Regine Eckardt, Göttingen

Speakers and Narrators

The paper investigates a seeming paradox. Formal theories of meaning standardly assume that sentences – in real communication or in fiction – are evaluated relative to utterance contexts. Therefore, such theories lead us to expect that texts in fiction always refer to speakers, which come as part of these utterance contexts. Intuition, on the contrary, tells us that many stories do not create the impression that there is someone who is telling us all this. I will discuss how this paradox can be resolved within formal semantics.

I Introduction

Who tells us the works of literary fiction? Since the invention of print, the tradition of story-telling in the marketplace has been replaced by stories in written form. But who is the person telling us such a story? Leaving folk epics aside, we can assume that every piece of fiction was written by an author. Does this mean that every such story has a narrator?

This does not follow, at least not in a simple manner. Clearly, the author and the person who tells us a work of literature cannot be simply equated. For instance, first-person narratives usually have a narrator who is clearly distinct from the author of the story. The *Sherlock Holmes* stories are told by Holmes's friend Dr Watson but were written by Arthur Conan Doyle, who is an altogether different person. But even stories without a first-person narrator can communicate beliefs and attitudes on the part of a narrating person which are not shared by the author. Hence, author and narrator must be kept separate.

Given that author and narrator are not necessarily identical, one might go further and ask whether it could happen that a story has an author but no narrator. Several answers to this question have been defended.

According to one position, *a story always has a narrator*. This position is not only preferable for the sake of generality. It also seems to find support by the following formal argument: Every sentence (in languages like English, German, French etc.) is specified for tense. Tense morphology of the verb relates the content of the sentence to utterance time. Past tense morphology codes that the events in question happened before utterance time, future forms code that they will take place later, and present tense morphology codes that events happen as they are being reported (historical present). All sentences of stories carry tense. Therefore, all stories need an utterance time to be interpreted. If there is an

utterance time, then there is an utterance situation. If there is an utterance situation, there is someone who makes an utterance. Hence we have a narrator.¹

According to an opposing position, *there could be stories without a narrator*. Intuitively, there are stories that give us no clue about the person who might be telling it. Such stories should most appropriately be classed as stories without a narrator. Hence, we should try to get rid of the idea that there is always a speaker.²

The question is whether and how these two positions can be reconciled. Proponents of narrator-free stories will have to spell out how an utterance can depend on utterance time without the notion of someone uttering. Proponents of uniform context dependence and ubiquitous speakers have a hard time explaining what kind of speaker we are supposed to imagine for a story that gives no clue of such a person, apart from the fact that it is told after some given series of events. In this paper, I propose how these two intuitions can be accounted for in a uniform theory of meaning for sentences and text.

The paper is structured as follows. In section II, I recapitulate the classical notion of sentence meaning in possible-world semantics in the philosophical tradition going back to Montague, Lewis, Stalnaker, and Davidson. This includes a model for information update which captures how the reader gains information while s/he reads a text. The model can be applied to fiction as well as informative text. In section III, this basic theory is extended by reference to utterances, making use of Kaplan's theory of context. I argue that the content of fiction must be modelled by what authors after Kaplan called the diagonal (of sentences/texts), or utterance content.³ This may be surprising, in view of the fact that texts of literary fiction, if anything, are hardly ever "uttered" in the physical sense of

¹ Restricting attention to formal semantic theory, this position is entailed by the joined body of literature about tense, context, and context dependence e.g. David Kaplan, "Demonstratives", in: John Almog/John Perry/Howard Wettstein (eds.), *Themes from Kaplan*, Oxford 1989, pp. 481–563; combined with Arnim von Stechow, *Tenses in Compositional Semantics* [Ms. University of Tübingen], 2009, <http://www.sfs.uni-tuebingen.de/~astechow/Aufsätze/Approaches.pdf> (accessed August 2013).

² This position has been defended e.g. in Andrew Kania, "Against the Ubiquity of Fictional Narrators", in: *The Journal of Aesthetics and Art Criticism*, 63/2005, pp. 47–54; Thomas E. Wartenberg, "Need there be implicit narrators of literary fictions?", in: *Philosophical Studies*, 135/2007, pp. 89–94; or Tilmann Köppe/Jan Stühling, "Against Pan-Narrator Theories", *Journal of Literary Semantics*, 40/2011, pp. 59–80.

³ Ede Thomas Zimmermann, "Kontextabhängigkeit", in: Dieter Wunderlich/Arnim von Stechow (eds.), *Semantics/Semantik. Ein Internationales Handbuch*, Berlin 1991, pp. 156–229; Ede Thomas Zimmermann, "Context dependence", in: Klaus von Heusinger/Claudia Maienborn/Paul Portner (eds.), *Semantics. An International Handbook*, Vol. 3, Berlin 2012, pp. 2360–2407.

the word. Section III concludes with a brief survey of different sources of information about the speaker. Section IV takes a closer look at a specific linguistic way to convey speaker information: speaker-oriented expressions. These expressions play an important role in literary texts. Their function will be illustrated in various modes of literary writing, we will look at their semantic values, and explain how they differ from traditional indexical expressions like *I*, *you* or *now*. Section V combines all these points and addresses the question “How much does the text tell us about the person speaking?” It illustrates speaker information in the present analysis, on basis of a range of sample text passages. We can trace sources of information about the speaker and the effects they have on the semantic representation of the text. Section VI, finally, returns to the speaker/narrator puzzle. We will combine two main slogans of the analysis: First, all sentence contents are represented by sets of *worlds/people/contexts* about which they are true. Second, sentences which cover many *worlds/people/contexts* are sentences which convey little information. This correlation between many options and little knowledge is a standard pattern used in theories of semantic content. Applied to the narrator/speaker question, we predict that many options for possible speakers reflect that we don’t know much about the speaker. The fiction of a narrator, however, requires some contentful restrictions on possible speakers. The greater the number of possible speakers, the less lively our sense of a narrator. In the extreme case, the fictional narrator (in the sense that the reader is led to imagine that someone is telling the story) can be completely absent.

II Worlds and Story Content

The present section introduces a first analysis of sentence content and story content based on possible worlds. We will also spell out how the literal story content interacts with the reader’s knowledge, and how stories are enriched by a reader’s expectations.⁴

What happens when we read and understand a novel, story or other piece of fiction? A gripping novel can lead us away from our own world, carry us into new worlds, and offer us new, exciting experiences. Many people describe this experience as the reader “immersing herself into the world of the story.”

⁴ The analysis goes back to Robert Stalnaker, “Assertion”, in: Peter Cole (ed.), *Syntax and Semantics*, Vol. 9, New York 1978, pp. 315–332; Robert Stalnaker, *Context and Content*, Oxford 1999; Robert Stalnaker, “Common Ground”, in: *Linguistics and Philosophy*, 25/2002, pp. 701–721; and was discussed for the case of fiction in David Lewis, “Truth in Fiction”, in: *American Philosophical Quarterly*, 15/1978, pp. 37–46.

For any given story, however, different readers can have different opinions on what the world of the story should look like. Different readers imagine parts of the story in very different ways, depending on their expectations and experiences. Usually, these different ways to imagine the story-world are equally well supported by the text. Stories can simply leave aspects of the world undetermined.⁵

In search of an intersubjective notion of sentence content, semanticists and philosophers put to use possible worlds in a different manner. They propose that sentence meanings are characterized by their truth conditions. For instance, a sentence like '*ravens are black*' is true in exactly those worlds where ravens actually are black.⁶ The meaning of a sentence can be characterized as the set of all those imaginable worlds in which the sentence is true. For instance, '*ravens are black*' denotes the set of all those worlds w where ravens are indeed black.⁷

Stalnaker extended this notion of sentence meaning to a model of information update and information increase.⁸ He argued that no human possesses complete knowledge about the actual world. This means that for many imaginable worlds w we cannot tell whether w is the actual world or not. At best, we can decide that w coheres with all we know about the actual world (it *could* be the real world), or that it doesn't. For any person **A**, the set of worlds which, according to **A**'s knowledge, could be the actual world, is also called the *epistemic alternatives* of **A**.

$$\text{Epi}(\mathbf{A}) = \{w \mid \mathbf{A} \text{ thinks that } w \text{ could be the world he really lives in}\}$$

A can acquire more information when he hears a sentence and believes that it is true. Imagine, for instance, that **A** so far has no beliefs about the color of ravens. His epistemic alternatives include worlds w_1 where ravens are red, worlds w_2

5 Sometimes, experts can argue that ways of reading a story are 'wrong' because the author could not have intended this way of enriching the plot. I will not attempt to spell out the exact borderline between permissible and false enrichments here. Note that Lewis, "Truth" emphatically sides with those who only allow interpretations which are in the spirit of the author.

6 Donald Davidson, "Truth and Meaning", in: *Synthese*, 17/1967, pp. 304–323.

7 My language use deviates from the philosophical standard. I will use the term *imaginable worlds* for philosophers' "possible worlds" and reserve the term *possible worlds* to talk about worlds in which a story *possibly* could take place. Of course, I do not suggest that the range of *imaginable worlds* depends on any specific person's power of imagination. They cover exactly what Lewis would call "possible worlds." Paul Portner, *What is Meaning? Fundamentals of Formal Semantics*, Malden, MA, 2005, offers an accessible introduction to the basic assumptions of truth conditional semantics; see also Irene Heim/Kai von Stechow, "Intensional Semantics", in: *Online Textbook*, 2007, <http://semantics.uchicago.edu/kennedy/classes/s08/semantics2/vonstechow+heim07.pdf> (accessed December 29, 2011).

8 See Stalnaker, "Assertion"; Stalnaker, "Common Ground".

where they are multi-colored like parrots, worlds w_3 where they are black, and so on. Now, an informed and trustworthy biologist **B** tells **A** the following:

Ravens are black.

B thereby invites **A** to *update* his epistemic alternatives with the set of worlds denoted by this sentence. If **A** believes **B**'s claim, he should restrict his epistemic alternatives to those worlds where *Ravens are black* is true.

Updated epistemic alternatives of **A**:

$$\begin{aligned}\text{Epi}(\mathbf{A})_{\text{new}} &:= \{ w \mid \mathbf{A} \text{ thinks that } w \text{ could be the world he really} \\ &\quad \text{lives in, and ravens are black in } w \} \\ &= \text{Epi}(\mathbf{A})_{\text{old}} \cap \{ w \mid \text{ravens are black in } w \}\end{aligned}$$

We intersect the epistemic alternatives of **A** with the set of worlds which represents the sentence *Ravens are black*. Update with sentences *S* for **A** is tantamount to restricting his former epistemic alternatives to those in which *S* holds true. The more knowledge **A** possesses, the *smaller* is the range of imaginable worlds w which he thinks could be the actual world. An increase in knowledge is reflected by a decrease of options. This pattern will be used at many later points in this paper.

Stalnaker was mainly concerned with information exchange and increase of knowledge. Yet, his analysis can be extended to fictional texts.⁹ The content of fiction is not supposed to increase our knowledge about the actual world. Hence, the content of sentences in fictional stories should not update the reader's epistemic alternatives (or else, we'd predict that the reader takes the content of the story as literal truth about the world). Fiction invites the reader to construct a new set of imaginable worlds, the worlds which are such that the story could have taken place in them. Let us assume that the reader, at the beginning of a story, supplies an initial set of worlds of the following kind:

$$\text{Story}(\mathbf{A})_0 = \{ w \mid \text{the present story could have happened in } w \}$$

Even the most widely conceived sets of story-worlds obey certain restrictions which reflect the expectations of **A** about the story. Usually, **A** will assume that the laws of physics hold true in all story worlds w (unless the story is science fiction). If the story was written 200 years ago, **A** will assume that the story worlds w do not contain things like cars, nuclear power plants or electric razors. These

⁹ See Lewis, "Truth".

restrictions are not imposed by the literal content of stories but by the reader's knowledge. We will likewise assume that **Story(A)₀** excludes the real world right from start. This ensures that fiction is never misinterpreted as information about the real world, even if the content of the story incidentally happens to match actual events in the real world.¹⁰

Different readers will start reading with different expectations. An expert on Indian jungles will read Kipling's *Jungle Book* with more specific knowledge than a reader who lacks this background. While **Story(A)₀** may not be too close to the belief worlds of the reader – after all, the reader expects fiction to deviate from the real facts of the world – it may become clear over time that the author intended some facts and generalizations to pattern with those of the real world. For instance, an author can write about a fictitious politician in such a way that typical psychological features of people in politics can best explain the character's actions, even though the story itself does not contain an explicit lecture on the psychological profile of politicians. This could be one of the ways in which we can find truth in fiction.

The reader **A** might need some time to figure out the intended possible worlds which the story starts from. I will use **Story(A)₀** in this more flexible manner as a label for the background information supplied by reader **A**, and allow that this information can become more specific as the story unfolds. For the sake of the present paper, it will be sufficient to assume that **Story(A)₀** can be corrected and adjusted over time. A detailed analysis of a story, sentence by sentence, might require a truly dynamic update of background and story content in alternation, a problem that will be left aside here.

Stalnaker's theory of information and information update can now be extended to the case of fictional texts. Once again, each sentence of the story denotes the set of worlds in which that sentence holds true. As before, the reader updates her information background with each sentence in sequence. The reader's set **Story(A)₀** is reduced step by step to those worlds where all sentences of the story hold true. However, it will always be clear that the set of worlds which emerges is *not* a set of worlds which is supposed to include the actual world. It is also understood that this set of worlds belongs to that particular story and is not supposed to be combined with the sets of possible worlds which belong to other stories.¹¹ At the

¹⁰ Lewis, *ibid.*, argues that the telling of fiction is an act of make-believe assertion: the speaker acts as if she was reporting real knowledge, which she isn't. He concludes that for this reason alone, the worlds described in fiction can never be the actual world, even if some story by chance actually happened in the way reported in fiction.

¹¹ Sequels are an obvious exception; see Lewis, "Truth" for careful ways to deal with inter-story inconsistencies and information transfer.

end of the story, the range of imaginable worlds that the story could be about has become much smaller than the original set **Story(A)₀**. Increase of information is again reflected by a decrease of options. The meanings of sentences are stable objects, no matter whether the sentence occurs as part of TV news or part of a novel. The content of sentences contributes to different sets of worlds; this is how the analysis captures the difference between news and fiction.

Finally, Stalnaker's account of information update can be generalized to a theory of information about discourse referents, as exemplified in Discourse Representation Theory,¹² File Change Semantics,¹³ and Dynamic Montague Grammar.¹⁴ In order to capture anaphoric cross-references in story texts, these frameworks propose that sentence and story content should be represented by sets of discourse referents.¹⁵ These are used to reflect possible groups of protagonists about which a sentence or text conveys true information. As in Stalnaker's account, an increase in information about discourse referents is reflected by a decreasing choice of imaginable people (and things) about which the story could truthfully be told. Example (1) illustrates the basic idea.

1. A woman arrived at Heathrow Airport. She was wearing red shoes.

The first sentence restricts the set of imaginable people to female persons who arrive at Heathrow Airport (at some salient time). This, obviously, leaves us with many possibilities. The second sentence adds more information about the person in question: She moreover wears red shoes. Taking these two properties together will lead to a much more restricted set of women the story could be about. Each subsequent sentence will offer further restrictions about the woman, her previous or subsequent behaviour, her clothes, other people she interacted with and so on. At the end of the story, we might still not know *the* one and only imaginable person that the story is about. But we have got a pretty specific idea what kind of

12 Hans Kamp, "A Theory of Truth and Semantic Representation", in: Jeroen Groenendijk et al. (eds.), *Formal Methods in the Study of Language*, Amsterdam 1981, pp. 1–41; Hans Kamp/Uwe Reyle, *From Discourse to Logic*, Dordrecht 1993.

13 Irene Heim, "File Change Semantics, and the Familiarity Theory of Definiteness", in: Rainer Bäuerle/Christoph Schwarze/Arnim von Stechow (eds.), *Meaning, Use and Interpretation of Language*, Berlin 1982, pp. 223–248.

14 Jeroen Groenendijk/Martin Stokhof, "Dynamic Montague Grammar", in: L. Kalman/L. Polos (eds.), *Proceedings of the Second Symposium on Logic and Language*, Budapest 1990, pp. 3–48, <http://dare.uva.nl/document/3702> (accessed March 30, 2013).

15 Formally, discourse semantic theories propose that the meaning of a sentence or text corresponds to the set of variable assignments which map the set of active discourse referents to a group of people/things such that the story holds true for this group.

person she might be. Once again, the increase in information is mirrored by a reduction of choices of imaginable persons who fit the content of the story.

I concede that our subjective impression, when reading textoids like (1) or even longer stories, is *not* that we are talking about sets of women. Every reader of (1) will imagine his or her own, more or less detailed picture of a woman arriving at Heathrow. Practically every reader will fill in details which are based on her own experiences and expectations. However, none of these details are suggested by the text alone. The literal content of (1) warrants very many subjective pictures of women at Heathrow Airport. This intersubjective variability in story content is captured by the possible-worlds framework.

III Sentences and Utterance Contexts

The preceding section introduced a basic framework to represent the content of sentences, of information update and of story content. The basic framework considered sentence content in a ‘lab situation’, so to speak, disregarding factors like utterance context, speaker, utterance time etc. Utterance contexts are obviously a necessary factor in order to understand the meaning of indexical words like *I*, *you* or *now*. Moreover, as I will argue in this paper, reference to utterance contexts is an indispensable ingredient of the meaning of fictional texts. Therefore, the present section will extend the basic framework by reference to utterance situations. We will start by considering the simple case of indexical words.

The meaning of indexicals like *I*, *you*, *here* depends on the situation in which they are uttered. The meaning and truth value of the following sentence depend on whether it is uttered by the author of the present paper or by the German chancellor in summer 2013.

2 My name is Angela Merkel.

Kaplan proposed that the meaning of linguistic expressions generally depends on utterance contexts *c*.¹⁶ An utterance context specifies at least a speaker *Sp(c)*, an addressee *Ad(c)*, a time *Time(c)*, a place *Loc(c)*, and a world *World(c)*. There are various ways to implement this idea. Following Kaplan, I will assume a domain of contexts *c* which are themselves without internal structure. There are functions *Ad*, *Sp*, *Time*, *Loc*, and *World*, which are defined for all contexts *c*. Each of these functions maps each context onto some value in the appropriate domain. For

¹⁶ See Kaplan, “Demonstratives”.

example, we might have a context c_5 which is mapped by S_P to Angela Merkel, which means that c_5 is a context where Angela Merkel is the speaker. Each context defines a unique speaker, addressee, time, place, and world. We will also assume that no two contexts may be different and yet share speaker, addressee, time, place, and world. Kaplan discusses a wide range of examples which illustrate this basic framework. For example, the sentence in (2) depends in meaning on the speaker. Once the speaker has been determined, we can decide in which imaginable worlds the sentence is true.

$c_1 (S_P(c_1) = A.M.):$

'My name is Angela Merkel' is true in all worlds w where the person A.M. carries the name 'Angela Merkel'.

$c_2 (S_P(c_2) = R.E.):$

'My name is Angela Merkel' is true in all worlds w where the person R.E. carries the name 'Angela Merkel'.

Kaplan proposed that semantic interpretation proceeds in two steps: First, we resolve context-dependent expressions. Second, we evaluate the resulting clause against all imaginable worlds and determine in which ones it is true. This second step carries us back to the notion of content proposed in section II.

However, when a speaker makes a claim in an actual utterance situation, there seems to be a connection between these two steps. In a context c where I stand in front of you and assert (wrongly) 'My name is Angela Merkel,' I'd invite you to believe the following:

the world of this context c is such that
the speaker of this context c (i.e.: R.E.)
carries the name 'Angela Merkel'
at the time of this context c .

The speaker makes a claim in context c *about* the world of context c . We could paraphrase this intention as "let me tell you something about the world in which we are." Scholars in analytic philosophy have isolated this special kind of meaning as the utterance meaning, and coined the term *diagonalization* for the way in which utterances *in* contexts are *about* contexts.¹⁷ To flesh out this idea, let me give some more examples.

3 My hair needs a cut.

¹⁷ See Philippe Schlenker, "Indexicals", in: Sven O. Hansson/Vincent F. Hendricks (eds.), *Handbook of Formal Philosophy*, Dordrecht 2010, for an excellent survey of data and proposals in the literature.

Imagine that I utter this sentence in a context c where I talk to my husband in Frankfurt on August 2, 2013. What I am communicating is this: We are in an utterance context c where the world of that context $\text{WORLD}(c)$ is such that the hair of $\text{Sp}(c)$ (i.e. myself) at $\text{TIME}(c)$ is too long. Re-rendering this in the language of sets, this is my message:

- 3.a We (I, the speaker, and you, the addressee) are in one of these:
 $\{c \mid \text{in } \text{WORLD}(c), \text{ the hair of } \text{Sp}(c) \text{ at } \text{TIME}(c) \text{ is too long}\}$

Note that our epistemic uncertainty remains. When I talk to you, we might have the feeling that we know very well in which local situation c we are. Still, we do not exactly know which imaginable world w is behind the horizon of our local speech situation. More formally, we don't know the value $\text{WORLD}(c)$ of c , and in this sense don't know what c we are in. Hence, sets like (3.a) will always contain many, possibly infinitely many contexts c , even though speaker, addressee, time, and place can be fixed, as in my conversation with my husband.

Next, imagine that I continue our conversation with (4)

- 4 I was at the hairdresser's 2 months ago.

Again, I communicate a restriction on our utterance contexts. This will add (4.a) to the restrictions in (3.a).

- 4.a We are in one of these:
 $\{c \mid \text{in } \text{WORLD}(c), \text{ Sp}(c) \text{ was at hairdresser's 2 months before } \text{TIME}(c)\}$

As in the preceding section, these two restrictions together will narrow down the range of imaginable contexts c in which we might be.

- 4.b $\{c \mid \text{in } \text{WORLD}(c), \text{ the hair of } \text{Sp}(c) \text{ at } \text{TIME}(c) \text{ is too long, and}$
 $\text{in } \text{WORLD}(c), \text{ Sp}(c) \text{ was at hairdresser's 2 months before } \text{TIME}(c)\}$

This set adds to the information that the interlocutors have about their utterance context at the beginning. If we call the latter **Context₀**, we can spell out my and my husband's common information state after the little exchange as follows.¹⁸

¹⁸ I assume that we both are fully aware of our identities, the place, and the time.

Context₀ = { c | $R.E.=Sp(c)$ and $M.K.=Ad(c)$, and $TIME(c)=2.8.2013$, and $PLACE(c)=FfM$, and laws-of-physics hold in $WORLD(c)$, and FfM is in Germany, and Germany is in Europe and ...}

Update by (4.b) yields

{ c | $R.E.=Sp(c)$ and $M.K.=Ad(c)$, and $TIME(c)=2.8.2013$, and $PLACE(c)=FfM$ laws-of-physics hold in $WORLD(c)$, in $WORLD(c)$, the hair of $Sp(c)$ at $TIME(c)$ is too long, in $WORLD(c)$, $Sp(c)$ was at hairdresser's 2 months before $TIME(c)$ }

In section II, we restricted attention to context-independent utterances and therefore used sets of imaginable worlds to model meaning. The present section takes into account the fact that content depends on speaker, hearer, time and place. These are added to the world dimension and generalize the account of section II. What remains stable is the idea that the more information we get, the narrower will be the range of imaginable contexts which could be the ones described by the text.

I took the case of real, direct face-to-face assertion as my starting case. What would happen if we took the little textoid (3)–(4) as part of a fictional story? Imagine a person **A** who reads (5) as part of a novel.

5 My hair needs a cut. I was at the hairdresser's 2 months ago.

Given that we are dealing with fiction, **A** will again assume that the set of utterance contexts under consideration cannot include the context c in which **A** actually finds herself. The intended contexts c are part of worlds $WORLD(c)$ which are *not* the real world, and they therefore cannot be the real utterance context. **A** is invited to *imagine* utterance contexts c . These are characterized by (5.a), the literal content of (5), as in our earlier example.

5.a { c | in $WORLD(c)$, the hair of $Sp(c)$ at $TIME(c)$ is too long and
in $WORLD(c)$, $Sp(c)$ was at hairdresser's 2 months before $TIME(c)$ }

These contexts are imaginary in that the worlds behind the horizon are *not* the actual world. Likewise, the update does not affect the set of contexts which **A** believes she is in. As in section II, this information will update the initial set which represents **A**'s expectations about the story he will read.

Story(A)₀ = { c | the present story could be told in c about $WORLD(c)$, by
 $Sp(c)$ at $TIME(c)$... }

If we compare the set of initial possible contexts **Context**₀ in real communication with this set **Story(A)**₀, we see that the set **Context**₀ is more limited in that all its contexts *c* single out the same speaker *SP(c)*, the same hearer *AD(c)*, the same time, and the same place. This reflects the fact that we normally know where we are, and who we are talking to.¹⁹ **Story(A)**₀ contexts may yield just any individual as potential speaker. After the update of **Story(A)**₀ with the content of (5), the choice of potential speakers is already more limited.

$$\begin{aligned} &\mathbf{Story(A)}_0 \cap (5.a) \\ &= \{ c \mid \text{the present story could be told in } c \text{ about } \mathbf{WORLD}(c), \text{ by} \\ &\quad \mathbf{SP}(c) \text{ at } \mathbf{TIME}(c) \text{ and} \\ &\quad \text{in } \mathbf{WORLD}(c), \text{ the hair of } \mathbf{SP}(c) \text{ at } \mathbf{TIME}(c) \text{ is too long and} \\ &\quad \text{in } \mathbf{WORLD}(c), \mathbf{SP}(c) \text{ was at hairdresser's 2 months before} \\ &\quad \mathbf{TIME}(c) \} \end{aligned}$$

This concludes our formal account of story content. Building on earlier theories, I assume that story content is interpreted as utterance content where the reader (of the story) has to make a guess about what kind of utterance situation she finds herself in. Sequences of sentences will serve to restrict the choice of possible utterance situations. For one thing, they will delimit the range of imaginable worlds *w* which could possibly be the world of the story. But they will also delimit the choice of possible speakers *SP(c)* who could be telling the story. As before, increased information is reflected by a decreased choice of contexts (incl. context worlds) the reader might be in. As in all fiction, it is clear that these worlds, and contexts, are not the actual world/context.

We have covered some examples which illustrated how the hearer or reader of a text can gather information about context. In the following section, I give a more comprehensive survey of sources of information about the utterance context, including information within and beyond language. Section IV will select one of the linguistic sources, speaker-oriented expressions, and investigate these in more detail.

III.1 Context Information

The present section surveys different ways in which texts give us information about the speaker in contexts *c*.

¹⁹ This is obviously not always so, e.g. in telephone calls, mail, in confessions and other anonymous communication. I am not concerned with assertion in anonymous speech situations.

True and direct communication. When a real person **A** reports to another real person **B** and both know each other, the utterance content of the sentences of the story will update a set **Context**₀ = {*c* | *Sp*(*c*) = **A** and *Ad*(*c*) = **B**, and *Time*(*c*) = **t**, and *Place*(*c*) = **p** ...} which ideally codes full information about the local speech situation. The report can convey information about the speaker **A** which **B** did not know before, but the identity of speaker and hearer is fixed from the very beginning.²⁰ Note that in spite of the fact that the local speech situation can be fully known to both interlocutors, **Context**₀ will nevertheless contain many contexts *c*, at least because **A** and **B** don't know which imaginable world is the one of 'their' context, i.e. the value of *World*(*c*). I adopt the common assumption in the literature that *World* is a function on contexts. Hence, each context *c* will map to exactly one imaginable world. If the same local speech situation could be part of two different imaginable worlds *w*₁ and *w*₂, this formally requires two contexts *c*₁, *c*₂ which have the same local speech situation but map on these two different worlds.

The context situation remains essentially the same if **A** tells **B** a fictional story. The initial set of contexts to be updated will be **Story**₀ = {*c* | *Sp*(*c*) = **A** and *Ad*(*c*) = **B**, and *World*(*c*) is not the actual world, and ...}. Hence, **A** and **B** engage in a kind of *as-if* play. They act as if they were in one of the worlds where the story is true. The present analysis entails that the immersion in other imaginable worlds start right at the beginning of fiction-telling, an assumption which was proposed informally by Lewis.²¹

First-person narrated fictional text. If a reader **B** reads and interprets sentences of a fictional first-person narrative, the text offers maximal and maximally explicit information about the speaker in contexts in which the text could be uttered. The reader will start with a comparatively large set of possible utterance contexts. These generalize the initial set of worlds that was discussed in section II.

$$\text{Story}(\mathbf{B})_0 = \{ c \mid \text{the story could be told in } c \text{ about } \text{World}(c), \text{ by} \\ \text{Sp}(c) \text{ at Time}(c) \dots \}$$

Story(**B**)₀ contains contexts with different speakers, i.e. there are *c*, *c'* in **Story**(**B**)₀ such that *Sp*(*c*) ≠ *Sp*(*c'*). By and by, the text eliminates many of the initial contexts

²⁰ The present proposal adopts the view that individuals have a world-independent existence (Kripke style approach) beyond possible worlds, even though their properties can change from world to world. An alternative way to model individuals-across-worlds makes use of a counterpart relation. In this spell-out, all speakers of contexts in **Story**₀ are counterparts of **A**, and all addressees are counterparts of **B**.

²¹ See Lewis, "Truth".

because their speaker value does not match the content of the text. For example, **Story(B)₀** will normally contain both contexts with a male speaker and contexts with a female speaker. If the story reveals that the first-person narrator is female, all contexts *c* for which $Sp(c)$ is male will be discarded. Nevertheless, even stories with a first-person narrator can leave the identity of the narrator underdetermined, and in this sense, even the final set of contexts can contain contexts with different speakers. Only very special kinds of stories will result in a set of contexts which all share the same speaker. For instance, if the narrator is supposed to be a real historical person, all contexts will arguably have that person as their speaker. In contrast, if the narrator carries a name but is not a historical person, it can be argued that there might be different individuals who could carry that name, and hence different contexts which fit the story could exhibit different values $Sp(c)$.

It might be useful to remind ourselves that even if the speaker $Sp(c)$ is fully determined by the story, and hence the same for all possible contexts, this does not entail that all properties of the speaker are also fixed, or that the speaker in the story can only have properties which the actual historical person had as well. The properties of $Sp(c)$ are still determined by the worlds of context $World(c)$, and these imaginable worlds can show us the speaker with fictitious properties, even if the fictitious speaker is supposed to be a historical person. (Among other things, the individual will have the property of telling a story which, in actual fact, the historical person never told.)

Speaker-oriented expressions. Even texts which don't use the first-person pronoun *I* can give information about the speaker. Emotive interjections (*thank heavens!*, *damn!*, *oh my god!*), evaluative adverbs (*luckily*, *sadly*), epistemic adverbs, modals, and particles (*perhaps*, *maybe*, *certainly*, *'wohl'*), emotive expressions (*at last!*, *X was SO smart/SUCH an idiot*, *What a dancer he was!*) as well as questions (*Was he mad?*) literally convey information about the person who counts as the speaker of a sentence. In literary texts, this can either be the speaker-in-context (and usually, such words quickly create the impression of a narrator) or it can be a protagonist, if the expression is used in direct speech or in (free) indirect discourse. Both cases will be discussed in more detail in section IV. The semantic analysis of these terms relates to the speaker, just as the semantic analysis of the pronouns *I* and *you* relates to the speaker and addressee in the utterance context. While we will sort out a few differences between these two kinds of indexical reference, all such expressions and words convey information about the speaker. For instance, consider a fictional story which starts with the following sentence.

6 How lovely Panama is!

A reader **A** of (6.) will initially maintain the set of contexts $\text{Story}(\mathbf{A})_0 = \{c \mid \text{the story could be told in } c \text{ about } \text{WORLD}(c), \text{ by } \text{SP}(c) \text{ at } \text{TIME}(c) \dots\}$ as usual. After interpreting (6.) and updating $\text{Story}(\mathbf{A})_0$ with its content, **A** will have reduced the set of all possible utterance contexts c to such ones where $\text{SP}(c)$ is a great fan of Panama in $\text{WORLD}(c)$. Similar considerations hold for other expressions, and the next section takes a closer look at them and the kind of information they contribute to the meaning of an utterance or text.

There is an intuitive difference between first-person narratives/texts and third-person narratives/texts where the speaker offers only indirect evidence about his emotions and attitudes. In the first case, the text establishes a *discourse referent* for the speaker, turns him officially into a member of staff, so to speak. In the latter case, the speaker is not a member of staff, but the set of contexts c which are such that the story could have been uttered there will reflect restrictions on the speaker value. This difference can be reflected in discourse representation theories, which were briefly mentioned at the end of section II.

Speaker-neutral fictional texts. Finally, there are fictional texts which use no words or expressions which refer to $\text{SP}(c)$ in any way. In this case, the set of contexts which reflect the content of the story can cover any value for $\text{SP}(c)$. More formally, neutral stories give rise to sets of contexts with the following property:

A story is speaker-neutral if for any context c in $\text{Story}(\mathbf{A})_n$ with speaker $\text{SP}(c) = \mathbf{x}$, and arbitrary other person $\mathbf{y}$, there is a context c^* in $\text{Story}(\mathbf{A})_n$ with $\text{SP}(c^*) = \mathbf{y}$, but $\text{WORLD}(c^*) = \text{WORLD}(c)$, $\text{TIME}(c^*) = \text{TIME}(c)$, $\text{AD}(c^*) = \text{AD}(c)$, $\text{PLACE}(c^*) = \text{PLACE}(c)$.²²

This property reflects the intuition that the person who tells us the story could be just any person at all. The text imposes no restrictions whatsoever on the possible options for a speaker. The final set of contexts $\text{Story}(\mathbf{A})_n$ is 'big' in the sense that for each of the worlds w which adhere to the story content, there are as many contexts which could be those where the story was told as there are people in the world – and each of these persons could have been the one who told us the story. Remember that, like at earlier points, many options mean little information. $\text{Story}(\mathbf{A})_n$ covers particularly many speakers simply because there is so little information in the text as to what kind of person the speaker might be.

²² Literally the condition is too liberal. It must be ensured that $\mathbf{y}$, like $\mathbf{x}$, lives at a time which coheres with the temporal information conveyed by in the story; i.e. $\mathbf{y}$ lives after the events of the story if the story is written in past tense, etc. For the sake of clarity, I leave out these requirements.

IV Information about the Speaker

This section takes a closer look at words and constructions which convey information about the speaker, and do so without making use of the first-person pronoun *I/ich*. We will first revisit a range of examples both in English and German. I will then argue that the information conveyed is part of the literal content of sentences. Finally, I sketch how they can be integrated into the analysis of story content in section III.

Consider the following first two sentences of Thomas Mann's novel *Joseph und seine Brüder* (*Joseph and his brothers*).

- 7 *Tief ist der Brunnen der Vergangenheit. Sollte man ihn nicht unergründlich nennen?*
(lit: Deep is the well of past. Should not one call it bottomless?)

This short passage conveys information about the speaker-in-context: Whoever it may be, the person establishes a bond with the addressee beyond mere information transfer. He asks a rhetorical question. A question indicates that there is someone who lacks knowledge and asks to be informed – even if only rhetorically. As a result, the reader immediately is invited to imagine a narrator. The novel consistently uses the pronoun *we* in reflective passages, which conveys the fiction that reader and narrator are engaged in a four-volume long reflection on the biblical story.²³ The English translation renders the question in a form that makes this reader-narrator company explicit from the very beginning:

- 8 Deep is the well of the past. Should not we call it bottomless?

With or without the use of a first-person pronoun, questions in texts convey the information that the speaker, whoever it may be, is asking for information.

Let us look at more indications for speakers' attitudes. The following sentences are the beginning of Selma Lagerlöf's tale *The Adventures of Niels* (*Niels Holgersson's underlige resa genom Sverige*) in the German and English version.

- 9 Es war einmal ein Junge. Er war ungefähr vierzehn Jahre alt, groß und gut gewachsen und flachshaarig.

²³ It remains to be discussed whether the reader, more concisely, imagines that *Thomas Mann* wants to engage in a four-volume discussion with her. Even this *Thomas Mann* will be hypothetical, though, a *Thomas Mann* in counterfactual worlds where true two-way exchange between reader and narrator is possible.

- 10 Once there was a boy. He was – let us say – something like fourteen years old; long and loose-jointed and towheaded.

Both passages express uncertainty about the age of the boy. The German passage uses *ungefähr* ('roughly'), which leaves it open whether the speaker has to guess the age or considers the exact age to be irrelevant for the reader. The English passage is more explicit in its reference to the speaker, as it uses the phrase *let us say*, which contains the pronoun *us* (even though it can be debated whether the pronoun actually refers to the speaker in a phrase like this). The entire novel conveys the sense that some narrating individual is responsible for the selection of information, comments, and evaluatives. Interestingly, the German version of the novel never uses second-person pronouns to locate the speaker, whereas the English translation occasionally refers to "*us*" or "*we*". The next passages offer more examples. In (11), both English and German implicitly refer to the speaker's epistemic background. In (12), the German version makes pronoun-free reference to the speaker whereas the same passage in English makes reference to the narrator explicit.

- 11 (*The boy watches his parents leave for church and believes that they congratulate themselves for having ordered him to read the Bible.*)
 German version: *Aber der Vater und die Mutter wünschten sich sicherlich nicht Glück, sondern sie waren ganz betrübt.*
 English version: *But his father and mother were certainly not congratulating themselves upon anything of the sort; but, on the contrary, they were very much distressed.*
- 12 (*Wild geese are flying over Niels Holgersson's farm.*)
 English version: *It was, as we have said, an uncommonly fine day, with an atmosphere that it must have been a real delight to fly in, so light and bracing.*
 German version: *Es war, wie gesagt, ein überaus schöner Tag, und die Luft war so frisch und leicht, dass es ein Vergnügen sein musste, darin zu fliegen.*

The German parenthetical in (12) does not use the pronoun *we* but nevertheless conveys a speaker commentary. The speaker refers back to an earlier statement and takes it up again to continue the story. The English phrase "*as we have said*" does the same, yet more explicitly.

A comparison with the Swedish original shows that Lagerlöf's own reference to the speaker was as implicit as in the German translation, at least in the passages in (10) and (12).

- 13 a. *Det var en gång en pojke. Han var så där en fjorton år gammal, lång och ranglig och linhårig, (...).*
 lit: "... he was so there some fourteen years old"

b. *Det var, som sagt, en ofantligt vacker dag med en luft, som det måtte ha varit en sann glädje att flyga i, så frisk och så lätt.*

lit: "It was, as said, an uncommonly fine day ..."

These examples suggest the following:

- Languages provide means to report comments by the speaker, with or without the use of first-person pronouns.
- Different languages can omit pronominal reference to different degrees. While German (and possibly, Swedish) often leaves the speaker implicit, English tends to use more pronouns.²⁴

But even English has speaker-oriented expressions which implicitly refer to the speaker. One type of writing where such expressions occur with high frequency is free indirect discourse. In such passages, speaker-oriented expressions report feelings and comments of the speaking/thinking protagonist. The following examples illustrate this type of speaker reference. All are taken from Katherine Mansfield's short story "Miss Brill." In each case, the speaker-oriented expression, used in free indirect speech, reflects Miss Brill's thoughts.

14 *She glanced, sideways, at the old couple. Perhaps they would go soon.*

→ epistemic modals

15 *Oh, how sweet it was to see them snap at her again from the red eiderdown!*

→ exclamatives

16 *[The nose] must have had a knock, somehow.*

→ epistemic modality

17 *Now there came a little "flutey" bit – very pretty! – a little chain of bright drops.*

→ interjective evaluatives

18 *The old people sat on the bench, still as statues. Never mind, there was always the crowd to watch.*

→ comparison, evaluation of contrasting assertions

19 *...and such a funny old man with long whiskers hobbled along in time to the music and was nearly knocked over by four girls walking abreast.*

→ evaluative *such*-constructions

²⁴ This is not intended as a characterization of English in general. What we can assess, at this point, are preferred choices of professional translators of literature as to the stylistically optimal way to render the content of a passage in English.

- 20 *No doubt somebody would have noted if she hadn't been there (...) No wonder! (...)*
→ commentary by local speaking protagonist: Miss Brill.

In Mansfield's short story, none of these expressions offer us information about the narrator of the story. They are all couched in free indirect discourse and are oriented to Miss Brill, the main character. This use is in fact quite typical for speaker-oriented expressions in literary texts. They are frequently used in indirect discourse in order to indicate that the passage has to be attributed to *some* speaker. Usually, the content of the expression suggests that this speaker is not supposed to be the narrator but, more plausibly, a protagonist. In free indirect discourse, speaker-oriented expressions convey information about *some* speaker, even though this is not the narrator. Yet, this confirms their function to convey information about the speaker of a sentence.

While we cannot go through a comprehensive survey of all speaker-oriented words of English, this list of examples should be sufficient to illustrate the phenomenon. Of course, the novel also uses meta-linguistic means to indicate passages reflecting Miss Brill's thoughts: lack of words (*a something, what was it?*), repetitions (*or even – even cupboards!*), self-corrections (*no, not sadness*), sound imitations (*tum-tum-tum tiddle-um!*) as well as Miss Brill's quoting other persons around her. I will leave such speaker indicators aside. While their *use* indeed can tell us something about the speaker (and in particular, indicate that the speaker is most likely *not* the narrator at that point), they do so at a meta-level and not by their literal meaning. For instance, there is nothing in the meaning of *a something* which indicates that the speaker is lacking words.

Could not the expressions in (14)–(20) be meta-linguistic signals? If they were, they might contribute to our knowledge about the utterance situation like other meta-linguistic or non-linguistic signals do. For instance, a smile or a frown can tell us a lot about the speaker's attitude, and thereby narrow down the range of possible utterance situations without counting as part of the utterance itself. However, I propose that the expressions in (14)–(20) are part of the utterance and add information to the utterance content:

- Many of them belong to grammatical categories which are well-integrated into clause structure: adverbials, modal verbs, wh-constructions, and others.
- They all take part in semantic composition insofar as they take arguments which are contributed by other parts of the sentence.
- Finally, they can be bound by quantifiers when they occur in examples like the following:

- 21 Every one of the children called and reported ...
that, perhaps, they would come home soon
how sweet it was to be here

that they must have gotten lost, somehow
 that there was such a funny old man
 that no doubt they would get home soon.

The examples show that the speaker parameter can be quantified over in the matrix clause. In (21.a), we understand that each child individually expresses epistemic uncertainty (*perhaps*), (21.b) attributes individual surprise at the sweetness, (21.c) individual epistemic necessity, etc. Quantificational dependencies offer evidence that the respective parameter must be accessible in the logical form of the sentence. We can conclude that the listed expressions refer to the speaker at a linguistic level.

As a final type of speaker-oriented items, let me add German particles (*‘Abtönungspartikeln’*). This category is widely used in German and more or less absent in Modern English. German particles code intricate beliefs on the part of the speaker about her own knowledge, the addressee’s knowledge, the current train of thought, and rhetorical intentions. The following examples illustrate this.

- 22 *Peter hat ja ein Auto. Er kann die Ski transportieren.*
 ‘Peter owns *ja* a car. He can transport the skis.’
 Contribution of *ja*:²⁵
- i. *epistemic*: the speaker believes that the addressee might know that Peter owns a car.
 - ii. *rhetorical intention*: the speaker draws attention to the (possibly known) fact in order to support an adjacent claim or proposal.
- 23 *Wieso will Peter die Ski nicht bringen? Er hat doch ein Auto.*
 ‘Why doesn’t Peter want to bring the skis? He owns *doch* a car.’
 Contribution of *doch*:²⁶

²⁵ The paraphrase summarizes various proposals in the literature, e.g. Angelika Kratzer, “Beyond ‘oups’ and ‘ouch’”. How Descriptive and Expressive Meaning Interact”, Talk, Presented at the Cornell Conference on Theories of Context Dependence, 1999. Available at *Semantics Archive*; Malte Zimmermann, “Zum Wohl: Diskurspartikeln als Satztypmodifikatoren”, in: *Linguistische Berichte* 199/2004, pp. 1–35; Malte Zimmermann, “Discourse Particles”, in: Klaus v. Heusinger/Claudia Maienborn/Paul Portner (eds.), *Semantics. An International Handbook*, Vol. 2, Berlin 2012, pp. 2012–2024; Regine Eckardt, “Particles as Speaker Indexicals in Free Indirect Discourse”, in: Lotte Hogeweg/Eric McCready (eds.), *International Journal for Language Data Processing: Formal Approaches to Discourse Particles and Modal Adverbs*, 1/2012, pp. 99–119; and more. For German particles in general, see Renate Pasch, *Handbuch der deutschen Konnektoren*, Berlin 2003.

²⁶ Drawing on Markus Egg, “Discourse Particles at the Semantics-Pragmatics Interface”, in: Werner Abraham/Elisabeth Leiss (eds.), *Modality and Theory of Mind Elements across Languages*, Berlin 2012, pp. 297–333; Zimmermann, “Particles”; Patrick Grosz, “German ‘doch’: An Element that Triggers a Contrast Presupposition”, in: *Proceedings of the Chicago Linguistic Society*, 46/

Speaker acknowledges that the host sentence 'Peter owns a car' conflicts with another salient proposition in context, here plausibly 'Peter doesn't want to bring the skis'.

In the final part of this section, I will briefly sketch semantic analyses for a few speaker-oriented items, and spell out in what respect such speaker-oriented items differ from personal pronouns (*I/ich, we/wir*), aiming at an integrated theory of reference to the speaker. As in earlier sections, I will assume a domain of contexts D_c and functions $SP, AD, NOW, HERE, WORLD$ that map each context c to its speaker, its addressee, its time, its place, and the world in which it occurs.

We will first consider speaker-oriented items which comment on the whole sentence S . (24) illustrates this with the German adverb *leider* and English *regrettably*.

- 24 *Leider* hat Peter kein Auto.
 Regrettably, Peter does not have a car.

The message of these sentences is two-fold. They are used to assert the content of S ("Peter does not have a car") and they report the speaker's momentary attitude about this fact.

- 24.a *Regrettably* + "Peter does not have a car"
 asserts:
 $WORLD(c)$ is such that Peter does not have a car at $NOW(c)$
 In $WORLD(c)$, $SP(c)$ regrets at $NOW(c)$ the content of "Peter does not have a car".

The same analysis applies to German *leider*. The contribution of (24.) to a story in German or English will hence consist in the following two restrictions on the sets of contexts in which that story might be told:

- { c | $WORLD(c)$ is such that Peter does not have a car at $NOW(c)$ }
 { c | In $WORLD(c)$, $SP(c)$ regrets at $NOW(c)$ the content of "Peter does not have a car". }

Epistemic modal expressions relate the content of the sentence S to the speaker's beliefs.

2010; Elena Karagjosova, "Adverbial 'doch' and the Notion of Contrast", in: B. Behrens/C. Fabricius-Hansen (eds.), *Oslo Studies of Language 1: Structuring information in discourse: The explicit/implicit dimension*, Oslo 2009, pp. 131–148; and others.

- 25 Perhaps, Peter does not have a car.
 Peter might not have a car.
 Vielleicht hat Peter kein Auto.

These sentences convey the following content:²⁷

- 25.a In $\text{WORLD}(c)$, the worlds which $\text{SP}(c)$ believes could be the real world include some w' where Peter does not own a car in w' .

This is the restriction on possible contexts that emerges:

- $\{c \mid \text{WORLD}(c) \text{ is such that the worlds which } \text{SP}(c) \text{ believes could be the real world include some } w' \text{ where Peter does not own a car in } w'\}$

Obviously, this information may trigger further inferences by the reader. If the speaker is reliable and trustworthy, the reader will infer that the context itself might be such that Peter does not own a car in $\text{WORLD}(c)$.

The contribution of *ja* illustrates that commentaries can also refer to the addressee, or at least the speaker's beliefs about the addressee.

- 26 *Peter hat ja ein Auto.*
 Peter has *ja* a car.

This is the content of (26.):

- 26.a $\text{WORLD}(c)$ is such that Peter owns a car at $\text{NOW}(c)$
 $\text{WORLD}(c)$ is such that $\text{SP}(c)$ believes that $\text{AD}(c)$ could know '*that Peter owns a car*'
 $\text{SP}(c)$ plans to use the fact '*that Peter owns a car*' to support another claim.

The particle *ja* expresses that the speaker is highly aware of the details of the on-going conversation, and is planning the course of argumentation. A wide range of particles in German convey the speaker's awareness of the utterance situation.

Emphatic intensifiers like *be so smart* or *be such an idiot* have been analysed as expressing the speaker's surprise or emotional involvement about qualities.²⁸ Let us take a look at an example:

²⁷ I am adopting a standard treatment of modality in terms of possible worlds. For an accessible introduction see Paul Portner, *Modality*, Oxford 2009.

²⁸ For a detailed analysis, the reader is referred to Jessica Rett, "A Degree Account of Exclamatives", in: *Proceedings of SALT*, 18/2008, pp. 601–618, and references therein.

27 Peter was SO smart.

The use of *SO* conveys the following content (which compositionally builds on the meanings of “Peter”, “was” and “smart”):

27.a *WORLD(c)* is such that Peter exhibits smartness to a high degree *d*
SP(c) believes that *d* is a high degree of smartness, and is emotionally moved
 in *WORLD(c)* at time *NOW(c)* about *d*.

Again, we see how sentence (27) contributes facts about the world in which the story is told (Peter is very smart) and at the same time facts about the interlocutors, specifically the speaker and her surprise.

Speaker commentaries in English have not so far been extensively discussed in the literature. Let me therefore speculate about the contribution of an expression like *no wonder* to the utterance “No wonder, *S*”. The message is more elaborate than a plain “*S* holds true, which does not surprise me”. This can be tested when we try to use ‘*no wonder*’ in various contexts.

Imagine that you are visiting England. You wake up in the morning and open the curtain. It is raining. You have strong beliefs about the English climate, and have been told that it rains most of the time. So the weather does not surprise you. Still, you could not in this situation utter (28).

28 #No wonder it is raining.

is used to mark that the sentence is grammatical but would be inappropriate in the described situation.

If we study the kind of situation in which *no wonder* can be used appropriately, the following picture emerges. The expression indicates a complex inferential process in the speaker’s mind. The speaker must have known *S* before (even if she may not have asserted it so far). The speaker has just learned *another* fact. And this new, other fact makes *S* much more plausible, less surprising than it was without this other piece of information. An appropriate situation to use (28) could look as follows: You have strong beliefs about the English climate and have been told that it rains all the time. You are on a trip around the world under the care of your well-organized partner. You have just landed in a country the name of which you don’t know. The weather, obviously, is awful. The pilot of the aircraft announces: “*We have now safely landed at Heathrow airport. Welcome to England.*” At this point, you put two and two together and utter (29).

29 No wonder it is raining.

The paraphrase in (29.a) spells out the content of “*no wonder*” in terms of speaker belief.

- 29.a *No wonder S*
 asserts:
S is true in the utterance context *c*.
 $\text{Sp}(c)$ knew that *S*.
 $\text{Sp}(c)$ has just learned or found out something *T* which makes *S* much more likely.²⁹

As (29.a) reveals, the English *no wonder* conveys very private information about the speaker and his increased knowledge. If a fictional text were to use the expression, the reader would immediately understand that the evoked utterance situation is one where the fictitious speaker gains new insights while speaking. This restriction on possible speakers in *c* could be sufficient to create the fiction of a narrator even without the use of pronouns *I* or *we*.

Pronouns like *I* and *we* have long been discussed as indexical expressions which refer to the speaker. In the present section, we have seen speaker-oriented expressions of a different kind. Even though they clearly refer to the speaker of the utterance (and sometimes even the addressee), they are not traditionally listed in the literature on indexicality. Let me end this section by comparing classical indexicals (in languages like English and German) to speaker-oriented items. One major property of indexicals in English and German is that they never shift reference. Wherever they are used, they always refer to the “true external” speaker.

- 30 *I* love chocolate.
 Peter said that *I* liked chocolate.
 Tom believes that *I* like chocolate.
 Sue claimed that Tom believes that *I* like chocolate.
 Sue hesitated. *I* seemed to like chocolate, she thought.

The pronoun *I* refers unambiguously to the person who utters the sentence. The same holds true for *you*, for the use of the third-person feature (and for utterance time, which, however, requires a more extensive discussion). The sole exception is posed by quoted speech – which, naturally, depicts the exact words of another person and hence constitutes a different case.

Matters are very different when we consider the use of a speaker-oriented commentary like *regrettably* in embedded speech contexts.

²⁹ Readers who love probability theory could spell this out as $P(S|T) > P(S)$: the conditional likelihood of *S*, given that *T*, is higher than the likelihood that *S*.

- 31 a. Peter said that, regrettably, I liked chocolate.
 b. Sue claimed that Tom, regrettably, believed that I like chocolate.
 c. Sue hesitated. Regrettably, I seemed to like chocolate, she thought.

What we observe is that the experiencer of regret can be Peter in (31.a), can be Sue in (31.b) and must be Sue in (31.c). Unlike in the case of true indexicals, the orientation of commentaries can shift away from the person who “really” utters the sentence.

If we want to devise a full analysis of how sentence content refers to the context of utterance, we have to implement context dependence in two different modes. One mode will have to follow Kaplan and others who restrict attention to indexicals which determine reference, based on utterance contexts.³⁰ This mode must always rigidly refer to the external utterance context. The other mode has to allow for more flexibility, including the option to shift reference to other speakers, and including the option to quantify over these speakers. The effect is once again illustrated in (32), whereas first person pronouns in (33) do not show this dependency.

- 32 *Every guest called and said that regrettably, his room was the worst.*
 = different regrets, different rooms for different guests
 33 *Every guest called and said that my room was the worst.*
 = only one room: the external speaker's

A full analysis of two modes of context sensitivity will require spelling out the exact treatment of context parameters at the interface between syntax and meaning. We need to ensure that eventually, all context-referring expressions are attributed to the correct kind of context, and offer information about the right kind of speaker. While such a treatment is feasible, its details do not contribute to the topic of the present paper.³¹

³⁰ See Kaplan, “Demonstratives”.

³¹ Examples of possible mappings from syntax to semantics/pragmatics have been spelled out in Philippe Schlenker, “A Plea for Monsters”, in: *Linguistics & Philosophy*, 26/2003, pp. 29–120; Philippe Schlenker, “Context of Thought and Context of Utterance. A Note on Free Indirect Discourse and the Historical Present”, in: *Mind and Language*, 19/2004, pp. 279–304; Yael Sharvit, “The puzzle of free indirect discourse”, in: *Linguistics and Philosophy*, 31/2008, pp. 353–395. Eckardt, “Particles”; the empirical challenges posed by shiftable indexicals are discussed in Philippe Schlenker, *Propositional Attitudes and Indexicality (A Cross-Categorical Approach)* [Diss. MIT], 1999; Schlenker, “Indexicals”; as well as in the literature on sign language where indexicals show patterns of use which differ substantially from those common in spoken languages.

V Story Contents

This section aims to illustrate the proposed analysis on the basis of a small sample of texts. We will look at texts without reference to the speaker, texts which use first-person pronouns, and texts which use speaker-oriented expressions. I will moreover illustrate how the same text differs in content, depending on whether it is part of a real face-to-face communication or read as fiction. Finally, we return to the question how the present analysis can reconcile the universal use of contexts with the observation that sometimes, but not always, a text creates the fiction of a narrator.

Our starting point will be another brief passage from the book *The Wonderful Adventures of Nils* (*Nils Holgersson's underbara resa genom Sverige*) by Selma Lagerlöf. The novel is written in a manner which creates the strong impression of being told by someone, without explicit reference to a narrator by first-person pronouns (at least in the Swedish original, as well as the official German translation). Going through the text, it becomes clear that Lagerlöf achieves this effect by the occasional use of speaker-oriented expressions. The basic passage in (34)/(35) was chosen because it contains two such expressions, both epistemic modals in English; the German version contains one modal particle and a modal adverb. The English passage provides more context.

34 [Nils Holgersson] waded forward between some white anemone-stems – which were so high they reached to his chin – when he felt that someone caught hold of him from behind, and tried to lift him up. He turned round and saw that a crow had grabbed him by the shirt-band. [a] He tried to break loose, but before this was possible, another crow ran up, gripped him by the stocking, and knocked him over. [b] If Nils Holgersson had immediately cried for help, the white goosey-gander certainly would have been able to save him; but the boy probably thought that he could protect himself, unaided, against a couple of crows. (*The Wonderful Adventures of Nils*, Episode April 13th. Selma Lagerlöf)

35 [a] [Nils Holgersson] versuchte, sich loszureißen; aber ehe ihm dies gelang, eilte noch eine Krähe herbei, biss sich in einem von seinen Strümpfen fest und riss ihn zu Boden. [b] Wenn der Junge sogleich um Hilfe geschrien hätte, wäre es dem Gänserich wohl gelungen, ihn zu befreien. Aber der Junge dachte wahrscheinlich, mit ein paar Krähen müsse er es allein aufnehmen können.

First, I render the restrictions on context which are conveyed by [a], a passage which does not contain speaker oriented expressions, and [b], a passage with two epistemic modals (English).

34.a {c | SP(c) lives at TIME(c), which is later than the events at hand;
in WORLD(c), N.H. tries to break loose;

in $WORLD(c)$, this is interrupted by a crow – different from any
aforementioned crow – who grips **N.H.** by the stocking
and knocks **N.H.** over}

- 34.b { c | $SP(c)$ lives at $TIME(c)$, which is later than the events at hand;
 $SP(c)$ in $WORLD(c)$ is certain that ‘if **N.H.** had called for help, **G.G.**
would have rescued him’;³²
 $SP(c)$ in $WORLD(c)$ holds it possible that **N.H.** thought in $WORLD(c)$
that he could fight a crow alone}

The literal content of passage [a] contains only minimal restrictions on possible speakers who could have told it. As the past tense indicates, the speaker must tell (and hence, live) after the time when Nils’ adventure takes place.³³ The content of passage [b] places more limits on the choice of possible speakers. In any world w which fits the story content so far, the range of possible speakers is delimited to those who maintain certain beliefs about Niels and hypotheses about counterfactual courses of events.

Let us next see how the content in [a] and [b] integrates with the reader’s or hearer’s knowledge in different situations of listening or reading the story. First, imagine that the story is being told by Selma Lagerlöf to her pupil Mats. This is what Mats knows about possible contexts at the outset.

Story₀(Mats) = { c | $SP(c)$ =**Selma.Lagerlöf** and $AD(c)$ =**Mats/myself** and
 $WORLD(c)$ is such that Sweden’s history and geography
are like in the actual world;
but $WORLD(c) \neq$ actual world }

At the point of the story in (34), Mats updates the current story set, a subset of **Story₀(Mats)**, by [a] and [b]. The set of contexts in [b] will contribute that Selma Lagerlöf has certain beliefs about the motives of Niels Holgersson and about counterfactual alternative courses of events. Yet, the choice of speaker is not further delimited at this point, because it has been clear all along that the speaker is Selma

³² Technically, this conditional requires another quantification over possible worlds: ‘In all worlds close to the speaker’s world where **N.H.** calls for help quickly, **G.G.** rescues **N.H.** with success’. See David Lewis, *Counterfactuals*, Oxford 1973; Kratzer, “Descriptive and Expressive Meaning”; Portner, *Modality*, for standard accounts of counterfactuals.

³³ This is consistent with Käte Hamburger’s observation that in fiction, the *past tense* morphology has lost its function to refer to the past: *past tense* is not related to any point in time which has to do with the *production* or *reception* of the text. *Past* means neither “before the author wrote this” nor “before you are reading this.” It is reduced to: “in whatever c this text could have been uttered, c must be after the events which are reported” – if *past tense* is used, as in our example.

Lagerlöf. Interestingly, both Mats and Selma Lagerlöf must be aware of the fact that they are in fact play-acting. Given that Selma Lagerlöf will be the inventor of the protagonist Niels, as well as everything that happens to him, she cannot really be uncertain about his motives. Both Mats and Selma act as if the story were real, and only under this assumption does it make sense for Selma to express the beliefs at hand.

Let us next consider the case where Mats is reading the story in a book. (Actually, the story was originally written as a geography textbook for Swedish school children.) Mats is aware of the fact that the book must have been written by someone. Yet, he also knows that schoolbook authors do not literally tell you stories. On the other hand, he may have been told that this was a book about Sweden's landscapes and history. Hence, this is his starting point.

$$\text{Story}_0(\text{Mats}) = \{ c \mid \text{WORLD}(c) \text{ is such that Sweden's history and geography} \\ \text{are like in the actual world;} \\ \text{but } \text{WORLD}(c) \neq \text{actual world} \}$$

At the point of the story presented in (34), Mats will again intersect his current story content with [a] and [b]. Both sets will further restrict the set of possible worlds $\text{WORLD}(c)$ about which the story is told. For instance, [a] restricts choices to those worlds where Niels is caught by crows – leaving aside worlds where he cries for help fast and is rescued, among others. Passage [b], however, imposes more restrictions on the speaker. Choices are delimited to those c where the speaker maintains two beliefs: First, that Niels *very likely* could have been rescued under certain circumstances. Second, that Niels *perhaps* acted wrongly out of pride. If Mats trusts the speaker, he can delimit the choice of possible contexts c further to those where Niels indeed could have been rescued, and where Niels indeed acted wrongly out of pride. But this will be Mats' inference from what has been presented as the beliefs of the speaker.

Finally, let us consider an alternative version of passage [b]. Information about the speaker could also be rendered by making use of the first-person pronoun. The passage might look like this:

- 36 If Nils Holgersson had immediately cried for help, I am certain that the white goosey-gander would have been able to save him; but I guess that the boy thought that he could protect himself, unaided, against a couple of crows.

The set of utterance contexts which is covered by (36) is more or less identical to the set (34.b). What is different, though, is the status of the narrator as a "member of staff." Dynamic theories of meaning capture members of staff as discourse referents. Discourse referents reflect which persons and things were explicitly

introduced in the story, and regulate anaphoric cross-referencing in texts. As soon as the text refers to $Sp(c)$ with the pronoun *I*, a discourse referent for the speaker will be generated and carried on in the (dynamic) representation of the text. In a dynamic version of the present account, we could hence track the differences in status of the speaker in (34) and (36) even though the content conveyed about the speaker remains the same in both versions.

Careful readers might point out that we have broadened the notion of ‘story content’. Asked about the content of *Niels Holgersson*, we would start listing his adventures and maybe also facts about Sweden’s geography and history – but we would not call ‘that someone is telling us this who is not sure about Niels’ motives sometimes’ part of the content, in the traditional sense of the word. In the present account, the content of fictional stories is taken in a broad sense: A story conveys its plot – events, actions, adventures, happenings – and, in addition, a more or less vivid story about the situation in which someone tells it to me. Sometimes it may be useful to distinguish between *broad* and *narrow* story content. I hope that no confusion arises in the present paper.

VI Speakers and Narrators

Let us return to the question where narrators come from. The present account of story content is based on the assumption that stories generate the set of possible contexts (including worlds, speakers, addressee, time, and place) in and about which they can truthfully be told.

Section II introduced the world-based version of the analysis that goes back to Stalnaker.³⁴ It assumes that sentence meanings are captured by the sets of imaginable worlds in which the sentence would be true. The content of a story is represented by the set of imaginable worlds where all sentences of the story (in the respective order) are true. The story hence restricts the set of all imaginable worlds to those which could be the world in which the story took place. The less specific the story, the more choices of worlds are left open. Section III extended the account to include context information. According to the new account, story contents are sets of possible contexts, namely those contexts c which are such that $Sp(c)$ could be the person who told the story, $Ad(c)$ could be the person who listens or reads, $Time(c)$ could be the time of the telling and $Place(c)$ could be the location where the story is being told. Moreover, c is linked to $World(c)$, the world about which the story is told – and hence, a world where the story is true.

34 Stalnaker, “Assertion”; Stalnaker, *Context*; Stalnaker, “Common Ground”.

The simpler account can be mapped onto this richer account, if we make the additional assumption that in a Stalnaker type theory, speaker A, addressee B, time T, and place P are known to both interlocutors. If this is the case, then there is a one-to-one correspondence between *possible story worlds* w in the simpler account, and possible story contexts in the richer account: Each world w in the simpler version of story content will correspond to the one context c where $SP(c) = A$, $AD(c) = B$, $TIME(c) = T$, $PLACE(c) = P$ (i.e. the local situation looks as both A and B know it), and $WORLD(c) = w$.

The richer account covers information about the speaker, including but not limited to first-person narrations. In section IV, we surveyed a range of known linguistic expressions that implicitly refer to the speaker. We restricted attention to *linguistic* reference to the speaker. What was systematically excluded from consideration are restrictions which draw on opinions expressed, positions defended, moral or religious content etc. While these add many more options to restrict the choice of possible speakers, it requires more than linguistic analysis to argue that they have this effect.

The account supposes that story content *always* covers a huge range of imaginable contexts, which all include a speaker. Doesn't this predict that stories always have a narrator? And is not this prediction falsified by the observation that many stories do not create the fiction of a narrator at all?

In fact, the account does not predict the *fiction* of a narrator. *Many options* always mean *little knowledge*, which can come close to *no knowledge at all*. If the content of a story could have been told by anybody, the illusion of a protagonist "somebody" simply does not arise.

The difference between *little* and *more* information can be illustrated on the basis of the passages by Selma Lagerlöf in section V. Consider once again the sets [a] and [b] which reflected speaker-neutral and speaker-referring content, respectively.

- [a] = { c | $SP(c)$ lives at $TIME(c)$, which is later than the events at hand;
in $WORLD(c)$, the person N.H. tries to break loose;
in $WORLD(c)$, this is interrupted by a crow – different from any
aforementioned crow – who grips N.H. by the stocking
and knocks N.H. over}
- [b] = { c | $SP(c)$ lives at $TIME(c)$, which is later than the events at hand;
 $SP(c)$ in $WORLD(c)$ is certain that 'if N.H. had called for help, G.G.
would have rescued him';
 $SP(c)$ in $WORLD(c)$ holds it possible that N.H. thought in $WORLD(c)$
that he could fight a crow alone}

These sets have distinct properties, and [a] is ‘richer’ than [b] in the following sense:

For any c_1 in [a], and any person **Z** who lives in $\text{WORLD}(c_1)$ around $\text{TIME}(c_1)$, there is a context c_2 in [a] such that $\text{SP}(c_2) = \mathbf{Z}$ and c_2 is like c_1 otherwise.

For a context c_3 in [b], however, we can find persons **Y** in $\text{WORLD}(c_3)$ who live at the right time but could not have been the speaker: i.e. there is no c_4 in [b] such that $\text{SP}(c_4) = \mathbf{Y}$ and c_4 is like c_3 otherwise. For instance, if a person **Y** does not believe in $\text{WORLD}(c_3)$ that Niels could have acted wrongly out of pride, **Y** is not an option as a speaker. In this sense, the set of contexts in [b] is more restricted than [a].

What we observe – and this is an empirical point, not a theoretical prediction – is that too few restrictions may be insufficient to create the fiction of a narrator, in spite of the fact that the formal speaker parameter is in play. The situation is, perhaps, comparable to *point of view* in movies. Each shot defines a position for the viewer and a direction in which he or she would look. In this sense, a *point of view* is always present throughout the movie. However, only very specific choices of subsequent points of view are suited to create the illusion that we are seeing the world through the eyes of one specific person, that there is a ‘seer’ in analogy to the narrator in texts.³⁵ In movies, putting your camera somewhere does not always create the fiction that someone was watching from this position. In narrations, linking the story to a speaker does not always create the fiction that someone is telling me this.

Does a story create the fiction of a narrator? This remains the volitional choice of the person who words the story: the author.

Works Cited

- Davidson, Donald, “Truth and Meaning”, in: *Synthese*, 17/1967, pp. 304–323.
- Eckardt, Regine, “Particles as Speaker Indexicals in Free Indirect Discourse”, in: Lotte Hogeweg/ Eric McCready (eds.), *International Journal for Language Data Processing: Formal approaches to discourse particles and modal adverbs*, 1/2012, pp. 99–119.
- Egg, Markus, “Discourse Particles at the Semantics-Pragmatics Interface”, in: Werner Abraham/ Elisabeth Leiss (eds), *Modality and Theory of Mind Elements across Languages*, Berlin 2012, pp. 297–333.

35 I want to thank Tilmann Köppe and Tobias Klauk who brought up this analogy in discussions.

- Groenendijk, Jeroen/Martin Stokhof, "Dynamic Montague Grammar", in: L. Kalman/L. Polos (eds.), *Proceedings of the Second Symposion on Logic and Language*, Budapest 1990, pp. 3–48, <http://dare.uva.nl/document/3702> (accessed March 30, 2013).
- Grosz, Patrick, "German 'doch': An Element that Triggers a Contrast Presupposition", forthcoming in: *Proceedings of the Chicago Linguistic Society*, 46/2010. Draft at <http://www.patrickgrosz.org/Publications.html> (accessed November 3, 2013).
- Heim, Irene, "File Change Semantics, and the Familiarity Theory of Definiteness", in: Rainer Bäuerle/Christoph Schwarze/Arnim von Stechow (eds.), *Meaning, Use and Interpretation of Language*, Berlin 1982, pp. 223–248.
- Heim, Irene/Kai von Fintel, "Intensional Semantics", in: *Online Textbook*, 2007, <http://semantics.uchicago.edu/kennedy/classes/s08/semantics2/vonfintel+heim07.pdf> (accessed December 29, 2011).
- Kamp, Hans, "A Theory of Truth and Semantic Representation", in: Jeroen Groenendijk et al. (eds.), *Formal Methods in the Study of Language*, Amsterdam 1981, pp. 1–41.
- Kamp, Hans/Uwe Reyle, *From Discourse to Logic*, Dordrecht 1993.
- Kania, Andrew, "Against the Ubiquity of Fictional Narrators", in: *The Journal of Aesthetics and Art Criticism*, 63/2005, pp. 47–54.
- Kaplan, David, "Demonstratives", in: John Almog/John Perry/Howard Wettstein (eds.), *Themes from Kaplan*, Oxford 1989, pp. 481–563.
- Karagjosova, Elena, "Adverbial 'doch' and the Notion of Contrast", in: B. Behrens/C. Fabricius-Hansen (eds.), *Oslo Studies of Language*, Vol. 1: *Structuring information in discourse: The explicit/implicit dimension*, Oslo 2009, pp. 131–148.
- Köppe, Tilmann/Jan Stühling, "Against Pan-Narrator Theories", *Journal of Literary Semantics*, 40/2011, pp. 59–80.
- Kratzer, Angelika, "Beyond 'oups' and 'ouch'. How Descriptive and Expressive Meaning Interact", Talk, Presented at the *Cornell Conference on Theories of Context Dependence*, 1999. Available at *Semantics Archive*.
- Lewis, David, *Counterfactuals*, Oxford 1973.
- Lewis, David, "Truth in Fiction", in: *American Philosophical Quarterly*, 15/1978, pp. 37–46.
- Pasch, Renate, *Handbuch der deutschen Konnektoren*, Berlin 2003.
- Portner, Paul, *What is Meaning? Fundamentals of Formal Semantics*, Malden, MA, 2005.
- Portner, Paul, *Modality*, Oxford 2009.
- Rett, Jessica, "A Degree Account of Exclamatives", in: *Proceedings of SALT*, 18/2008, pp. 601–618.
- Schlenker, Philippe, *Propositional Attitudes and Indexicality (A Cross-Categorical Approach)* [Diss. MIT], 1999.
- Schlenker, Philippe, "A Plea for Monsters", in: *Linguistics & Philosophy*, 26/2003, pp. 29–120.
- Schlenker, Philippe, "Context of Thought and Context of Utterance. A Note on Free Indirect Discourse and the Historical Present", in: *Mind and Language*, 19/2004, pp. 279–304.
- Schlenker, Philippe, "Indexicals" [Ms, NYU], 2010, forthcoming in: Sven O. Hansson/Vincent F. Hendricks (eds.), *Handbook of Formal Philosophy*, Dordrecht, Draft at <https://files.nyu.edu/pds4/public/> (accessed November 3, 2013).
- Sharvit, Yael, "The Puzzle of Free Indirect Discourse", in: *Linguistics and Philosophy*, 31/2008, pp. 353–395.
- Stalnaker, Robert, "Assertion", in: Peter Cole (ed.), *Syntax and Semantics*, Vol. 9, New York 1978, pp. 315–332.
- Stalnaker, Robert, *Context and Content*, Oxford 1999.

- Stalnaker, Robert, "Common Ground", in: *Linguistics and Philosophy*, 25/2002, pp. 701–721.
- Stechow, Arnim von, *Tenses in Compositional Semantics* [Ms. University of Tübingen], 2009, <http://www.sfs.uni-tuebingen.de/~astechow/Aufsätze/Approaches.pdf> (accessed August 2013).
- Wartenberg, Thomas E., "Need There Be Implicit Narrators of Literary Fictions?", in: *Philosophical Studies*, 135/2007, pp. 89–94.
- Zimmermann, Ede Thomas, "Kontextabhängigkeit", in: Dieter Wunderlich/Arnim von Stechow (eds.), *Semantics/Semantik. Ein Internationales Handbuch*, Berlin 1991, pp. 156–229.
- Zimmermann, Ede Thomas, "Context Dependence", in: Klaus von Heusinger/Claudia Maienborn/Paul Portner (eds.), *Semantics. An International Handbook*, Vol. 3, Berlin 2012, pp. 2360–2407.
- Zimmermann, Malte, "Zum Wohl: Diskurspartikeln als Satztypmodifikatoren", in: *Linguistische Berichte*, 199/2004, pp. 1–35.
- Zimmermann, Malte, "Discourse Particles", in: Klaus v. Heusinger/Claudia Maienborn/Paul Portner (eds.), *Semantics. An International Handbook*, Vol. 2, Berlin 2012, pp. 2012–2024.