be explained. Peter's familiarity with the "name-*form*" alone will not account for the communicative success. Recall that it was part of the description of this case

that if Peter were asked "who is Paderewski, the famous politician?" – while pointing at a crowd, Peter would correctly point out the right individual.

Introducing modes of presentation to explain why (c) and (d) *as entertained by Peter* express one and the same proposition requires that modes of presentation don't only allow the relevant individuations to be more fine-grained but also more *coarse*. This could go roughly like this: since Peter doesn't associate[326] distinct modes of presentations with (c) and (d) respectively, they express the *same* proposition out of his mouth. It must be pointed out that we would generally refrain from allowing modes of presentations to play such a role. Consider:

(j) Elms are beautiful trees.

(k) Beeches are beautiful trees.

Putnam is on record for confessing that he associates the same "modes of presentation" with the words 'elm' and 'beech'.[327] Are we willing to accept that Putnam's utterances of (j) and (k) express the same assertion? I think not. Let's not overlook how Putnam's observation is backed by a strong belief that whatever kind of tree 'elm' and 'beech' pick out, they denote *different* trees. He believes that (j) and (k) do *not* make the same claim, although he is unable to specify the difference between elms and beeches. Putnam accounts for the distinction between (j) and (k) with the *division of linguistic labor*, which allows individual speakers to defer to the relevant experts to enrich their individual conceptions.[328] I do not doubt that we intricately rely on more informed speakers (perhaps even in the way just described), but let's also not overlook that our

[326] How modes of presentation are associated with or related to propositions will turn out to be different depending on the particular theory endorsed.
[327] Putnam (1975, 1975): pp. 226-227.
[328] The expert conceptions (proposals for how to pick out the objects in question) are not entirely fixing the extension of the relevant terms, since even experts can be wrong about the nature of the subject under discussion.

deferring to other speakers may be hindered in the given case, if Putnam were to think that 'elm' and 'beech' behave like 'hedgehog' and 'urchin'.[329]

Solutions to the puzzle that rely on modes of presentation don't have the resources to distinguish Peter's situation in the inverse Paderewski puzzle from Putnam's elm-beech thought example. I conclude that those approaches therefore are not providing us with the appropriate solution to the puzzle.

5.4 Assignments to the rescue?

According to Fiengo and May, the name-type 'Paderewski' is a lexical item, but doesn't refer, since it can name any number of things.[330] Their notion of *names* is better identified with what Kaplan calls 'generic names' or what I refer to as 'name forms'. They argue that only expression-types in which names occur refer. Expressions are "employed lexical items […that occur] as constituents of sentence types". [331] Numerical subscripts distinguish among distinct expression types (such as 'Paderewski$_1$' and 'Paderewski$_2$').

At the heart of Fiengo and May's account lies the idea

> that having distinct conceptions does not imply that distinct things are being conceptualized. Someone may think of Paderewski in this way or that, but from this nothing follows whether he is thinking of one person or two.[332]

This is explained by reference to assignments. *De dicto* use of assertions contain assignments in addition to propositions. "Assignments are relations between syntactic expressions and values

[329] 'Hedgehog' and 'urchin' are co-referential.
[330] Fiengo and May (2006): pp.146-147.
[331] Ibid. p.146.
[332] Fiengo and May (2006): p. 143.

[…]".[333] For example, "John believes [[Aristotle$_1$ was the teacher of Alexander the great] and ["Aristotle$_1$" has the value Aristotle$_1$]].[334]

Fiengo and May distinguish between Aristotle, the philosopher, and Aristotle, the shipping magnate through different indices.

> [T]he natural syntactic interpretation of indices is that coindexed expression-tokens (in a discourse) are occurrences of the same expression-type, but that noncoindexed expression-tokens are not. So occurrences of "[$_1$Aristotle]" and "[$_2$Aristotle]" are occurrences of distinct expression-types, for each of which there is a distinct Assignment;[335]

In this case the individual employs *different* expression-types containing the same name and these expressions denote *different* objects. Perception and reality coincide: there are two people named 'Aristotle' and the individual thinks of them as distinct individuals.

Fiengo and May adhere to a principle called 'Singularity', which states that a speaker *S* can have two expression types e_1 and e_2 that contain the same name, if and only if *S* believes that e_1 and e_2 do *not* corefer.[336, 337]

Peter's naming practices adhere to the Singularity Principle: He employs two distinct 'Paderewski' expression types: 'Paderewski$_1$' and 'Paderewski$_2$' that happen to corefer, even though Peter believes them *not* to. When Peter informs us: "Paderewski has musical talent" and

---

[333] Ibid. p.14.

[334] Ostertag raises concerns as to how to negate *de dicto* beliefs of the conjunctive form Fiengo and May put forth. Suppose person *S* believes [P & A], where A is an assignment. A denial of such a report should simply consist in a denial of belief in P & A. "But whereas one can use a negative *de dicto* belief report to deny belief in P, one cannot use a negative belief report to deny belief in A (to do this is incompatible with the report's being *de dicto*". Ostertag concludes that the conjunctive analysis is incorrect; assignments are not part of what is being asserted (Ostertag (2007)).

[335] Fiengo and May (2006): p. 17.

[336] Ibid. pp. 70-75, p. 147.

[337] According to Fiengo and May, what they call 'names' do not refer. Their terminology 'name' is better understood to stand for what Kaplan calls 'generic names', which allows us to say such things as: "many people share the name 'David'". I am interested in when speakers share the same name and how this facilitates their communication about the same topic. It thus seems that neither, what Fiengo and May call 'expressions' nor what they call 'names' is well suited to find answers to my questions. The former are too fine-grained and the latter don't refer. If the speaker believes that the name 'Paderewski' has different *assignments*, then she believes that 'Paderewski$_1$' and 'Paderewski$_2$' behave like 'Cicero' and 'Tully', and she (falsely) believes that 'Paderewski$_1$' and 'Paderewski$_2$' do not co-refer.

tells us at a different occasion "Paderewski has no musical talent", then the *de dicto* content of those beliefs expressed is the following:

(1) [Paderewski$_1$ has musical talent] and ["Paderewski$_1$" has the value Paderewski$_1$].

(2) [Paderewski$_2$ has no musical talent] and ["Paderewski$_2$" has the value Paderewski$_2$].

We ascribe to Peter the following beliefs:

(3) Peter believes [[Paderewski$_1$ has musical talent] and ["Paderewski$_1$" has the value Paderewski$_1$]].

(4) Peter doesn't believe [[Paderewski$_2$ has musical talent] and ["Paderewski$_2$" has the value Paderewski$_2$]].

Due to distinct assignments, Peter thus doesn't contradict himself with (1) and (2), nor do *we* ascribe to Peter contradictory beliefs.

Ostertag Points out that (3) and (4) contradict Fiengo and May's principle of *Singularity*. As the enlightened speakers we cannot employ different assignments for the expressions that denote Paderewski the pianist and Paderewski the politician – precisely because we believe them to be one and the same.[338] Fiengo and May are fully aware of this problem. They write:

> Consider, in the context of Singularity, a sincere utterance of "Paderewski$_1$ is Paderewski$_2$." The immediate conclusion we come to is that *no speaker is in a position to assert this consistently with his beliefs*.[339]
>
> Rather the attribution [of assignments] is to the second party—that is, normally to the addressee. It is as if the speaker, in uttering "Paderewski is Paderewski" *de dicto* says in effect to the hearer, *your* "Paderewski"-expressions corefer, and this can be asserted consistently with the speaker's beliefs about how many people are named "Paderewski".[340]

[338] Ostertag (2007).
[339] Fiengo & May (2006): p. 148.
[340] Ibid.: pp.152-153.

It is *ad hoc* to suppose that enlightened speakers simply *mimic* Peter's state of mind in their belief ascriptions. Ostertag asserts that it is equally *ad hoc* to suppose that in an enlightened context we *mimic* Peter's assignments to reflect *his* usage of the name 'Paderewski'.[341]

What Fiengo and May's account explains well is Peter's motivation to merge his assignments into one. Once Peter believes that there is only one Paderewski and that he mistakenly thought there were distinct men bearing that name, he must, in order to conform to *Singularity*,

> give up [the belief] that there are two distinct "Paderewski"-expression types, and replace it with the belief that there is just one.[342, 343]

What remains unclear to me is why Peter must adopt the informant's expression 'Paderewski', or that of his linguistic community (from the perspective of Fiengo and May's account). Peter clearly shares the name (form) 'Paderewski' with his linguistic community – but that name (form) can denote a variety of individuals. What remains unexplained is that he *shares* with his community and interlocutor the same expression 'Paderewski' – the one denoting Paderewski the Polish statesman and pianist. Coreference cannot be sufficient for sharing the same expression, or else 'Paderewski_1' and 'Paderewski_2' would be tokens of one and the same expression. So, given that Peter doesn't use the same expression as his informant, how is he able to understand her utterance: "Paderewski is Paderewski"? Supposedly, the informant mimics Peter's expressions 'Paderewski_1' and 'Paderewski_2'. Through the Gricean maxim of relevance,[344] Peter knows that what the speaker means to say is that Paderewski_1 is Paderewski_2. Even though Peter

[341] Ostertag (2007).
[342] Fiengo & May (2006): p.154.
[343] In my eyes, Fiengo and May don't sufficiently recognize that such a change in a person's conceptual apparatus is going to be a gradual process.
[344] Grice (1989): p. 27.

and the informant's expressions in question share the same name (form) and are coreferential it is unclear why the informant is in a position to *mimic* Peter's expressions.[345]

Furthermore, why is Peter committed to adopt the informant's expression 'Paderewski' upon hearing the identity statements?[346] Fiengo and May want to argue that when Peter utters:

(5) Paderewski has red hair.

and we compare this with the following utterance of an enlightened speaker

(6) Paderewski has red hair.

then the *non-de dicto* content of (5) and (6) are identical. It is merely the *de dicto* contents of (5) and (6) that diverge.

(5') [Paderewski$_1$ has red hair] and ["Paderewski$_1$" has the value Paderewski$_1$].

(6') [Paderewski has red hair] and ["Paderewski" has the value Paderewski].

So far, so good! But what do Peter and the enlightened speaker share in virtue of which (5) and (6) have the same content? Again, it cannot be the name (form) – since name (forms) don't refer.

Suppose someone utters "Aristotle is a man" with Aristotle the philosopher in mind:

(7) Aristotle is a man.

The same person then utters in a different context "Aristotle is a man" with Aristotle the shipping magnate in mind.

(8) Aristotle is a man.

Clearly, (7) and (8) do not state the same *de dicto* or *non-de dicto.*

(7') [Aristotle$_1$ is Athenian] and ["Aristotle$_1$" has the value Aristotle$_1$].

(8') [Aristotle$_2$ is Athenian] and ["Aristotle$_2$" has the value Aristotle$_2$].

[345] In my own account I will attempt to bridge that gap. It is admittedly a tricky problem!
[346] I will also attempt to explain this commitment in my own account.

But how can we show this, if the assignments only enter into the *de dicto* content?

It seems to me that the *non-de dicto* content of the utterances must be the following:

(7") Aristotle$_1$ is a man.

(8") Aristotle$_2$ is a man.

What do the subscripts designate? "They do not, on our view, use one tool "[Aristotle]" on some occasions to refer to the philosopher and on others to the shipping magnate."[347] But which are the different tools that they are using? Are they using distinct expression types or name types? It seems that neither fits the bill. Certainly not different name (forms); but do they symbolize different expression types? If so, then Peter's belief *non-de dicto* is also different from the enlightened speaker, since Peter is using different expression types than his informant. Let's designate the enlightened speaker's expression with subscript 'E':

(5") Paderewski$_1$ has red hair.

(6") Paderewski$_E$ has red hair.

Some third entity is needed that Peter and the enlightened speaker can share – in (5) and (6) – in virtue of which their expressions contain the same name (form) 'Paderewski' and in virtue of which their expressions are coreferential. This entity is not contained in 'Aristotle$_1$' and 'Aristotle$_2$'. What Fiengo and May seem to be saying is that understood *non de dicto*, Peter and the enlightened speaker's *expressions* 'Paderewski$_1$' and 'Paderewski$_E$' constitute performances of the same expression types because they corefer *and* contain the same name (form) – whereas this cannot be said of 'Aristotle$_1$' and 'Aristotle$_2$' (they fail to corefer). In contrast, in the *de dicto* "mode" the two expressions 'Paderewski$_1$' and 'Paderewski$_E$' do *not* constitute

[347] Fiengo & May (2006): p. 17.

performances of the same expression type. I struggle to make sense of expressions being identical under one mode, but not under another.

In the previous section I criticized solutions to the Paderewski puzzle that employ modes of presentation because they cannot handle the inverse puzzle. Fiengo and May recognize the need to deal with the flip-side of the Paderewski puzzle and offer the following solution:

> The hearer believes there is only one person named "Paderewski," so the only relevant identity statement he would hold to be true would be "Paderewski$_1$ is Paderewski$_1$," a sentence of the form $\lceil \alpha = \alpha \rceil$. But then the initial reaction of the hearer to the speaker's utterance of "Paderewski isn't Paderewski" is that it is the negation of this, that the speaker had uttered a logical falsehood. […The hearer] could recognize that it is possible that the speaker's intention was to utter a sentence of a different form, namely, "Paderewski$_1$ isn't Paderewksi$_2$".[348]

Again, the informativeness of the identity statements or negated identity statement is explicated in terms of a mimicking of one's interlocutor's assignments in conjunction with the Gricean conversational maxim of relevance.[349] Again, I am not sure whether their solution to the inverse scenario can explain why Peter is committed to adopting the informant's epistemically superior assignments, given that now their expressions don't corefer.

Fiengo and May's proposal is a novel approach to the Paderewski puzzle. In my eyes what they get right is that we have to take seriously Peter's *de lingua* beliefs, that is, how he keeps track of his own words (i.e. 'Paderewski') over time and across contexts. In the following section I will propose a solution where Peter's expressions flat out constitute different performances of the name form 'Paderewsi' than those performances of the enlightened speakers.

[348] Fiengo & May (2006): p.155.
[349] Grice (1989): p. 27.

5.5 Registers to the rescue!

I agree with Fiengo and May that the distinct ways that Peter employs to think about Paderewski are part of what he asserts in (i) and (ii). In my account there are no assignments or modes of presentation that are added to the content of what Peter asserts. The difference in how Peter conceptualizes Paderewski, the pianist, and Paderewski, the politician, is reflected in the comparative content of Peter's names 'Paderewki$_{I1}$' and 'Paderewki$_{I2}$' respectively. I ascribe the speaker the kind of linguistic control that renders Peter's beliefs fine-grained in such a manner that allows us to distinguish beliefs attributed to Paderewski *under the guise* "politician" from those attributed to him *under the guise* "pianist". Modes of presentation accounts fail to *contain* the multiplicity of beliefs that result from different conceptions associated with an individual or group of things – thereby threatening successful inter-subjective communication in general. My own account sets *limits* to the fine-grainedness of beliefs across speakers. Words are *not* individuated in terms of the content of the corresponding registers that contain the relevant individual conceptions. Unless the speaker falsely believes one public word to be two, or two public words to be one, the content of the register is determined by what the public word expresses.[350, 351] Finally, my own account is designed to deal with the original, as well as the inverse puzzle.

In context A, Peter is exposed to 'Paderewski$_{P1}$', which he intends to repeat and produces 'Paderewski$_{I1}$'. In context B, Peter is exposed to 'Paderewski$_{P2}$', which he also intends to repeat and produces 'Paderewski$_{I2}$'. The two performances 'Paderewski$_{P1}$' and 'Paderewski$_{P2}$' are instances of one and the same name (in the thick sense), that is, competent speakers will

[350] Assuming the speaker doesn't opt out of the respective public language practice.
[351] The content of a register not linked to a public repetition tree is simply what the idiolectal word expresses (of a matured register).

recognize them as performances of the same name. Peter, on the other hand, will judge his own performances 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' to not constitute repetitions of one another.

I will argue that the individual correctly judges her own words 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' to represent *different* words. Peter split one public name into two separate mental registers. This constitutes a failure to repeat the public word 'Paderewski$_{P1}$' when he utters: 'Paderewski$_{I1}$'. Nevertheless, *both* registers are *linked* to one and the same public word 'Paderewski$_{P}$' (of which 'Paderewski$_{P1}$' and 'Paderewski$_{P2}$' are instances). This link is not merely assumed in principle, but can be tested: if Peter finds out that there is only one Paderewski and that 'Paderewski$_{P1}$' and 'Paderewski$_{P2}$' denote that Paderewski – will he as a result of that information *merge* his registers 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' into one? If so, then Peter's words 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' were all along both linked to 'Paderewski$_{P}$' – without his having been *adequately* competent in 'Paderewski$_{P}$' (since competence is incompatible with his labelling distinct registers 'Paderewski$_{P}$').

Let's indulge in drawing a metaphoric picture to make this distinction more palatable. Imagine that we have mental registers for proper names and kind terms.[352] So, everything a person knows about arthritis goes into the arthritis-register. If a person were to at one point believe that arthritis can afflict body-parts other than the joints and then learn that this is not so, she will do a little "house-cleaning" and dispose of or adjust the related beliefs. We can also add beliefs to a register thereby leading to a more complete conception regarding a given topic. The individuation of registers is not determined by the particular attitudes or dispositions that the speaker associates with it (at any given time). What makes a register persist throughout the changes in conceptions such as to allow for it to pertain to the same topic? The topic of each

[352] The register analogy is a popular one that has been used by several proponents in varying ways. To name a few: Bach (1987); Evans (1973); Evans (1982); Forbes (1989); Schroeter (2008).

register is anti-individualistically individuated – given that the speaker intends to participate in the relevant public practice. It is not up to the arthritis patient to define what the 'arthritis' label of her respective register is about – external conditions, such as social and physical environmental factors will individuate that register (in part). The existence of mental registers is compatible with anti-individualism. At the same time the content of each register presents itself to the agents as pertaining to the same topic, while the content of distinct registers each present themselves as pertaining to different topics.

Peter's registers are *all-along* connected to the public name 'Paderewski'. Peter intends to successfully communicate with his environment and is thereby committed to merge his two separate Paderewski registers – when faced with the appropriate evidence (stemming from a source Peter trusts). He has (unknowingly) adopted a split *public* concept. When we say: "Peter believes that Paderewski has red hair", we speak truly (in a sense), since Peter believes of Paderewski that he has red hair. The distortion of our translation is rendered plain when we have to say things such as: "Peter believes that Paderewski is not Paderewski." We can only do full justice to Peter's belief thus: Peter has a *split* Paderewski-concept that divides into Paderewski *under the guise of pianist* and into Paderewski *under the guise of politician.*[353] He believes that Paderewski under the guise of pianist [split concept 1] is not Paderewski under the guise of politician [split concept 2]".

Can a word *w* be linked to a public repetition tree without *w* constituting a repetition of the corresponding public word? Remember how I argued in chapter four that somebody who does not yet know anything about lions may utter "lion" in a question such as "what is a lion?" without thereby yet being competent in the use of the public word 'lion'. In that case the quasi-

[353] I am borrowing terminology from Soames and Salmon here, but my proposal is significantly different from theirs.

word ‘lion’ is linked to the public word ‘lion’ without it constituting a performance of that public word (yet) – provided that she uses future repetitions of that word to denote lions (and only lions).[354]

The realization that ‘Paderewski$_{I1}$’ and ‘Paderewski$_{I2}$’ ought to be merged represents a *substantial change* in the structure of the cognitive apparatus and will affect his inferential behavior.

Am I not facing the same problem as Fiengo and May? How can Peter understand the identity statement “Paderewski is Paderewski” uttered by an enlightened speaker – if Peter and the enlightened speaker do not share the same word ‘Paderewski’? Fiengo and May argued that the informant mimics Peter’s assignments. Ostertag accused their solution of being *ad hoc*. Does my account fare any better? I claim both of Peter’s words ‘Paderewki$_{I1}$’ and ‘Paderewski$_{I2}$’ to be *linked* to the public repetition tree ‘Paderewski$_{P}$’. Peter is thus already committed to *merge* his split registers into one pubic ‘Paderewski$_{P}$’ register – if presented with the appropriate evidence (i.e. he could follow Paderewski, the politician around and come to realize that he is in fact Paderewski the pianist[355]). Can Peter *understand* the testimony: “Paderewski is Paderewski”? The informant tells Peter that Paderewski$_{P}$ is Paderewski$_{P}$. In order for this identity statement to be informative it must have an effect on how Peter organizes his registers. That is, Peter must “take it in” as ‘Paderewski$_{I1}$ is Paderewski$_{I2}$’. But for Peter to *understand* the testimony, it would seem that he must share the word ‘Paderewski$_{P}$’ which would render the identity statement trivial.

[354] Recall Mr. Anchovy whom I introduce in chapter four. He goes on to use ‘lion$_{A}$’ to denote anteaters; his word is *not* linked to the public repetition tree ‘lion$_{P}$’.
[355] Unfortunately this is no longer possible, since Paderewski is long dead.

Are there other examples of words that share the same denotation, but for one reason or other, need to be treated (in schemata) as distinct words? Do we communicate successfully when these distinct but co-referential terms are used in conversations?

Words and their pejorative counterparts come to mind such as: 'woman' and 'broad', 'homosexual' and 'fag', 'obese' and 'fat', 'German' and 'Kraut'… Clearly, there is a big difference between somebody calling me 'Kraut' versus 'German'. It gets a little more complicated if we imagine a person who severely dislikes Germans uttering: 'Nathalie is German'. This person may have shied away form the explicit insult and opted for the implicit one. On the other hand, a close friend may say jokingly: 'Nathalie is a Kraut' – using the relevant term in an ironic and figurative manner and drawing attention to the fact that I like sausages and *Sauerkraut* without intending to insult me. Such use would be comparable to feminists referring to themselves as 'bitches'. Perhaps several years down the line, this new use of 'bitch' becomes well engrained in ordinary practice and will enter the English lexicon as a new literal use for the word-form 'b-i-t-c-h' – as happened with the word-forms 'l-i-v-i-d' and 'g-a-y'. For the time being, however, 'bitch' is *literally* insulting.

I don't want to argue that it is some kind of *conceptual overlap* between the 'Paderewski_{I1}' (or 'Paderewski_{I2}') register and the informant's 'Paderewki_{P}' register that allows Peter to understand the identity she attempts to convey, since this would make it impossible for me to solve the inverse puzzle using the same explanation. So, I am not sure whether pejorative terms are of any help here.

Peter could understand the identity statement to express that

i) Paderewski_{I1} is Paderewski_{I1}.

ii) Paderewski_{I2} is Paderewski_{I2}.

iii) Paderewski$_{I1}$ is Paderewski$_{I2}$.

Only iii) renders the identity statement informative to Peter. According to the Gricean maxim of relevance this is how Peter ought to interpret the utterance. According to my own proposal, Peter is not able to *fully* understand the informant's utterance, since (at the time) they don't share the same words. 'Paderewski$_{I1}$' is a distinct word (in the thick sense) from 'Paderewski$_{P}$'. The obvious alternatives that the "message" can be gotten across on the basis of an overlap of conceptions or merely by a shared word-*form* – I already rejected. Nevertheless, once Peter realizes that 'Paderewski$_{P}$' denotes the same person as both of his words 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' and he trusts the informant, he will merge his two registers into *one*. For the purpose of logical consistency, it is advisable to do so, independently of public usage. Once Peter merges those two registers, the *new* resulting register will then be individuated in terms of the public repetition tree 'Paderewski$_{P}$'. Part of the motivation for this merger is that this is Peter's way of *fixing* his failed repetition attempts of 'Paderewski$_{P}$'. Of course, he won't think in terms of registers clashing with public usage – but one ought not be using two distinct idiolectal words, to correspond to one public word.

Peter must retire his registers 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' – from his new semantic point of view they are confused and thus of no more use. As a result, Peter cannot reflect on his previous belief that "Paderewski$_{I1}$ is not Paderewski$_{I2}$" and state it to be false, since he can no longer use those words. According to my account, the change in belief is one in *structure*, not (merely) content.

Even though a transfer of knowledge is not possible until Peter merges the registers in question and thereby "*unlocks*" previously held beliefs about Paderewski as knowledge. Prior to this merging procedure, when Peter hears from another speaker: "Paderewski has red hair" I am

forced to argue that Peter doesn't acquire knowledge from that testimony, since he cannot understand the word 'Paderewski$_P$' contained in the utterance. I want to argue that the *semantic link* between 'Paderewski$_P$' and 'Paderewski$_{I1}$' allows for Peter to acquire some information – *conditional* knowledge, that is. If Peter follows through on his commitment to use his words 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' effectively when talking to other speakers (of his linguistic community), then if he is presented with evidence that his registers *conflict* with how the word he intends to repeat 'Paderewski$_P$' is used by the competent speakers, he will make the necessary changes to live up to the *responsibilities as a language consumer*.

In the inverted case, Peter has one register where he ought to have two. We should thus ascribe to Peter an *amalgamated* public concept corresponding to the two names 'Paderwski$_{P3}$' and 'Paderewski$_{P4}$'. His thoughts are somewhat challenging to make sense of. When Peter splits his one Paderewski$_I$-register into register: Paderewski$_{P3}$ and register: Paderrewski$_{P4}$, he won't know which beliefs apply to Ian and which ones apply to Ignacy – or to both of them.

Peter's belief of "Paderewski has red hair" will have to be translated as: "Peter believes of Paderewski, the musician, and Paderewski, the politician, that they *as one unit* have red hair." If we were to opt for the simpler claim: "Peter believes that Paderewski, the musician, and Paderewski, the politician, both have red hair" – we thereby ascribe to him the belief that there are *two* such individuals (a belief he lacks).[356] Regardless of the complexity of our translation manual, there is a clear link from Peter's 'Paderewski$_I$' to the *public* concepts, which in principle allow us to express Peter's beliefs.

[356] For the same reason we cannot simply claim that his name 'Paderewski$_I$' is ambiguous. If a speaker is competent in the use of an ambiguous word-form, he knows that the various instances can apply to different reference groups or individuals. Peter is altogether unaware of the need for disambiguation.

Admittedly, the original and the reverse case are different: nothing speaks against having two names for one and the same person – as long as one remembers that whatever one predicates of the one name, one must predicate of the other (which Peter fails to do). However, one ought not have a single proper name for two distinct individuals to begin with.

I have shown how a register analogy can be compatible with anti-individualism. In both puzzles, the original and the inverse, it is not up to Peter what changes to implement in his conceptual apparatus. Once faced with reliable information challenging his confused register system pertaining to Paderewski or "the Paderewskis" he must

- in the original example: merge his 'Paderewski_{I1}'-register with his 'Paderewski_{I2}'-register.
- in the inverse case: split his one 'Paderewski'-register into two separate registers.

If not, he opts out of successfully communicating with others via the word that the utterance "Paderewski" stands for. It is the link between Peter's and the public registers that allows us to convey to Peter in the actual world: "Paderewski, the pianist, is Paderewski, the politician"[357] and to Peter in the counterfactual world: "Paderewski_{P3}, the pianist, is not Paderewski_{P4}, the politician."

The register analogy is a useful tool in explaining why some identity statements seem trivial to us, such as "Hesperus is Hesperus", while others have informational content such as "Hesperus is Phosphorus". All the content within a register represents itself to us as "blatantly", that is, without reflection, about one and the same topic. In the long run we should structure our representative system such that co-referential terms become alternative labels for one and the

[357] Due to the Gricean maxim of relevance telling Peter: "Paderewski is Paderewski" is probably sufficient to convey the identity we intend to. Peter will assume that we didn't mean to inform him of something trivial such as: 'Paderewski_1 is Paderewski_1' or 'Paderewski_2 is Paderewski_2'.

same register. We then wouldn't even have to reflect on transferring beliefs from the Hesperus to the Phosphorus register, as well as the Venus register. It all depends on how well-engrained the identity between Venus, Phosphorus and Hesperus is within an individual's cognitive system to lead to unified record-keeping. Depending on how long we were engaged in Venus discussions and Hesperus conversations, prior to discovering the related identity, we may keep two registers: one for Venus and one for Hesperus. As a result, considerations such as 'if Venus is a planet, then Hesperus must be a planet also' may take some reflective effort. Whereas 'if Venus is a planet, then Venus is a planet' is implicit in our record keeping. Given that representative systems may vary across individuals, *identities that are trivial to some may be informative to others*.

If we were to communicate to Peter: "Paderewski, the pianist is the same individual as Paderewski, the politician" – it is correct that this utterance is just as trivial and uninformative to *us* as "Hesperus is Hesperus". Assuming that we are enlightened and know that Paderewski, the pianist is the same person as Paderewski the politician, we only have one Paderewski register and only one name associated with that register. To Peter, however, this statement is *informative*. It tells him that his two Paderewski-registers have the same *topic* – they "pick out" the same person. If he wants to effectively communicate with us he should eventually "melt" the two distinct Paderewski registers into one. As a result, all the modes of presentation associated with 'Paderewski$_{I1}$' ought to transfer to 'Paderewski$_{I2}$' and *vice versa*. In the end having two distinct but homophonic and homographic names 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$' for one and the same individual will become redundant and only one public name will remain – thereby rendering the statement 'Paderewski is Paderewski' also uninformative to Peter. Once Peter realizes that Paderewski$_{I1}$ is the same individual as Paderewski$_{I2}$ there is no more reason to keep

the two registers separate – without the registers there is no more reason for the distinct names, which in turn necessitates one and the same sign for 'Paderewski' in the formalization of his inferences. I thus view Fiengo and May's principle of *Singularity* as appropriate. I furthermore agree with Fiengo and May's bald bi-conditional:

> We will find that the *beliefs* of individuals will be critical in the statement of those identity conditions on words: two occurrences of words are occurrences of the same word for an individual if and only if that individual believes they are the same.[358]

In other words, individuals cannot be wrong about the individuation of their own words – even though they may be mistaken about the individuation of pubic words.

Imagine you have been reading some impressive papers by a philosopher named 'Sarah Connor'. Over the past year, her publications ceased and another author captures your attention: Sarah Jameson. One day you mention to a colleague that the most promising work you have been reading in your field lately was published by two women: Sarah Connor and Sarah Jameson. Your friend then informs you that Sarah Connor has recently gotten married and changed her name to 'Sarah Jameson'. Knowing that those two names belong to one and the same person should have you ascribe whatever you believed of Sarah Connor to Sarah Jameson and *vice versa* (resolving all conflicting beliefs) and thereby rendering the identity statement uninformative to you over time. It takes a while for these identities to become automatic assumptions within your cognitive apparatus, because, until they are engrained in our cognitive apparatus, such changes take some reflective effort.

Identity statements such as 'Hesperus is Phosphorus' – involving distinct public names – may in some circumstances be uninformative to a person. I have in mind cases where one name is introduced in terms of the other. Suppose Peter knows that Hesperus is the bright celestial object visible in the evening sky. He is introduced to the name 'Phosphorus' thus: "Hesperus is

[358] Fiengo & May (2006): p. 53.

Phosphorus." It seems to me that Peter would simply add the newly acquired name 'Phosphorus' as an alternative label to his Hesperus register, since he doesn't really know anything else about Phosphorus. A third person then (not having heard the first identity statement) says: "Hesperus is Phosphorus." This last statement should have no informational content for Peter (unless he didn't trust the first source on the matter). He lacks the resources to pick out the referent of 'Philosophorus' independently of his conceptions regarding the referent of 'Hesperus'.

In most cases, if names are distinct in form such as 'Hesperus' and 'Phosphorus' we will acquire information about what they denote independently (regardless of their co-reference). This probably explains why we have distinct registers for 'Hesperus' and 'Phosphorus' even when we know that they co-refer. For some it may require reflection to see that they are identical – which then explains why usually 'Hesperus is Phosphorus' does not have the same cognitive significance as 'Hesperus is Hesperus'.[359]

The distinction between informative and uninformative identity statements can be made without mention of modes of presentation. I do, however, admit of a correlation between registers and modes of presentation. If a person associates *distinct* modes of presentation with 'Hesperus' and 'Phosphorus' this indicates distinct registers for those words – thereby rendering 'Hesperus is Phosphorus' informative. As discussed earlier, distinct registers or names don't necessitate distinct modes of presentation (recall Putnam's reflections on 'elm' and 'beech').

[359] The register analogy allows us to account for why for *some* individuals identity statements such as 'Hesperus is Phosphorus' have the same cognitive significance as 'Hesperus is Hesperus'. This happens when 'Hesperus' and 'Phosphorus' are simply alternative labels for the same register. My account also allows for *gradations* in informativeness. Identity statements are more ore less informative, depending on the strength of cognitive links between the two registers in question (i.e. the Hesperus-register and the Phosphorus-register). Fiengo and May's account does not allow for such gradations and I don't see how they could explain why for some speaker's identity statements such as 'Hesperus is Phosphorus' are trivial. 'Hesperus' and 'Phosphorus' are expressions that contain different expression-forms (I would say name-forms) and will thus by default lead to *different* Assignments – thereby rendering the identity statement informative (Fiengo & May (2006): chapter 3).

The *correct* belief that 'elm' and 'beech' denote different kinds in the public language may suffice to keep the respective registers separate. If, however, the speaker believes that 'elm' and 'beech' behave like 'hedgehog' and 'urchin', such that those names are alternative labels for one and the same register, we may face a similar situation as in the inverse Paderewski puzzle – such that the speaker's 'elm' and 'beech' cannot be taken to be the same as our identically spelled words (even though they are linguistically linked to one another). This shows that the speaker's own (usually implicit) linguistic beliefs cannot be ignored in individuating her words.[360]

My proposal clearly is not a novel one. Kripke himself envisaged such a response in the original paper introducing the Paderewski puzzle:

> One might argue that Peter and we do speak different dialects, since in Peter's idiolect 'Paderewski' is used ambiguously[361] as a name for a musician and a statesman (even though these are, in fact, the same), while in our language it is used unambiguously for a musician-statesman. The problem then would be whether Peter's dialect can be translated homophonically into our own.[362]

Kripke entertains the possibility that Peter speaks a different dialect from ours, but ultimately rejects it. I believe that his preference for a homophonic translation stems from his epistemic commitments. I offered a translation manual for Peter's words that still allows for communicative success in situations where we are trying to pass on information about the world. But the passage goes on:

> Before he hears of 'Paderewski-the-statesman', it would appear that the answer is affirmative for his (then unambiguous) use of 'Paderewski', since he did not differ from anyone who happens to

[360] I hope that this approach may also lead to new insights into to the question of what implications anti-individualism has on self-knowledge – the knowledge an individual has about the content of her own thoughts.

[361] Kripke opted for a different notion for name than I. Where Kripke describes Peter's word 'Paderewski' as ambiguous (either referring to Paderewski as musician or statesman, but not both), I simply say that there are *two* distinct names. In the present context this difference is purely verbal and can be set aside.

[362] Krike (1979): footnote 37, p. 279.

> have heard of Paderewki's musical achievements but not of his statesmanship. Similarly for his later use of 'Paderewski', if we ignore his earlier use.[363]

Kripke's reflections illuminate a potential threat to my view. Prior to Peter reading the biography, he only has one Paderewski-register. At this point Peter is conceptually on par with the speaker who only knows of Paderewski as a musician. This lack of information would not prevent us from translating this speaker's name 'Paderewski_{I}' as the public name 'Paderewski_{P}' – thereby taking it at *face value*. Why should Peter's case then be any different, prior to his reading the biography that leads him to develop a second Paderewski_{2}-register and acquire a *new* name. So far, I assumed that none of Peter's names 'Paderewski_{I1}' or 'Paderewski_{I2}' ought to be translated as the public 'Paderewski_{P}'. If '$\text{Paderewski}_{I1\text{-at time t1}}$' is the same name as '$\text{Paderewski}_{P}$', while '$\text{Paderewski}_{P1\text{-at time t2}}$' is *not* the same name as 'Paderewski_{P}' it follows that *unbeknownst* to Peter '$\text{Paderewski}_{I1\text{-at time t1}}$' is not the same name as '$\text{Paderewski}_{I1\text{-at time t2}}$'. This creates new potential for Peter to make invalid inferences or believe contradictions, since Peter will be inclined to schematize two *distinct* names as one and the same – thereby being willing to infer conjunctions that do not logically follow (remember the reverse Paderewski puzzle). As a result our entire efforts would have been in vain!

I am not convinced that our initial intuitions for either taking both names 'Paderewski_{I1}' and 'Paderewski_{I2}' at face-value or *none* is well-founded. I wonder whether this doesn't stem from the consideration that the corresponding modes of presentations entertained by Peter "overlap" equally with the fully informed user's modes of presentation associated with 'Paderewski_{P}'. I have advertised against a notion of names in which modes of presentation individuate names as a solution to the Paderewski puzzle. I see no risk in proposing to translate the name that Peter

[363] Kripke (1979): footnote 37, p.279.

acquires *first* at face-value, which would avoid Kripke's objection, but not do so for the name or register he acquires later.

5.6 Conclusion

I wanted to do justice to the following intuitions: (A) Peter doesn't violate his *rationality* by asserting (i) and (ii), although only one of the two assertions can be true; (B) we *understand* what Peter wants to express with the utterances in question and Peter *understands* our public word 'Paderewski'. I propose a special normative connection between our public name 'Paderewski' and Peter's names 'Paderewski$_{I1}$' and 'Paderewski$_{I2}$'. Peter is obliged to merge his split concept into one, when he finds out that they both denote the same individual. My anti-individualist version of the register analogy led me to conclude that speakers are always right about the individuation of their own proper names. It remains to be shown whether this conclusion also applies to natural kind and artifact terms. If this hypothesis holds for words generally, we may have a new solution to the *prima facie* conflict between anti-individualism and self-knowledge – the knowledge an individual has about the content of her own thoughts. I will analyze this solution in the following chapter.

CHAPTER 6

ARE INDIVIDUAL SPEAKERS RELIABLE WHEN IT COMES TO INDIVIDUATING THEIR OWN WORDS? TEST-CASE: SLOW SWITCHING

6.0 Introduction

Oscar, an individual, ignorant of the chemical composition of water, is unknowingly switched from Earth to Twin Earth.[364] We are to assume the two planets physically (and historically) identical with one exception: we substitute Earth's water, a collection of H_2O molecules with twin-water, the chemical formula of which is entirely different and will be symbolized by 'XYZ'.[365] It is generally agreed that the Earth English word '$water_E$' denotes all and only water (H_2O) and that the Twin Earth English word '$water_{TE}$' denotes all and only twin water (XYZ).

With few exceptions[366], anti-individualists agree that while on Earth Oscar's utterance u_1 "$water_1$ quenches thirst" expresses that *water* quenches thirst, after some adequate amount of time spent on Twin Earth, Oscar's utterance u_2 "$water_2$ quenches thirst" expresses that *twin water* quenches thirst. Burge argues that judgments of our own (present) thoughts are self-verifying[367] - regardless of empirical knowledge of our social and physical environment. Oscar can reason as follows: "I know that I am (right now) thinking that water quenches thirst." If

---

[364] In order for there to be consistency in Oscar's phenomenal experiences (non-intentionally described) he must find on Twin Earth a "replica" of all his family and other social bonds. Oscar is to find a twin wife, twin children, twin friends, twin colleagues and twin acquaintances, who all have a history mimicking the one Oscar had with his actual wife, children, friends, colleagues and acquaintances. We will thus assume that Twin Oscar has been simultaneously switched from Twin Earth to Earth.

[365] We are to ignore complications such as the physical compositions of humans on Twin Earth. If there is no water on Twin Earth, then Twin Earthians couldn't possibly be physically identical to Earthians. It is also hard to see how the developmental histories of Earth and Twin Earth could have been even remotely similar. Such scientific concerns are to be set aside.

[366] Falvey (2003) argues that on Twin Earth Oscar's word 'water' denotes *zwater* (a merged concept of the concepts *water* and *twin water*). My own view is similar to Falvey's. The difference is that I claim Oscar's '$water_O$'-register denotes zwater on both planets.

[367] Burge (1988)

Oscar is on Earth, then his first-order thought, as well as his second-order thought will be about water. On Twin Earth Oscar will, in contrast, be thinking that twin-water quenches thirst and also interpret his thought as pertaining to twin-water. The content of both thoughts are determined by the *same* environment and are thus guaranteed to be about the *same* kind of stuff[368] – thereby rendering such judgments *authoritative*. There is disagreement about whether Oscar should have the ability to recall his earlier utterance u_1 and thereby be in a position to compare what u_1 asserts with his present utterance u_2. Oscar will assess the two utterances to express the *same* state of affairs, even though they express distinct thoughts.[369]

Boghossian has argued that the inability to recognize that some past thought τ_1 and some present thought τ_2 have different contents, threatens the individual's ability to draw *valid* inferences from τ_1 and τ_2[370] without first investigating her environment accordingly.[371] Social anti-individualism attempts to explain the transfer of knowledge across speakers. Taschek explains how

> there is no making sense of persons, in the relevant sense, without the possibility of assessing their rationality. This means at the very least, being able to assess their beliefs for logical consistency, to evaluate their inferences for validity, and to identify intersubjective agreements and disagreements between their beliefs and those of others.[372]

If Boghossian's objection is justified, then social anti-individualism fails to be able to explain how inter-subjective communication is possible in cases where the speaker uses a word that is intended as a repetition of nodes that belong to distinct repetition trees.

---

[368] It is generally agreed that it takes some time for Oscar to adapt to the new linguistic community and physical environment.
[369] I am taking for granted that there is no *distortion* of content between the content of thought and what (content) the corresponding utterance Oscar uses to convey his respective thought expresses. Presumably Oscar doesn't need language to access his own thoughts. Utterances are simply my means to talk about Oscar's thoughts as well as Oscar's means to convey his (more complex) beliefs to others.
[370] Any argument containing *both* thoughts are premises would threaten the individual's ability to reason discursively.
[371] Boghossian (1992).
[372] Taschek (1998): pp. 323-353.

Burge describes the main motivation for individualism as follows:

> [A]n individual's representational content depends (partly) on a web of inferential connections with other representational contents. The idea is that the constitutive conditions for understanding a concept cannot outrun the network of inferences that the individual can draw.[373]

Given that (social) anti-individualist's believe that a speaker can be partly under- or mis-informed about what her words denote[374], it follows that the content of the individual's concept can outrun the network of inferences that the individual can draw. What needs to be explained is how – in the face of switching scenarios – this doesn't threaten the *a priority* of her logical reasoning abilities.

Suppose that while water and twin-water dissolve salt, vitamin A is only soluble in twin-water.[375] Oscar is being told on Twin Earth: "Vitamin A is soluble in water"[376], which leads him to *recall* an utterance he heard on Earth: "Water dissolves salt."[377] He then reasons as follows:

(i) $Water_1$ dissolves salt.[378]

(ii) $Water_2$ dissolves vitamin A.

(iii) Thus, $water_3$ dissolves vitamin A and salt.[379]

Oscar's ability to draw a valid inference from (i) and (ii) pre-supposes that he is able to *keep track* of his word '$water_O$' throughout his reasoning process. It would suffice if *some* external

[373] Burge (2007): p. 176.
[374] Burge (1989, 2007): p. 261.
[375] On Earth vitamin A stands out for not being soluble in water, but oil.
[376] Oscar has not been exposed to utterances suggesting the contrary to Oscar on Earth.
[377] It is assumed that Oscar at the time of the original acquisition of τ_1 and τ_2 respectively was justified in trusting the corresponding testimonies.
[378] Let's assume that twin-water also dissolves salt. Nevertheless, what (ii) expresses doesn't follow from the original testimony, which only pertained to water and not twin water.
[379] Oscar will detect contradictions where there aren't any. If he were (unlike I assumed in the above example) to remember that he was once told "Vitamin A cannot be dissolved in water" when he is being told on Twin Earth "Vitamin A can be dissolved in water", then Oscar will falsely conclude that he has to choose between which of the two statements he believes to be true.

mechanism ensured that Oscar uses the *same* word 'water' in (i) through (iii). Gibbons, on the contrary, argues that the content of memory is *preserved* throughout the switch such that Oscar retains the ability to generate the concept *water* in addition to having acquired the *new* concept *twin water*[380] – henceforth referred to as 'the dual concept theory' (DC). Ludlow, Tye and Falvey contend that Oscar looses the ability to generate the concept *water* (after having spent an adequate amount of time) on Twin Earth[381] – which I name 'the replaced concept theory' (RC). While DC ascribes invalid inference reasoning to Oscar, RC avoids such a charge, since Oscar's word 'water' is assumed to denote the same substance from (i) to (iii). RC allows for Oscar to keep track of his own words throughout the argument, but we are left puzzling over how he would simply forget such an ordinary word as '$water_E$' (without any corresponding change in Oscar's cognitive apparatus).

Another strategy, pursued by Schiffer and Burge, combines the best of DC and RC. After he encounters *twin water*, Oscar is assumed *generally* capable of generating both concepts: *water* and *twin water*. At the same time throughout his argument (i) through (iii) it is also assumed that his word 'water' *anaphorically*[382] preserves reference (AR).[383] If Oscar's word 'water' can denote water in one context, but twin water in another, then Oscar *fails to keep track* of his word

---

[380] Gibbons (1996); Boghossian (1992) merely demonstrates some implications of this view; he doesn't endorse DC; Burge assumes that Oscar has the ability to generate the concepts *water* and *twin water*, but he denies that this means Oscar therefore draws invalid inferences (such as in (i) through (iii)).
[381] Ludlow (1995); Tye (1998); Falvey (2003).
[382] Usually the word applies to pronouns. Schiffer (1992) doesn't use the word 'anaphoric' to describe his proposal. Goldberg describes Schiffer's proposal as one involving "*anaphoric reference preserving*" (Goldberg (2007b): p.185). Schiffer imagines that an instructor believes two students with the generic name 'Yolanda' to be one and the same. One always attends his classes on Tuesday, the other on Thursday. "Ralph is thinking of the Thursday Yolanda under some vivid memory mode of presentation and, because he confuses her with the Tuesday Yolanda, thinks that she asked a very good question on Tuesday. To express this belief, he says: 'On Tuesday, Yolanda asked a very good question' […We] may say that the spoken token of 'Yolanda' refers unambiguously to the Thursday Yolanda" (Schiffer (1992): p.33). It is this process of reference transfer that I described as 'anaphoric'.
[383] Burge (1998); Schiffer (1992).

'water'. I have assumed that the reference of a given word *w* remains the same across time (and across individuals).

Goldberg explains how DC and AR (the same goes for RC) all threaten an individual's ability to reason discursively – either (a) the justificatory relation between premises and conclusion is broken, or (b) the epistemic relation between at least one premise and the experience through which it originally acquired justification is broken (i.e. a first-person perceptual experience or testimony).

I observe that the root of the epistemic implications is that DC, RC and AR all violate Oscar's ability to correctly individuate his own words across time and varying contexts, which also entails that speakers don't forget words altogether[384] – especially when it comes to such common words such as 'water'.

I will propose a solution that allows for Oscar to reliably keep track of his word '$water_O$' across time and varying contexts – from now on referred to as 'the rigid[385] register theory' (RR). Even though RR is compatible with the *a priority* of logical abilities, RR comes with its own negative epistemic implications. I conclude that we should favor RR over DC, RC and AR.

6.1 Unwarranted inferences

It is not sufficient for the content of Oscar's memories regarding (Earth) water to have been preserved throughout the switch to Twin Earth (and after), he also must be able to *access* them as memories pertaining to water. Some assume that the content of the second-order thought is

[384] Unless neurological conditions caused the loss of memory (i.e. Alzheimer's disease or amnesia).
[385] I am using the word 'rigid' to stress my very strong claim that words of a particular individual *cannot* under any circumstances change in meaning across time or contexts. I am not saying that the individual's word *w* has the same reference in all metaphysically possible worlds – since she may keep track of the various tokens in question (the ones that look like tokens of *w*) differently.

automatically the same as the first-order thought. When Oscar thus recalls his memory "water dissolves salt", he correctly remembers it as a memory about water (not twin-water).

> [We] can surely count on there being *syntactic* mechanisms which ensure that if a token of sentence S enters the belief box, then *ceteris paribus*, so does a token of the sentence 'I believe that *S*'. And so, provided that the embedded and unembedded tokens of *S* may always be counted upon to express the same content, the second order belief will always report correctly on the content of the first-order belief […].[386]

As a result, Oscar's utterances "$water_1$" and "$water_2$" will constitute performances of *different* words (at least if you assume, like I do, that in order for two non-indexical words *w* and *w'* to be one and the same, they must at least have the same reference). The token in premise (i) and the token in premise (ii) will not denote the same kind of stuff. In premise (i) Oscar recalls a memory from Earth that pertains to water. Since according to DC the content of memories is preserved (unless there is an accompanying cognitive change within the individual's mental apparatus), (i) is correctly remembered as a thought about water. Premise (ii), however, is a new belief the content of which is determined by Oscar's new environment and therefore pertains to twin water. Assuming that Oscar can adequately express his own thoughts, (i) is a claim about water, while (ii) is a claim about twin water. Oscar, however, is incapable of recognizing the difference between water and twin water thoughts. As a result, he judges utterances "$water_1$", "$water_2$" and "$water_3$" to constitute performances of the *same* word '$water_O$'. If we were to ask Oscar to formalize his argument[387] he would present us with something equivalent to the following:

Let 'Wx' stand for 'x is water', 'Ax' stand for 'x is vitamin A', 'Sx' stand for 'x is salt' and 'xDy' stand for 'x dissolves y':

---

[386] Boghossian (1992): p.15. Boghossian merely demonstrates the implications of this view, but does not endorse it himself; Gibbons (1996).

[387] We will assume Oscar competent in symbolic logic (I believe this assumption won't distort the outcome of this analysis).

(i') $(\forall x)(\forall y)((Wx \wedge Ay) \supset xDy)$

(ii') $(\forall x)(\forall y)((Wx \wedge Sy) \supset xDy)$

(iii') $(\forall x)(\forall y)(\forall z)((Wx \wedge Ay \wedge Sz) \supset (xDy \wedge xDz))$

(iii') states that all the stuff that Oscar calls 'water' dissolves vitamin A *and* salt. If (i') through (iii') is one of the correct ways of schematizing Oscar's argument in question, then he draws a *valid* inference.

When Oscar utters: (i) "water dissolves salt", he thereby meant to *repeat* the testimony ***T*** (heard on Earth) through which he acquired the corresponding belief τ. It was not one of Oscar's own observations of water (i.e. upon dropping a spoon of water and watching the crystals dissolve) that gave rise to τ. This means that the content of (i) is stipulated as not epistemically connected to Twin Earth experiences. Oscar never deliberates about how his current experiences may confirm his belief in (i) – such as when he drops salt into boiling water in preparation for cooking pasta. Oscar simply doesn't spend much time in the kitchen or wondering about food science. In addition, Oscar never learned on Earth that vitamin A doesn't dissolve in water. Hence, when he hears the testimony ***T'*** through which he forms the corresponding belief τ expressed by (ii), he doesn't have to negate any of his beliefs previously held true. He simply adds a new conception or attitude to his 'water_O'-register ("dissolves vitamin A").

If DC is correct in arguing that Oscar's memories of water are preserved and that he can access them as thoughts about water, then he commits the fallacy of equivocation. "True premises", Boghossian explain, "conspire, through a fallacy of equivocation that Peter is in principle not in a position to notice, to produce a false conclusion."[388]

[388] Boghossian (1992): p.22.

Ascribing characteristic *A* to entity *a* and characteristic *B* to entity *b*, does not imply that there is some entity *c* that displays characteristics *A* and *B*. This would only follow, if we knew entities *a* and *b* to be one and the same individual, or if they were groups with all the same individuals as members.[389] It is stipulated that this is not the case for water and twin water. If any amount of liquid is water, then it is not twin water and *vice versa*.

According to DC Oscar *should* have instead reasoned as follows:

> Let 'W_1x' stand for 'x is water'[390], 'W_2x' stand for 'x is twin water', 'Ax' stand for 'x is vitamin A', 'Sx' stand for 'x is salt' and 'xDy' stand for 'x dissolves y':
>
> (i") $(\forall x)\,(\forall y)\,((W_1x \wedge Ay) \supset xDy)$
>
> (ii") $(\forall x)\,(\forall y)\,((W_2x \wedge Sy) \supset xDy)$
>
> (iii") $(\forall w)\,(\forall x)\,(\forall y)\,(\forall z)\,((W_1w \wedge W_2x \wedge Ay \wedge Sz) \supset (wDy \wedge xDz))$

From (iii″) it does *not* follow that there is some substance that dissolves vitamin A and water, even though there may be such a substance (i.e. some oils or twin water).

If we interpret Oscar's word '$water_1$' to denote water and Oscar's word '$water_2$' to denote twin water, then (i″) and (ii″) are apt schemata for Oscar's premises (i) and (ii) respectively. His argument thus looks like the following:

(i") $(\forall x)\,(\forall y)\,((W_1x \wedge Ay) \supset xDy)$

(ii") $(\forall x)\,(\forall y)\,((W_2x \wedge Sy) \supset xDy)$

(iii*) $(\forall x)\,(\forall y)\,(\forall z)\,(((W_1x \vee W_2x) \wedge Ay \wedge Sz) \supset (xDy \wedge xDz))$

---

[389] Since I have assumed that twin-water dissolves salt and vitamin A, Oscar's conclusion may be true (if it indeed were to pertain to twin water). The question is of course not whether Oscar's conclusion happens to be true, but rather whether his premises (i) and (ii) jointly imply (iii).

[390] Oscar's second premise is assumed epistemically linked to a testimony he heard on Earth. As I will discuss in a moment, some philosophers argue that Oscar's epistemic link has been broken (even though it was "intact" while he lived on Earth). According to DC, Oscar's *knowledge* that water dissolves salt appropriately relies on the original testimony and this epistemic link is *not* broken by his change in location.

(iii*) doesn't follow from (i″) and (ii″). Oscar equivocates "apples with oranges", or water with twin water and thus infers a conclusion that reaches beyond the information provided by the premises. Sometimes we reason hastily, but this doesn't explain Oscar's mistake. No amount of careful reflection and reconsidering each step of the argument is going to put him in a position to see that he should have reasoned differently. Furthermore, the argument in question is not terribly complex to begin with – quite the contrary!

If we were to present Oscar with the alternative schemata (i″) to (iii″), he would tell us that we can formalize his sentences as proposed, but this *hides* the information that '$(\forall x)\ (W_1x \leftrightarrow W_2x)$'; consequently, the distinction between the two symbols 'W_1x' and 'W_2x' is unnecessary. Short of telling Oscar about his having been switched from Earth to Twin Earth, Oscar is unlikely to be convinced of the need for the above distinction. He never learned on Earth that water doesn't dissolve vitamin A. Hence, he has no experiences that would lead him to suspect that there are two types of water.

Unless somebody informs him of the switch in location, Oscar is simply not going to challenge his beliefs about the words 'water_E' and 'water_{TE}', due to the considerable similarities between water and twin water. He is going to continue to believe that the two words are one and the same and are correctly to be translated as his word 'water_O'. If Oscar is unable to even minimally assert the distinction between what 'water_1' and 'water_2' denote respectively, can we interpret (i) and (ii) as schemata (i″) and (ii″)? Recall Putnam's self-proclaimed ignorance as to what characteristics distinguish an elm from a beech tree. Putnam was not able to specify how elms are different from beeches, but he was, nevertheless, minimally aware of a distinction between those two types of trees. I argued that the content of his 'elm'-register and the content of his 'beech'-register are identical – in the sense that he ascribes the same characteristics to each

type of tree. Putnam nonetheless keeps information about what these words denote separate. This warrants interpreting his words 'elm_P' and '$beech_P$' as *distinct* words. Oscar, on the other hand, keeps only one '$water_O$' register stored in his cognitive apparatus.

DC may simply insist that sometimes speakers are wrong about how they store information pertaining to the entities they talk about. They may further point out that this whole reference to registers is merely an explanatory tool anyways. It is quite unlikely that we would ever be able to scientifically prove that they exist. We are able to demonstrate that some activities consistently activate one part of the brain (or several specific areas of the brain), while other activities consistently activate different parts of the brain. It is doubtful that we could fine-tune this method to apply it to word-registers. Perhaps such a spatial correlation with registers doesn't even exist. That said, I stipulated the existence of registers on the basis of the speaker's intra-subjective word individuations – these intra-subjective judgments DC violates regardless of whether registers exist in the cognitive apparatus of the individual. Unless Oscar is generally correct about how he keeps track of his own words, we can't get word types off the ground. Recall my observations in chapter one, as well as Kaplan's claims that the graphic and phonetic features of utterances aren't enough to group word utterances into word-types. Our investigation of how we communicate inter-subjectively pre-supposes the existence of word-types. We thus concluded that we have to rely on the speaker's own intra-subjective judgments about how she keeps track of her words by grouping utterances. This means that we cannot suddenly question this ability at a later stage without offering reasons why those particular judgments of Oscar's are *false* without having this skepticism permeate *all* his intra-subjective word individuations – thereby undermining our entire project. Does the fact that Putnam's thought-example is far-

fetched to begin with contain this skepticism? According to Ebbs such skepticism about self-knowledge undermines our ability to even describe such Putnam style thought examples.[391]

Either way, mistakes in formal reasoning should not be explained by reference to false beliefs regarding the empirical world.

6.2 Monist proposals

By arguing that Oscar cannot generate thoughts about both entities at any given time, RC blocks the attribution of invalid reasoning to Oscar. The ability to produce thoughts about water is lost when he becomes competent in the concept *twin water*[392] – throughout Oscar's conceptual apparatus the concept *twin water* replaces the concept *water*. RC entails that Oscar's '$water_O$'-register changes topic and either abruptly or somehow suddenly switches its reference from water to twin water.

When Oscar is being told on Earth "water dissolves salt" he forms the belief that *water* dissolves salt through what was asserted by the original testimony *T*. After several years of immersion in the Twin Earth English linguistic community, Oscar hears "water dissolves vitamin A." RC agrees with DC that Oscar is linguistically competent in the Twin English word '$water_{TE}$' denoting twin water and thereby comes to store the belief that twin water dissolves vitamin A. The content of his memory which he expresses with "water dissolves salt" switches, once Oscar acquires the concept *twin water*, his "memory" pertains to twin water.[393]

[391] Ebbs (2003).
[392] Ludlow (1995); Tye (1998).
[393] It is puzzling how two thoughts with significantly different contents are supposed to represent the same memory.

RC avoids ascribing bad inference reasoning to Oscar, since his utterance 'water' is only linked to *one* public word repetition tree at any given time, thereby rendering the schematizing of his argument apt. According to RC Oscar correctly schematizes his reasoning as presented by (i') through (iii'). Furthermore, Oscar doesn't make false comparative judgments of his words '$water_E$' and '$water_{TE}$', since once he is able to perform an instance of the word '$water_{TE}$' he thereby looses the ability to utter "$water_E$". Without the ability to bring to mind both words, he is unable to make a false comparative judgment as to whether they constitute the same or different words.

Heal points out how Oscar will have many referential intentions that *conflict* with twin water being the referent for his word '$water_O$'.[394] Oscar may say things such as "water is the stuff I swam in as a child" (and twin water is not the stuff Oscar swam in as a child). However, I have pointed out in chapter three, how such conflicting referential intentions cannot always break the link to the public repetition tree (here, the public repetition tree being '$water_{TE}$'), or else we risk to undermine anti-individualism (*à la* Burge) altogether. An important ingredient to anti-individualism is that speakers can be wrong about some of the characteristics of the things their words denote. Consequently, Oscar may simply be wrong about having swum in twin water as a child. Still, Heal's concerns give rise to the question as to why we should *prioritize* the referential intentions that would select twin water as the referent for Oscar's word '$water_O$' over those that would select water as the correct referent. Should we value referential intentions less because they were formed further back in time? Perhaps. But what if the individual recalls those characteristics she learned *first* about what she interprets to be one topic more vividly?

[394] Heal (1998).

RC comes with significant costs. It is unclear whether we can ever make sense of some of RC's implications. Even though, according to RC, Oscar doesn't make a false judgment as to whether two of his words are one and the same, something even more puzzling happens: Oscar completely forgets a rather ordinary word, such as '$water_E$'. There are many words that we haven't used for years, but they don't thereby entirely slip our cognitive apparatus. Certainly, if we were presented with a relevantly similar word-form, it would lead us to recall the word in question. When I argued that speakers reliably keep track of their own words (unless some neurological damage has occurred), this also entails that speakers do *not* randomly forget words. In order for individuals to keep track of their own words, they must be able to keep track of the corresponding registers. Words are labels for registers.

The topic of Oscar's register '$water_O$' *changes* independently of any cognitive alterations of the register (non-intentionally described). Registers change over time because we refine and enrich them by adding, removing and ideally correcting our conceptions of the respective entities they are about. With the exception of adding the belief that "water can dissolve vitamin A" he didn't change his beliefs about $water_O$.[395] The corresponding register changes what it pertains to – it changes the *topic*. If registers change what they pertain to over time, that is, switch whatever stuff they were denoting from reference group A to reference group B (the two reference groups being distinct), then w_1 and w_2 could constitute performances of the same word, even though w_1 denotes a different reference class than w_2. Non-indexical words ought to be interchangeable in a sentence without changing the truth-value of the sentence as a whole. Dropping such a fundamental assumption will have far-reaching consequences, which we ought to avoid.

[395] This of course means we ignore his extended experiences with *new* particulars that he identifies as water: such as 'this glass is filled with water'. Oscar doesn't think that this much extends his knowledge of "water", since he already keeps identifying the same kind of stuff as water over and over again – which of course is false. Oscar is identifying glasses filled with twin water as "water".

When Oscar looses the ability to generate the concept *water*, he thereby also *forgets* all his thoughts that pertained to water. At the same time Oscar is *unaware* of having forgotten anything, since component parts of these memories were substituted with *twin water*. As a result, Oscar's memories turn into "mis-memories". When Oscar utters: "I used to swim in water as a child" he is now stating that he used to swim in twin-water as a child – which is *false* (he used to swim in water). The same goes for all his thoughts that constitute "memories" about particulars of that physical world on Earth. He somehow cannot entertain any thoughts that pertain to individual objects (including sentient and rational "objects") of his past habitat on Earth. The truth-value of a large portion of his thoughts (that were at some time about life on Earth) *changed* from *true* to *false*. Only memories that don't pertain to particulars (including the particular planet) or to the natural kind water still remain unaffected by the switch. It is very hard to come to terms with how a change in location – one Oscar is unaware of – can have such drastic effects on Oscar's mental life. How can we explain Oscar's *collective* forgetting of all thoughts that pertain to the concept *water* and the respective individuals from Earth?

Since RC agrees that Oscar used to know that water dissolves salt, how can he cease to have this knowledge without having made any changes in his corresponding 'water'-register?[396] Nevertheless, it is agreed that if we could point at Oscar having acquired *new* relevant beliefs, then he may have changed his mind about a previously held belief. Falvey suggests:

> [N]ew information can lead one to question a previous judgment that was in fact warranted.[397]

---

[396] Boghossian (1992).
However, it is unclear whether Oscar forgets anything – his memory's content *switches* without him noticing the change. I agree with Boghossian that something puzzling is happening with Oscar's cognitive apparatus. I am not sure whether we can argue that Oscar doesn't forget anything, if he indeed ceases to be capable to generate the concept water and thus looses a good deal of knowledge.
[397] Falvey (2003): p. 223.

As Falvey himself notes, this observation doesn't seem to help us much with the puzzle at hand, since Oscar certainly doesn't identify his *new* experiences with twin water as *new* information. It brings us closer to the key question: does Oscar's change in location have any significant effects on Oscar's cognitive apparatus? What supposedly happens is not that he adds new conceptions to his 'water'-register, it is that the *topic* of his register changes altogether. This would of course explain the loss of knowledge – if Oscar cannot generate thoughts about water, then he also couldn't possibly have knowledge of the corresponding substance. I explained above how this approach is conceptually dangerous.

Goldberg demonstrates how approaches like RC that allow for Oscar to keep track of his word '$water_O$' in (i) through (iii) threaten Oscar's ability to reason on the basis of *discursive justification* – "the kind of justification that accrues to a belief acquired via inference." RC entirely focuses on whether Oscar's reasoning is valid, but thereby overlooks dangerous epistemic implications.

> If I [= Oscar] am not presently [= at t_r] suffering from a relevant cognitive malfunction and have not acquired any relevant new evidence regarding any of my premise-beliefs since the time I acquired them, the degree of justification of my belief in the conjunction of the premises at t_r is exclusively determined by the degrees of justification enjoyed by each of my premise-beliefs, individually, *at the time each was acquired*.[398, 399]

In Oscar's case, however, the epistemic link between premise (i) and the experience that gave rise to premise (i) is *interrupted*. At some point, Oscar's belief τ expressed by "water dissolves salt" has adequate epistemic justification through testimony *T*. This epistemic link is mysteriously broken, without any corresponding change in Oscar's cognitive apparatus (nothing changes "inside" his mind).

---

[398] Goldberg (2007b): p. 189.

[399] Goldberg seems to take for granted that the belief in question cannot change its epistemic warrant over time. I may, for instance, believe that the birth-rate in Germany is 0.8 per person (and that this is not going to change for several years). This is the kind of belief that may cease to be true over time. Germany's government could adopt drastic measures to increase the birth-rate of its citizens.

> [I]nsofar as the hearer aims to acquire knowledge of how the world is *through* that testimony, he (the hearer) must recover how that testimony has represented the world to be. And this, of course, is another way of saying that the hearer must have understood the testimony.[400]

According to RC, when Oscar first is exposed to the testimony *T*: "water dissolves salt" (on Earth) he *understands* the claim *T* makes about the world and about water specifically. From the time Oscar acquires competence in our word 'water$_E$' he remains competent in 'water$_E$' throughout his stay on Earth.[401] Upon hearing *T* Oscar forms the adequate belief that water dissolves salt (τ). *T* thus *transfers its epistemic warrant* onto τ – since Oscar understands how *T* has presented the world to be and adequately stores *T*'s content as τ.[402] Premise (i) is meant to simply restate τ. Here we are facing an identity question: should we say that (i) mis-represents τ, or instead claim that (i) correctly states τ's content (but the content of τ changed)? Either way the epistemic link to the original testimony *T* is *interrupted*, since Oscar is unable to recall its content (without the ability to generate thoughts about water, he cannot recall the content of *T*).

Goldberg seems to assume that τ and what (i) states are one and the same thought (but that the thought changed its content over time). Of course, if we cannot keep track of the content of our own thoughts in this manner, such that at t_1 τ claims the word to display state of affairs s_1, but that at t_2 τ claims the world to display state of affairs s_2 (and s_1 is a different state of affairs from s_2), then it shouldn't come as a surprise that the degree of justification of τ would quite likely have changed from t_1 to t_2 (without the individual having acquired relevant new evidence). To me the core of the problem is whether we accept the possibility that an individual's mental

[400] Goldberg (2007a): p. 42.
[401] It is not entirely clear when Oscar is supposed to loose his competence in 'water$_E$'. We do know it supposedly happens some time between his leaving Earth and his acquiring competence in the word 'water$_{TE}$'.
[402] Oscar doesn't possess knowledge that would increase the degree of justification of what *T* expresses.

register can *change the topic* from water to twin water[403,404] as assumed by RC. I maintain that a mental register doesn't change in content (over time or across contexts).

Of course, one could take the converse approach to RC and argue that Oscar instead never acquires the concept *twin water* and only ever stores beliefs about water in his '$water_O$' register. The advantage of attributing such adaptation failure to Oscar (AF) is that we don't have to explain Oscar's excessive loss of knowledge and the undetected substitution of the effected beliefs by false beliefs. According to AF Oscar's reasoning from (i) to (iii) is valid.

As anti-individualists we would have a hard time explaining why Oscar never picks up the public concept *twin water* no matter how much time he spends on Twin Earth. How can Oscar directly perceive twin water and appear to engage successfully in substantial conversations about twin water, without having picked up the corresponding public concept? Again, this reference failure would pertain to all names for particulars (i.e. the name of Oscar's twin wife[405]) in his environment that constitute twins to particulars Oscar had already been acquainted with on Earth. Oscar's success in deferring to his Earth (English) community is in accordance with anti-individualism; his later (partial) failure to defer to his Twin Earth linguistic community seems puzzling. In chapter four I offered an example where my prior membership in the German linguistic community interfered with acquiring competence in the ordinary English word 'angina'. What makes my own case different from Oscar's – even though we both intended to defer to the respective linguistic community in question and still failed – is that I had *significant* beliefs about the characteristics of what my word 'angina' denotes that are incompatible with my

[403] Or as later discussed change content from water to zwater.
[404] I suppose it would also be possible for the water-register to cease to exist altogether and be replaced simultaneously by a twin-water register. It seems to me that in such a case there would be no epistemic continuity between the original testimony T and what premise (i) states. Oscar would simply *forget* τ and form a *new* belief τ^*.
[405] The twin wife is really Twin Oscar's wife. Oscar and the twin wife never exchanged marriage vows.

having spoken about angina. I was furthermore aware of my interacting with *distinct* linguistic communities: the German and the English ones. Oscar isn't even aware of the *possibility* of distinct words being in use because he believes he is engaging with the *same* linguistic community throughout the time span of the thought example. He furthermore doesn't seem to have false referential beliefs about twin water. Since AF doesn't seem a viable proposal for anti-individualists, I consider it a dead end.

Falvey offers a clever version of RC[406], where Oscar's water concept doesn't get replaced by the concept *twin water*, but instead by a disjunctive concept (RC') *zwater*.[407] He proposes that Oscar's "judgment purports only to be about a kind of stuff, and there are disjunctive kinds of stuff, even if they are not natural kinds (jade is a frequently cited example)."[408] Oscar's utterance "this is water" is true if and only if what he points at is either water or twin water. We cannot attribute to him a disjunctive claim such as "this is water or twin water" since this would entail that Oscar is minimally aware of the two kinds – this is why I prefer to attribute to Oscar a *merged* "concept". Oscar is not even minimally aware of a distinction between the two kinds of water.

Just like RC, RC' allows us to interpret Oscar's argument as *valid*, since all utterances of 'water' therein are instances of the same word '$water_O$' denoting *zwater*. This, however, implies that somehow Oscar lost track of his own word '$water_O$', in the sense that at some point in the past (on Earth) it denoted water and at present (on Twin Earth) it means *zwater* – while Oscar is unaware of any change of reference. Since he lost the ability to produce an instance of his

[406] Falvey (2003) and I disagree on Oscar's memories from Earth such as "water quenches thirst" pertaining *only* to water. I will argue that if Oscar's '$water_O$'-register matures as one that pertains to water *and* twin-water, then we have to apply this to all of Oscar's past utterances.
[407] I suppose 'zwater' stands for 'zwitterwater', where 'Zwitter' is a German word to be translated as 'hybrid' into English.
[408] Falvey (2003): p.229

original word 'water$_O$' denoting water, I can't argue that he incorrectly judges instances of the two words to be performances of the *same* word. Without the ability to produce instances of both words, Oscar is not in a position to make such comparative judgments.

Falvey takes our word 'jade' to be evidence for the possibility of such disjunctive concepts.[409] Many people may be under the false impression that 'jade' denotes *one* natural kind whereas it denotes *two* different natural kinds: jadeite and nephrite (distinct chemical substances). Jade is nevertheless a homogenous kind of a sort: it is (typically) a green gemstone often carved into small to medium artifacts. Without instruments enhancing the human senses, jadeite is indistinguishable from nephrite. Suppose Patrick, a person who is oblivious to the two types of jade, were to observe a sample of jadeite and come to conclude: "jade has about the same hardness as quartz" (Patrick assumes that jade is a natural kind). As it happens, this property is not shared by nephrite (which is slightly softer than jadeite). Patrick's sample was biased. In anticipation of concerns that Patrick's word 'jade' only denotes jadeite (and not nephrite), let's also stipulate that Patrick also interacted with samples of nephrite. Suppose his wife has a necklace made out of nephrite beads and Patrick has not dared to test those for their hardness properties. Patrick would describe the piece of jewelry as a "jade necklace".

Patrick reasons as follows:

(iv) Jade has about the same hardness as quartz.

(v) Jade is a metamorphic rock.

[409] I concur.

(vi) Thus, jade is a metamorphic rock with about the same hardness as quartz.[410]

I think most of us would agree that Patrick simply has the *false* belief that *all* jade has about the same hardness as quartz, but that his inferences are valid. This is merely supposed to show that it is quite possible for a person to store information *about* two distinct natural kinds in one register – even though it is not stored *as* information about two distinct kinds. That is, from the perspective of Patrick's registers, there is only *one* kind of jade.

The same goes for Oscar's argument. The conclusion follows from the premises, but it is false that all zwater dissolves salt and vitamin A (since the latter is not soluble in water). Of course, Oscar's situation is more complicated than Patrick's. The public word 'jade' (in English) also expresses such a disjunctive concept and thus when Patrick uses his word '$jade_P$' we don't have to explain failure to defer to competent speakers of the public word 'jade'.

Even though RC' allows for some continuity between the '$water_O$'-register Oscar has on Earth with the one he has on Twin Earth (once he has spent an adequate amount of time on this planet), there is still a change in topic from *water* to *zwater*. At least RC' allows for an *overlap* between what the register pertains to on Twin Earth and what it pertains to on Earth, I discussed how problematic it is to allow for a mental register to change its topic. According to my own theory such a radical change – a change in topic, rather than in the conceptions stored in the register – causes for a *new* register to be formed, while the respective "muddled" register gets retired.

[410] If we were to teach Patrick first order logic, this is how Patrick would formalize his belief: Let's assume 'Jx' stands for the predicate 'x is jade'; 'Qx' stands for the predicate 'x is quartz'; 'Rx' stands for the predicate 'x is a metamorphic rock'; 'xHy' stands for the relation that 'x has the same hardness as y'.

(v) $(\forall x)\,(\forall y)\,(Jx \wedge Qy \supset xHy)$
(vi) $(\forall x)\,(Jx \supset Rx)$
(vii) $(\forall x)\,(\forall y)\,(Jx \wedge Qy \supset xHy \wedge Rx)$

The resulting break in continuity of Oscar's memories pertaining to water still remains puzzling. At some point Oscar had the experience of swimming in water. This event brings about the later memory: I used to swim in water as a child. While Oscar resides on Earth (and doesn't forget the event in question non-intentionally described), he *knows* that he used to swim in water as a child. Once the substitution of the concepts in question occurs in Oscar's cognitive apparatus, his utterance: "I used to swim in water as a child," expresses that Oscar used to swim in *zwater* as a child. Falvey admits that

> [it] is not correct to say that [Oscar] remembers that [he swam in] zwater [as a child], since he did not know this at the time. Partly for this reason, and the fact that the change in concept is undetectable [to Oscar], we may not want to say that [his] present judgment represent knowledge of [the bathing event].[411]

RC' like RC implies that Oscar suffers from a puzzling loss of knowledge, which we can't find a satisfying explanation for. Falvey moderates this objection by observing how Oscar's "error is mitigated by the fact that *zwater* is a broadening of the concept that did figure in [his original] thought."[412]

It is not very believable, however, that a change in location would bring about such selective amnesia. This psychological condition is usually caused by some traumatic experience or the result of physical injury. The switch couldn't have been traumatic for Oscar, since he wasn't even aware of it. It also doesn't seem part of the story that Oscar is injured during the process. It remains puzzling why Oscar would loose such large amounts of knowledge.

While in RC the epistemic link to the original testimony about water "water dissolves salt" is entirely broken, since Oscar fails to generate thoughts about water on Twin Earth altogether, according to RC' the link remains partly intact, since Oscar's word '$water_O$' on Twin Earth means *zwater*. Oscar is not able to generate thoughts about water specifically, but he is still able

[411] Falvey (2003): p.240.
[412] Ibid.

to generate thoughts that are partially about water – he can generate thoughts that are about water or twin water without awareness of the latter two being distinct entities. This seems to be an advantage RC' has over RC.

Even though RC' improves the epistemic link between the testimony *T* and what premise (i) states, there is still a loss in degree of justification between the belief formed at the time when Oscar first hears *T* and the time he reasons through the argument in question (even though he doesn't acquire new relevant information). When Oscar thinks τ for the first time, τ is justified through testimony *T*. On Twin Earth Oscar ceases to understand $T - \tau$ somehow switched content and turned into a "mis-memory" (it now pertains to zwater, whereas it pertained to water at some earlier time). Even though it happens to be true that water and twin water both dissolve salt, premise (i) is *not* formed on the basis of an observation Oscar made while interacting with twin water; it's justificatory status epistemically depends on a testimony heard on Earth.

According to RC' Oscar can still generate the concept *water*[413], he just can't generate it in isolation from *twin water*. Even though RC' has the same problematic implications as RC in kind they are somewhat more palatable in their severity. There is *somewhat* less of a memory loss and there is at least a partial epistemic link to the original testimony that premise (i) links to.

We cannot, however, overlook a newly introduced challenge by RC'. Since Oscar's word 'water_O' on Twin Earth denotes zwater, while his interlocutors' words 'water' denote twin water[414], Oscar is constantly talking past his interlocutors – at least in part. Premise (ii) is actually *false* and was never justified. Oscar meant to repeat a testimony *T'* he heard on Twin

[413] The concept *water* that picks out H_2O.
[414] We assume that Oscar is the only one that underwent the switch from Earth to Twin Earth. Since we assumed that Twin Oscar simultaneously was switched from Twin Earth to Earth and thus presumably, according to RC', would also acquire the concept *zwater*, Twin Oscar would be the only one who could understand Oscar. Of course, they are stipulated to never meet…

Earth. Oscar, however, never understood that *T'* asserted twin water to dissolve vitamin A. Oscar interprets *T'* to make the claim that zwater dissolves vitamin A – which is false (since not all zwater dissolves vitamin A).

Schiffer and Burge propose an anti-individualist solution that avoids some of the problems RC and RC' introduce, without accusing Oscar of an invalid inference (AR).[415] It is argued that Oscar's word '$water_O$' denotes the same throughout his reasoning (i) to (iii), due to an anaphoric reference preservation within a given context. According to Schiffer, "[w]hatever mode of presentation is operative for [Oscar] in the first utterance will also be operative with him in the second: for his intention is to be speaking about the same [stuff]."[416]

As per to AR Oscar can generally generate both concepts *water* and *twin water* (even though he is not aware of their distinctness). We thus don't have to explain a general loss of knowledge of thoughts pertaining to water or individuals from Earth. Oscar still has the memory "I swam in water as a child" and it states correctly that Oscar swam in water as a child. It is only when Oscar thinks of twin water within the same context that he gets his "memories" mixed up. This means that Oscar's '$water_O$'-register is ambiguous: depending on context it either pertains to water or twin water.

Schiffer has us imagine a person, Ralph, being unable to keep two students separate. Ralph has beliefs about them stored in the same mental register labeled 'Yolanda'. Since proper names can only have one referent (at a time), we are supposed to conclude that when Ralph utters 'Yolanda' he only refers to one student (not both). For the puzzle in question, it doesn't matter which individual Ralph's name 'Yolanda' refers to, since all Schiffer needs to show is that Oscar's word '$water_O$' denotes the same stuff throughout his reasoning process.

---

[415] Schiffer (1992); Burge (1998).
[416] Schiffer (1992): p.33.

An odd result of this proposal is that what premises (i) and (ii), as well as the conclusion (iii) express depends on the order within which the premises are presented. Surely, this order is arbitrary and shouldn't settle how we ought to interpret what the argument as a whole pertains to.

AR still implies a loss of epistemic connection to the experience that originally lent justification to one of the premises.[417] This time the culprit is premise (ii), since the occurrence of 'water_O' anaphorically receives its reference from premise (i). Premise (ii) lost its epistemic link to testimony *T'*, the content of which gave rise to the belief τ that (ii) expresses.

AR is furthermore incompatible with my assertion that speakers can reliably individuate their own words (across time and varying contexts). The reference of Oscar's word 'water_O' changes across contexts sometimes denoting water, at others twin water. Perhaps Schiffer would reply that Oscar still keeps track of the same word – but it happens to be an ambiguous word. If that is so, Oscar should be *aware* of using an ambiguous word, since his word 'water_O' will not be substitutable across contexts without changing the truth-value of the statement that 'water_O' is a component part of.

6.3 Merged registers

Unless the individual suffers from some cognitive malfunctioning, the degree of justification of a conclusion-belief is derived from the degree of justification of the conjunction of the respective premise-beliefs. The degree of justification of each premise-belief goes back to some event of acquisition – some event that caused the individual to assert the premise-belief. In the case where the premise-belief was acquired as a result of another person's testimony, the premise belief τ is justified to the degree the original testimony *T* was justified assuming that the

[417] Goldberg (2007b).

individual understood T (and assuming that the individual has not and does not acquire any additional relevant evidence that bestows greater epistemic warrant on the premise belief). According to Goldberg, the degree of justification of τ ought to remain constant across time (unless the individual acquires new relevant evidence).[418] If the content of τ is allowed to *change* over time, without the individual's notice (without some conscious change in the individual's cognitive apparatus) – then it is quite likely that a belief's justification is not preserved through memory. In other words, the epistemic link of the belief to its justificatory experience is broken. If we want to argue for the *a priority* of logical abilities, we must not only take into account the validity of an argument, but also consider that the premises are warranted due to some original experience. If the epistemic link to that experience is interrupted (i.e. Oscar ceases to understand the testimony), then the respective premise ceases to be justified.

I will offer a *new* solution that will neither ascribe invalid reasoning to Oscar, nor argue that his premise-beliefs fail to preserve their justification through the switch from Earth to Twin Earth. I will achieve this by arguing that Oscar keeps track of his own word 'water_{O}' – that is, Oscar does *not* forget his word 'water_{O}' which he acquired on Earth and he is able to individuate his own utterances "water" correctly across time (and across varying contexts). At the same time, Oscar doesn't fail to link his 'water_{O}'-register to the 'water_{TE}'-repetition tree (the word 'water_{TE}' on Twin Earth). Finally, Oscar's 'water_{O}'-register doesn't change reference (over time and across contexts) – it claims registers to be rigid (RR). However, the only reason why the degree of justification of Oscar's premise-beliefs remains constant over time is that, since Oscar never really understood the respective premises, the justificatory link between testimony T and

[418] Goldberg (2007b).

Oscar's resulting belief τ was always defective – τ was never adequately justified to begin with (this is true for both premise-beliefs (i) and (ii)).

You may wonder whether this means that Oscar was never able to communicate with his fellow Earthlings about water and whether he ever effectively communicates with his fellow Twin Earthlings about twin water – using the various performances of his word 'water_O' (which are all linked to the *same* register). It is Oscar's commitment to repeat our word 'water' and his willingness to make the adequate changes to his 'water'-register – if he finds out about the switch and the two kinds of water – that establish a semantic link between 'water_E' and 'water_O'. Again, I must stress that this semantic link is weaker than identity – this does not make Oscar's performance 'water_O' a successful repetition of our word 'water_E', precisely because the same commitment exists with regard to the Twin English word 'water_{TE}' (denoting twin water). I assume that Oscar's word 'water_O' cannot constitute a repetition of nodes stemming from distinct repetition trees. At the same time Oscar's word 'water_O' is a *merged* concept of two *public* repetition trees – he is not free to use it however he wants.

Once Oscar finds out about the two distinct liquids water and twin water and the corresponding public words 'water_E' and 'water_{TE}' Oscar must *split* his 'water'-register into two separate registers, such as to avoid drawing invalid inferences. This re-organizing of his registers is what allows Oscar to *unlock* the original testimonies (that gave rise to his beliefs (i) and (ii)). His later split concepts 'water_{O1}' and 'water_{O2}' are each respectively repetitions of the public words 'water_E' and 'water_{TE}'. There is semantic continuity between 'water_O' and 'water_{O1}' (as well as 'water_{O2}'). Some of the conceptions previously associated with 'water_O'

were *transferred* to '$water_{O1}$' (the same ones or others get transferred to '$water_{O2}$').[419] Nevertheless, until Oscar performs this splitting of his '$water_O$' register, he is talking past his interlocutors when uttering "water". While testimonies acquired through his informants' use of the word '$water_E$' or '$water_{TE}$' are defective – the information contained in those testimonies can be *unlocked* as described above.

I thus agree with DC and RC' that at the time he reasons through his argument in question Oscar indeed is linked to both public repetition trees '$water_E$', as well as '$water_{TE}$' – but I clearly identify it as a *defective* link. RR furthermore concurs with RC' that it is one and the same '$water_O$' mental register that is linked to *both* repetition trees *at the same time*. As a result, Oscar's formal logical reasoning isn't faulty. Given that all of Oscar's utterances "water" are linked to the *same* register[420], Oscar draws a *valid* conclusion (iii). His word 'water' denotes water as well as twin water without distinguishing the two entities as distinct kinds.

What distinguishes RR from RC' is that RR doesn't allow for Oscar's memories to switch in content – even those that contain the word '$water_O$'. So, Oscar's register '$water_O$' doesn't undergo a change in what it denotes. Since I also want Oscar to be (at least) "partially successful" at communicating with his Twin Earth community when it comes to conversations about twin water, I will have to argue that Oscar's word '$water_O$' means *zwater* throughout – on Earth, as well as Twin Earth. From that it follows that Oscar's word '$water_O$' already denotes twin water (as well as water) prior to his encountering twin water. Admittedly, RR appears *prima facie* unbelievable – perhaps even crazy. Please bear with me; I will address those concerns in a moment.

---

[419] This transferring of conceptions is tricky to do correctly. Oscar may not know which conceptions correctly belong to which register. Peter didn't have this problem when merging his '$Paderewski_1$' and '$Paderewski_2$' registers.

[420] At least up to the point when he finds out that his word '$water_O$' is really linked to words denoting distinct natural kinds.

Oscar's mistake is not to be located in his inference reasoning; he simply comes to a *false* conclusion because he assumes a *false* premise. What Oscar predicates in (ii) of the set of things his word 'water' denotes is only true of a subset thereof. Even though the corresponding original testimony that Oscar meant to repeat in (ii) is true, (ii) is false.

> [I]nsofar as the hearer aims to acquire knowledge of how the world is *through* that testimony, he (the hearer) must recover how that testimony has represented the world to be. And this, of course, is another way of saying that the hearer must have understood the testimony.[421]

Since Oscar fails to distinguish between water and twin water and the original testimony claimed that vitamin A is soluble in twin water, Oscar has misunderstood what the original testimony meant to convey. The same goes for premise (i) – even though (i) happens to be true, it is based on a testimony Oscar failed to understand.

The downside of RR is that Oscar never manages to successfully communicate with his linguistic community about water or twin water (at least not until he finds out about the switch in location); this is true on Earth as well as Twin Earth. When Oscar utters "water", he is talking about zwater, while others are either talking about water or twin water. Neither in Earth English nor in Twin Earth English does a word exist that expresses the disjunctive concept that Falvey and I intend to ascribe to Oscar. My assumptions also commit me to this communication failure – but I argue that it can be rectified through Oscar's re-organizing of his registers accordingly (that is, he must split his 'water_O'-register).

The upside is that since the topic of Oscar's register 'water_O' doesn't change, (a) Oscar's reasoning is valid; (b) there is no loss of knowledge related to his respective memories (i.e. "I swam in water as a child"); (c) Oscar doesn't forget the word 'water_E' (since he was never linguistically competent in its use); (d) Oscar reliably individuates his utterances "water" across

[421] Goldberg (2007a): p. 42.

time and across contexts; and (e) the degree of justification of his premise-beliefs is preserved across time.

It is difficult to make believable that Oscar's word '$water_O$' is "linked" to the '$water_{TE}$' repetition tree, even before he interacts with people who have experienced twin water and before Oscar directly perceives twin water. Keep in mind that from first exposure to full-fledged competence, the register is one and the same. While the register's content is growing it remains the *same* register and it pertains to the same content throughout its development.

In chapter four I explained how a speaker, who encounters a word for the first time will create a mental register with an accompanying label. Let's first analyze the non-problematic case, where the individual is not unknowingly switched to a twin-environment. Let's call this individual 'Kevin'. When Kevin first encounters the word '$water_E$' at t_0, he creates a register with the label '$water_K$'. Kevin may have picked up the word, without having any idea what it denotes (except for inferring from its grammatical role that it is a kind word – perhaps even a mass noun). I want to allow for Kevin to be using the same word '$water_K$' from first exposure to the public word '$water_E$' to the point (t_n) where he has filled the register with *substantial information* such that his register has matured. Even though at t_0 Kevin's register doesn't contain any information that would single out water as the correct referent and we cannot consider him yet a competent participant of the '$water_E$'-speaking practice, we interpret his word as denoting water. It is his future commitment to that practice that *retroactively bestows reference* onto Kevin's word '$water_K$' at t_0[422]. I want us to recognize that even in normal cases, we already allow for the future to partly individuate an individual's words.

[422] Kevin's word '$water_K$' does *not* undergo a change in reference from first exposure to '$water_E$' which leads him to create a '$water_K$'-register to performances of '$water_K$' when his register has matured. The difference between the early and the later performances is that Kevin acquires competence in the word

Oscar's case, even though it strikes us *prima facie* as puzzling, is not really extraordinary. I am not violating anti-individualism, since Oscar's word 'water_O' is linked to some *public* repetition tree – what is unusual, is that his word is connected to *more than one* repetition tree. Oscar "owes" commitment to *both* linguistic communities, which partially explains why he is not entirely failing in his communicative attempts – he successfully coordinates his behavior with his fellow speaker through his word 'water_O' and interpreting their words 'water_E' and 'water_{TE}' to constitute the same word as his word 'water_O'. As stated above, my account implies that Oscar is (from an epistemic point of view) *not* successfully communicating about a single topic with his interlocutors. When Oscar talks about zwater, others talk to him about water (or twin water).

Once Oscar finds out about the different kinds of water and the switch in location, he *will* split his merged register into separate registers. This *commitment* is what makes him a member of *both* public practices – Oscar would regard himself as a member of both linguistic communities, if he were to be informed about there being two linguistic communities linked to his word 'water_O'.

Some may find it unconvincing that in the two parallel worlds – one where Oscar is never switched to Twin Earth W_1 and another, where Oscar does get switched to Twin Earth W_2 – the two Oscars will have had identical experiences up to a certain time *t* (prior to the switch), they even live in the same environment, but the Oscar in W_1 denotes water with his word 'water_O', while Oscar in W_2 denotes zwater with 'water_O'. As an anti-individualist, I would question

---

'water_K' *gradually*. That is, only once his register is matured is he entitled to make substantial claims about water using 'water_K'. Since in the normal case, Kevin is assumed to develop competence in the *public* word 'water_E', his performances of the word 'water_K' also constitute performances of the word 'water_E'.

whether the appropriate individuation of words is *determined* by mental states *plus* the individual's physical and social environment.[423]

RR still has some unpalatable implications, but it seems to fare better than any existing alternative (as shown by the following table):

	DC	RC	RC'	AR	RR*
Valid reasoning	–	+	+	+	+
No puzzling *loss of knowledge*	+	–	–	+	+
No puzzling *loss of words*	+	–	–	+	+
Reliable *word-individuations*	–	+	+	–	+
Degree of justification constant	+	–	–	–	+
Intact *epistemic link* to testimonies	+	–	–	–	–

6.4 Conclusion

My analysis lends further evidence in support of Goldberg's observation that anti-individualism introduces serious epistemic costs that need to be considered. Assuming anti-individualism, there is no straightforward solution that keeps Oscar's discursive reasoning intact without introducing other significant costs to his way of acquiring knowledge of his environment. My own view is that we cannot even begin to make sense of how we acquire

[423] Ebbs (2009) goes so far as to argue that *nothing* makes our word-individuations true apart from our actual interpretations of words (PJSSs).

knowledge of the external world without first taking for granted that we can keep track of our own words (across time and contexts). RR is the only solution to the problems posed to anti-individualism by switching scenarios that takes seriously Oscar's ability to individuate and not inexplicably forget his own words.

BIBLIOGRAPHY

Alward, Peter –

(2005) "Between the Lines of Age: Reflections on the Metaphysics of Words," in *Pacific Philosophical Quarterly*, Vol. 86: pp. 172-187.

Bach, Kent –

(1987) *Thought and Reference*, Clarendon Press.

Begby, Endre –

(Forthcoming) "Concepts: their alleged role in cognition and communication." Retrieved from http://www.google.com (with title in the search-line).

Bilgrami, Akeel –

(1992) *Belief and Meaning – The Unity and Locality of Mental Content*, Blackwell Publishers.

Boghossian, Paul –

(1992) "Externalism and Inference," in *Philosophical Issues*, Vol. 2: pp.11-28.

Buchanan, Ray –

(2005) co-authored with Ostertag, Gary "Has the problem of incompleteness rested on a mistake?" in *Mind*, Vol. 114, No. 456: pp. 889-913.

Burge, Tyler –

(1979, 2007[424]) "Individualism and the Mental," reprinted in *Foundations of Mind*, Oxford University Press: pp. 100-150.

(1986, 2007) "Intellectual Norms and Foundations of Mind," reprinted in *Foundations of Mind*, Oxford University Press: pp. 254-274.

(1988) "Individualism and Self-Knowledge," in *The Journal of Philosophy*, Vol. 85, No. 11: pp. 649-663.

(1989, 2007) "Wherein is Language Social," reprinted in *Foundations of Mind*, Oxford University Press: pp. 275-291.

(1998) "Memory and Self-Knowledg," in *Externalism and Self-Knowledge*, (eds.) Ludlow, Peter and Martin, Norah, CSLI Publications: pp. 351-370.

(2007) "Postscript to "Individualism and the Mental"," in *Foundations of Mind*, Oxford University Press: pp. 151-181.

Cappelen, Herman –

(1999) "Intentions in Words," in *Noûs*, Vol. 33, No. 1: pp. 92-102.

[424] If two dates are listed, the first denotes the date of the original publication and the second, denotes the date of the listed reprint.

Crimmins, Mark –

(1989) co-authored with Perry, John "The Prince and the Phonebooth," in *The Journal of Philosophy*, December: pp: 685-711.

(1992) *Talk about Beliefs*, MIT Press.

Donnellan, Keith –

(1966, 2008) "Reference and Definite Descriptions," reprinted in *The Philosophy of Language 5th edition*, (ed.) Martinich, Aloysius.

Ebbs, Gary –

(1998) "Bilgrami's Theory of Belief and Meaning," in *Philosophy and Phenomenological Research*, Vol. 58, No. 3: pp. 613-620.

(2002) "Learning from Others," in *Noûs*, Vol. 36, No. 4: pp. 525-549.

(2003) "A Puzzle about Doubt," in *New Essays on Semantic Externalism and Self-Knowledge*, (ed). Nuccetelli, Susan, MIT Press: pp. 143-168.

(2009) *Truth & Words*, Oxford University Press.

Evans, Gareth –

(1973, 1985) "The Causal Theory of Names," reprinted in *Collected Papers*, Oxford University Press.

(1982) *The Varieties of Reference*, (ed.) Mc Dowell, John, Oxford University Press.

Falvey, Kevin –

(2003) "Memory and Knowledge of Content," in *New Essays on Semantic Externalism and Self-Knowledge*, (ed.) Nuccetelli, Susan, MIT Press: pp. 219-240.

Ferbezar, Ina –

(2008) co-authored with Stabej, Marko "Razumeti razumevanje," in *Jezik in slovstvo*, 53: pp. 15-31.

Field, Hartry –

(1980) *Science without Numbers*, Princeton University Press.

Fiengo, Robert –

(2006) co-authored with May, Robert, *De Lingua* Belief, The MIT Press.

Fine, Kit –

(2009) *Semantic Relationism*, Wiley Blackwell.

Frege, Gottlob –

(1892) original title "Über Sinn und Bedeutung," in *Zeitschrift für Philosophische Kritik*, 100. Reprinted (1947) as "On Sense and Nominatum," Feigl, Herbert (transl.), Feigl, Herbert and Sellars, Wilfrid (eds.), in *Readings in Philosophical Analysis*, Appleton-Century-Crofts.

Forbes, Graeme –

(1989) "Cognitive Architecture and the Semantics of Belief," in *Midwest Studies in Philosophy* 14: pp. 84-100.

(1990) "The Indispensability of Sinn," in *The Philosophical Review*, XCIX, 4: 535-563. (2006) *Attitude Problems – An Essay on Linguistic Intensionality*, Oxford Press.

George, Alexander –

(1990) "Whose Language is it Anyway? Some Notes on Idiolects," in *The Philosophical Quarterly*, Vol. 40, No. 160: pp. 275-297.

Gibbons, John –

(1996) "Externalism and Knowledge of Content," in *The Philosophical Review*, Vol. 105, No. 3: pp. 287-310.

Goldberg, Sanford –

(2007a) *Anti-Individualism*, Cambridge University Press.

(2007b) "Anti-Individualism, Content Preservation, and Discursive Justification," in *Noûs*, Vol. 41. No. 2: pp.178-203.

Golding, Andrew –

(1999) co-authored with Roth, Dan "A Winnow-Based Approach to Context-Sensitive Spelling Correction," in *Machine Learning*, 34: pp. 107-130.

Goodman, Nelson –

(1947) co-authored with Quine W. V. O. "Steps towards a Constructive Nominalism," in *Journal of Symbolic Logic*, Vol 12, No. 4: pp. 105-122.

(1970) "Seven Strictures of Similarity," in *Experience and Theory*, ed. Foster, Lawrence and Swanson, J. W., University of Massachusetts Press: pp. 19-30.

(1976) *Languages of Art*, Hackett Publishing company.

Grice, Paul –

(1956) co-authored with Strawson, Peter "In Defense of a Dogma," in *The Philosophical Review*, Vol. 65, No. 2: pp. 141-158).

(1961) "The Causal Theory of Perception," in *Proceedings of the Aristotelian Society*, vol. xxxv: pp. 121-153.

(1989) *Studies in the Way of Words* Harvard University Press.

Haslwanter, Thomas –

(1996) co-authored with Moore, Steven "A Geometric Basis for Measurement of Three-dimensional Eye Position Using Image Processing," in *Vision Res.*, vol. 36, no. 3: pp. 445-459.

Hawthorne, John –

(forthcoming) co-authored with Lepore, Ernest "On Words." A version of this paper has been presented at the Pacific American Philosophical Association meetings in April 2010. Retrieved from http://ruccs.rutgers.edu/faculty/lepore/OnWords.pdf

Heal, Jane –

(1998) "Externalism and Memory," (Part II), in *Aristotelian Society Supplementary Volume*, Vol. 72, Issue 1: pp.95-109.

Heck, Richard –

(1995) "The Sense of Communication," in *Mind*, vol. 104, no. 413: pp. 79-106.

(2006) in *Content & Modality: Themes from the Philosophy of Robert Stalnaker*, eds. Thompson and Byrne, Clarendon Press.

Janssen, Maarten –

(2004) co-authored with Visser, Albert "Some Words on *Word*," in *La Nuova Critica*, Nr: 43-44: pp. 71-95.

Kaplan, David –

(1989) "Afterthoughts," in (eds.) Almog, Perry, and Wettstein *Themes from Kaplan*, Oxford University Press: pp. 598-599.

(1990) "Words," in *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vo. 64: pp. 93-119.

King, Jeffrey –

(1998) "What is a Philosophical Analysis," in *Philosophical Studies*, Vol. 90, No. 2: pp. 155-179).

Kripke, Saul –

(1972) *Naming and Necessity*, Harvard University Press.

(1979), "A Puzzle About Belief," in (ed.) Margalit, A., *Meaning and Use*, Dortrecht: D. Reidl: pp. 239-283.

Kukich, K. –

(1992) "Techniques for automatically correcting words in text," in *ACM Computing Surveys*, 24 (4): pp.377-439.

Lepore, Ernest –

(forthcoming) co-authored with Hawthorne, John "On Words." A version of this paper has been presented at the Pacific American Philosophical Association meetings in April 2010. http://ruccs.rutgers.edu/faculty/lepore/OnWords.pdf

Loar, Brian –

(1985) "Social Content and Psychological Content," in *Contents of Thought*, (eds) Grimm, Robert and Merrill, Daniel, The University of Arizona Press: pp. 128-129.

May, Robert –

(2006) coauthored with Fiengo, Robert, *De Lingua* Belief, The MIT Press.

Mercier, Adèle –

(1993) "Normativism and the Mental," in *Philosophical Studies*, Vol. 72, No.1: pp. 71-88.

Moore, Steven –

(1996) co-authored with Haslwanter, Thomas "A Geometric Basis for Measurement of Three-dimensional Eye Position Using Image Processing," in *Vision Res.*, vol. 36, no. 3: pp. 445-459.

Neely, Erica –

(2005) *Dissertation: Competence, Expertise and Linguistic Communities*, University of Illinois at Urbana-Champaign.

Ostertag, Gary –

(2005) co-authored with Buchanan, Ray "Has the problem of incompleteness rested on a mistake?," in *Mind*, vol. 114, no. 456: pp. 889-913.

(2007) "Review of Robert Fiengo, Robert May, *De Lingua* Belief," in *Notre Dame Philosophical Reviews*, Vol. 9: http://ndpr.nd.edu/review.cfm?id=10985

Owens, Joseph –

(1989) "Contradictory Belief and Cognitive Access," in *Midwest Studies in Philosophy*, Vol. 14: pp. 289-316.

(1995) "Pierre and the Fundamental Assumption," in *Mind & Language*, Vol. 10, No. 3: pp. 250-273.

Pagin, Peter –

(2006) "When does communication succeed?," in (ed.) Kotatko, P., *The Constraints on Interpretation*, Cambridge Scholar Press, Cambridge.

Paul, Matthias –

(1999) *Success in Referential Communication*, Kluwer Academic Publishers.

Peirce, Charles Sanders –

(1903,1998) "Nomenclature and Divisions of Triadic Relations, as Far as They Are Determined," reprinted in *The Essential Peirce*, Indiana University Press: p. 291.

(1908, 1998) "Excerpts from Letters to Lady Welby," in *The Essential Peirce*, Indiana University Press.

(1931-58) *Collected Papers of Charles Sander Peirce*, Hartshorne and Weiss (eds.), Harvard University Press: vol. 4, par. 537).

Perry, John –

(1989) co-authored with Crimmins, Mark "The Prince and the Phonebooth," in *The Journal of Philosophy*, December: pp: 685-711.

Putnam, Hilary –

(1973, 1975) "Explanation and Reference," reprinted in *Mind Language and Reality – Philosophical Papers Volume 2*, Cambridge University Press: pp. 196-222.

(1974, 1975) "Language and Reality," reprinted in *Mind Language and Reality – Philosophical Papers Volume 2*, Cambridge University Press: pp. 272-290.

(1975, 1975) "The Meaning of Meaning," reprinted in *Mind Language and Reality – Philosophical Papers Volume 2*, Cambridge University Press: pp. 215-271.

(1981) *Reason, Truth and History*, Cambridge Press.

(1988) "Meaning, Other People, and the World," in *Representation and Reality*, MIT Press.: pp.19-41.

(2001) "Meaning, Other People, and the World," in *Representation and Reality*.

Quine, W. V. O. –

(1947) co-authored with Goodman, Nelson "Steps towards a Constructive Nominalism," in *Journal of Symbolic Logic*, Vol 12, No. 4: pp. 105-122.

(1960) *Word and Object*, MIT Press.

(1992) *Pursuit of Truth*, Harvard University Press.

Recanati, François –

(1993) *Direct Reference: From Language to Thought*, Blackwell Publishers.

(2004) *Literal Meaning*, Cambridge University Press.

Rosch, Eleanor –

(1975) "Cognitive representation of semantic categories," in *Journal of Experimental Psychology* 104(3): pp. 192-233.

Roth, Dan –

(1999) co-authored with Golding, Andrew "A Winnow-Based Approach to Context-Sensitive Spelling Correction," in *Machine Learning*, 34: pp. 107-130.

Salmon, Nathan –

(1986) *Frege's Puzzle*, MIT Press.

(1989) "Illogical Belief," in *Philosophical Perspectives*, 3: pp. 243-285.

Sas, Jerzy –

(2006) "Handwritten Recognition Accuracy Improvement by Author Identification," in *ICAISC* 2006, LNAI 4029: pp. 682-691.

Schiffer, Stephen –

(1977) "Naming and Knowing," in *Midwest Studies in Philosophy*, II: pp: 28-41.

(1987) "The 'Fido'-Fido Theory of Belief, in *Philosophical Perspectives*, I: pp. 455-480.

(1992) "Belief Ascriptions," in *The Journal of Philosophy*, Vol. 89, No. 10: pp. 499-521.

(forthcoming) "Propositions, What Are They Good For?," R. Schantz (ed.), *Current Issues in Theoretical Philosophy: Prospects for Meaning Vol. 3*. Walter de Gruyter.

Schroeter, Laura –

(2008) "Why be an Anti-Individualist?," in *Philosophy and Phenomenological Research*, July Vol. 77, No. 1: pp. 105-141.

Soames, Scott –

(1987) "Direct Reference, Propositional Attitudes, and Semantic Content," in *Philosophical Topics*, 15: pp. 47-87.

(1988) "Substitutivity," in JJ. Thomson (ed.), *On Being and Saying: Essays in Honor of Richard Cartwright*, MIT Press.

(2002) *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*, Oxford Press.

Strawson, Peter –

(1956) co-authored with Grice, Paul "In Defense of a Dogma," in *The Philosophical Review*, Vol. 65, No. 2: pp. 141-158).

Tarski, Alfred –

(1994) "The Semantic Conception of Truth: and the Foundations of Semantics," in *Philosophy and Phenomenological Research*, vol. 4, no.3: pp. 341-379.

Taschek, William –

(1998) "On Ascribing Beliefs: Content in Context," in *The Journal of Philosophy*, No. 7 July: pp. 323-353.

Tye, Michael –

(1998) "Externalism and Memory," (Part I) in *Aristotelian Society Supplementary Volume*, Vol. 72, Issue 1: pp. 77-94.

Verdonik, Darinka –

(2010) "Between understanding and misunderstanding" in *Journal of Pragmatics* 42: pp. 1364-1379.

Verschueren, Jeff –

(1999) *Understanding Pragmatics*, Hodder Arnold Publication.

Visser, Albert –

(2004) co-authored with Janssen "Some Words on *Word*," Maarten, in *La Nuova Critica*, Nr: 43-44: pp. 71-95.

Weigand, Edda –

(1999) "Misunderstanding: The standard case" in *Journal of Pragmatics* 31: pp. 763-785.

(2000) "The Dialogue Action Game," in *Dialogue Analysis VII. Working with Dialogue*, Coulthard, M. (ed.), Niemeyer: pp. 1-18.

Wetzel, Linda –

(1993) “What are occurrences of expressions,” in *Journal of Philosophical Logic*, 22: pp. 215-220.

(2000) “The Trouble with Nominalism,” in *Philosophical Studies*, 98: pp. 361-370.

(2006) “Types and Tokens,” entry in the on-line *Stanford Encyclopedia of Philosophy*.

(2009) *Types and Tokens*, MIT Press.