Emar Maier*

Attributions of form and content

A discourse-structural account of reporting

Abstract: I propose a discourse-level analysis of report constructions. Indirect dis- 5
course, quotation, free indirect discourse, and attitude ascriptions are all analyzed in
terms of a discourse relation of Attribution connecting two propositional discourse units:
a frame (*he said*, *she dreamed*) and a (possibly complex, multi-sentence) report unit
(*"I'm an idiot"*, *(that) she was president*). I provide an underspecified semantics for
the discourse relation of Attribution that involves a flexible notion of 'characterization'. 10
A discourse unit may characterize a speech event by reproducing its linguistic surface
form (as in quotation) or its propositional content (as in indirect speech and attitude
reports), or both. I formalize this unified discourse-level Attribution approach within the
general framework of SDRT, and apply it to direct, indirect, and free indirect reports
that extend beyond the single embedded or quoted clause. This novel account is the first 15
to do justice account for the complex internal dependencies within stretches of reported
discourse that have thus far been mostly ignored by formal semanticist.

1 Introduction: discourse, coherence, and reporting

20

A correct interpretation of a multi-sentence discourse includes more information than is
contained in the interpretations of its individual sentences taken in isolation. Take the
mini-discourse in (1).

(1)   John was biking home late. A police officer stopped him. She give him a fine. 25
        His lights were off.

We naturally infer that a police officer stopped John *while* John was biking home
late *and then* the police officer gave John a fine *because* John's lights were off. The
individual sentences themselves describe states and events, which we as interpreters try

---

*****Corresponding author: Emar Maier,** University of Groningen

to combine into a coherent discourse by inferring various causal, temporal and other relations between these states and events (Hobbs 1979). These coherence inferences are defeasible and constrained by rationality, world-knowledge, a finite inventory of potential discourse relations (Narration, Background, Elaboration, Explanation, etc.),

5 and linguistic cues (an overt connective like *and then* would signal Narration, *because* would signal Explanation).

Now say the story continues with a question like (2).
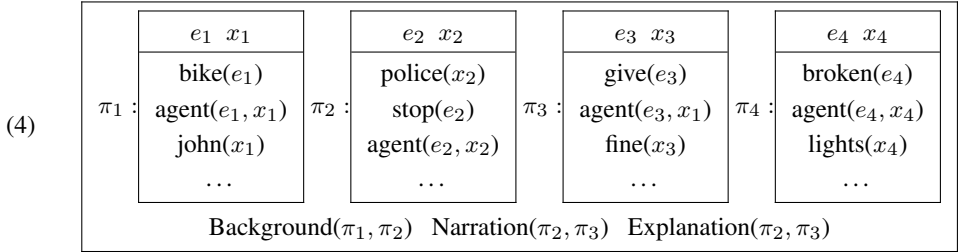
(2)　　What was he thinking?

In principle, (2) could represent a genuine question of the writer to the reader (perhaps

10 in a creative writing assignment), but, in the given narrative context, a much more likely interpretation is that this is rather a report of a question that one of the characters is asking. It could be the police officer reprimanding John by asking, somewhat sarcastically, "What were/are you thinking?", or perhaps it's John reflecting on his own actions, thinking to himself "What was I thinking?". In this paper I propose to account for these

15 different readings in terms of differences at the level of discourse structure. My proposal will be couched in the general framework of Segmented Discourse Representation Theory (SDRT, Asher and Lascarides 2003). Crucially, my analysis revolves around a dedicated discourse relation called Attribution. I will provide a very general semantics for Attribution in terms of an underspecified notion of characterization that covers the

20 full range of reporting types, from verbatim direct quotation to the paraphrasing of propositional content in attitude ascriptions.

2 Modeling coherence in SDRT

SDRT treats each individual sentence as a discourse unit, and formulates a number of axioms that model the establishment of discourse relations, like Narration, Result,

25 Contrast, and Elaboration, between semantic representations of discourse units. Other than competing theories of discourse structure it gives these relations a model-theoretic semantic interpretation. For instance, the story in (1) contains four elementary discourse units, typically labeled π_1, π_2, etc.

(3)　　π_1 : John was biking home late.

30 　　　　π_2 : A police officer stopped him.

　　　　π_3 : She gave him a fine.

　　　　π_4 : His lights were broken.

SDRT is compatible with any dynamic semantic interpretation for the individual discourse units, but I'll use DRT (Kamp and Reyle 1993), and extend its box-style notation to SDRSs as a whole, as illustrated in (4):[1]

(4)

| $\pi_1:$ | $e_1 \ x_1$ bike(e_1) agent(e_1, x_1) john(x_1) . . . | $\pi_2:$ | $e_2 \ x_2$ police(x_2) stop(e_2) agent(e_2, x_2) . . . | $\pi_3:$ | $e_3 \ x_3$ give(e_3) agent(e_3, x_1) fine(x_3) . . . | $\pi_4:$ | $e_4 \ x_4$ broken(e_4) agent(e_4, x_4) lights(x_4) . . . |

Background(π_1, π_2)   Narration(π_2, π_3)   Explanation(π_2, π_3)

Abstracting away from the semantic contents of the elementary units we can visualize [5] just the global coherence structure of the discourse as a graph:

(5)

$$\pi_1 \xrightarrow{\text{Background}} \pi_2 \xrightarrow{\text{Narration}} \pi_3$$
$$\pi_3 \xrightarrow{\text{Explanation}} \pi_4$$
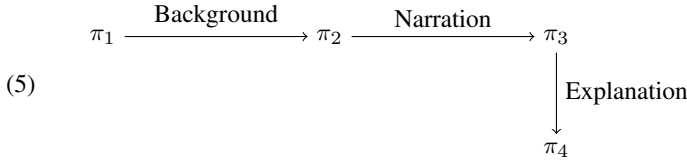
In these diagrams we stick with the standard SDRT convention of horizontal edges visualizing *coordinating* discourse relations, i.e. discourse relations like Narration and Background that in some intuitive sense move the story forward and change the active [10] topic, and vertical edges visualizing *subordinating* relations, i.e. relations like Explanation or Elaboration that don't move time and instead explore subtopics of the "dominant" node.[2]

The two main questions for a formally and practically usable discourse semantics are: how do we derive a graph representation like (5) or (4) from a discourse like (1), [15] and how exactly are we to interpret such formal structures? The SDRT framework provides two formal systems to answer these two questions. To start with the latter, the model-theoretic interpretation of an SDRT graph representation extends the standard DRT semantics for the graph's π_i-labeled DRS nodes with interpretation rules for the various discourse relations like in (6). Notation: K_{π_1} denotes the DRS unit that is labeled [20]

---

1 To avoid formal clutter in notation I leave π_i discourse referents out of the DRS universes and ignore the top-level π_0 altogether. In the examples I discuss these can always easily be reconstructed unambiguously.

2 The main advantage of this convention is to visualize the so-called Right Frontier Constraint that relates anaphora resolution to discourse structure. In this paper we are not concerned with anaphora resolution so we'll skip over this (Asher and Lascarides 2003)

with proposition label π_1; e_{π_1} denotes the main eventuality introduced in the universe of DRS unit labeled π_1; $[\![K]\!]$ denotes the dynamic semantic interpretation of a DRS (i.e. a context change potential, defined as a function from information states to information states, representing how an utterance affects an input context, à la Groenendijk and
5   Stokhof 1991); $\circ$ denotes function composition (i.e. the dynamic semantic analogue of conjunction); the symbol $\bigcirc$ in a DRS condition denotes temporal overlap between eventualities; $\prec$ denotes immediate temporal precedence (the second eventuality occurs right after the first):

(6)   a. $[\![\text{Narration}(\pi_1, \pi_2)]\!] = [\![K_{\pi_1}]\!] \circ [\![K_{\pi_2}]\!] \circ [\![e_{\pi_1} \prec e_{\pi_2}]\!]$
10   b. $[\![\text{Explanation}(\pi_1, \pi_2)]\!] = [\![K_{\pi_1}]\!] \circ [\![K_{\pi_2}]\!] \circ [\![\text{cause}(e_{\pi_2}, e_{\pi_1})]\!]$
         c. $[\![\text{Background}(\pi_1, \pi_2)]\!] = [\![K_{\pi_1}]\!] \circ [\![K_{\pi_2}]\!] \circ [\![e_{\pi_1} \bigcirc e_{\pi_2}]\!]$

In words, (6-a) says that a Narration relation between two discourse units means that we have to update the context with the contents of both discourse units, in order, and moreover the main eventuality described by the second unit, follows immediately after
15   the event described by that of the first.

      Now for the first question, how to derive a discourse structure representation like the graph (5) and ultimately the full SDRS (4) from a sequence of utterances? Let's assume that the elementary discourse units are already identified and assigned DRS representations by the standard DRS construction algorithm (see Kamp and Reyle 1993).
20   Now, SDRT's so-called Glue Logic provides defeasible inference rules that specify what discourse configurations trigger what discourse relations. For instance, a sequence of two discourse units where the first contributes a state and the second an event licenses the inference that they are connected by a Background relation – unless the resulting graph leads to an inconsistent or not maximally coherent final output representation. Similarly,
25   a sequence of two eventive units defeasibly triggers ($\rightsquigarrow$) a Narration connection.

(7)   a. $\text{state}(e_{\pi_1}) \wedge \text{event}(e_{\pi_2}) \rightsquigarrow \text{Background}(\pi_1, \pi_2)$
         b. $\text{state}(e_{\pi_1}) \wedge \text{event}(e_{\pi_2}) \rightsquigarrow \text{Narration}(\pi_1, \pi_2)$

We will not go into the formal details of either model theory or Glue Logic, nor into the presupposed DRS construction algorithm and dynamic semantics in terms of context
30   change potentials. I trust the above examples suffice to illustrate the SDRT framework to the uninitiated and defer to Asher and Lascarides (2003) for all details. In the following I provide an account of reported speech and quotation in this general framework, treating reporting as a discourse phenomenon, i.e., analyzing its semantic effects in terms of a semantically interpreted discourse relation of Attribution (Hunter 2016; Cumming
35   2020).

3 Indirect discourse as content attribution

3.1 From operators to event modifiers

Attitude and speech reports have occupied a central position in semantic theory from its very beginnings (Frege 1892). In contemporary possible worlds semantics, the intensional operator approach (Hintikka 1969) and its descendants (Kaplan 1989; Schlenker 2003) are still dominant. Recently, there's been a rise in event-based versions, where the attitude or speech verb introduces an event of thinking, speaking, hoping, and the complement clause specifies the content of that event (Kratzer 2006; Hacquard 2010) (notation: $^\wedge\varphi$ refers to the possible worlds proposition expressed by φ, which is just a traditional Montagovian way of dealing with intensionality).

(8)   a.   Mia said Don is a phony.
        b. $\exists e[\mathrm{say}(e) \wedge \mathrm{agent}(e, \mathrm{mia}) \wedge \mathrm{content}(e, {}^\wedge\mathrm{phony}(\mathrm{don}))]$
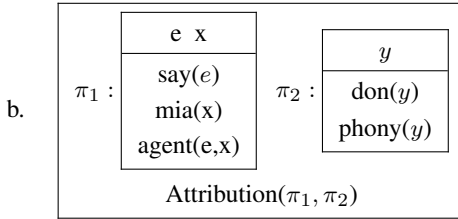
Such an analysis fits neatly in the general neo-Davidsonian framework by treating subject and complement uniformly as event modifiers. Instead of treating speech and attitude verbs as special operators it relies on the idea that there are certain events that have propositional contents. In this section I'll adopt the event-based approach but move it from the syntax–semantics interface into the discourse/pragmatics level.

3.2 From clausal complements to discourse units

When we look at a report like (8-a) from the perspective of discourse structure, the first question that arises is whether we are dealing with a single elementary discourse unit ('Mia said Don is a phony') or with two units ('Mia said something', 'Don is a phony') connected by a discourse relation. Hunter 2016 provides strong arguments for the latter, using a discourse relation of Attribution in SDRT to connect attitude and complement. Combining Hunter's Attribution with the event-based analysis of speech and attitudes gives the following semantic analysis:[3]

(9)   a. π_1 :Mia said. π_2 :Don is a phony.

---

3 An advantage of using the event-based approach here over the classic intensional operator approach that Hunter uses is that a unit of the form 'Mia said' in (9-a), without a grammatical object, is semantically speaking completely well-formed and interpretable in Neo-Davidsonian event semantics, while Hunter needs to stipulate some kind of existential closure in the logical form.

b.

$$\pi_1 : \boxed{\begin{array}{c} \text{e \ x} \\ \hline \text{say}(e) \\ \text{mia(x)} \\ \text{agent(e,x)} \end{array}} \quad \pi_2 : \boxed{\begin{array}{c} y \\ \hline \text{don}(y) \\ \text{phony}(y) \end{array}}$$

$$\text{Attribution}(\pi_1, \pi_2)$$

Following Hunter, Attribution is a non-veridical discourse relation, i.e. its truth does not presuppose the truth of both arguments. Specifically, π_2 serves just to characterize what Mia said, not what the world is actually like. We build this into our semantics as follows:

5 (10) $[\![\text{Attribution}(\pi_1, \pi_2)]\!] = [\![K_{\pi_1}]\!] \circ [\![\text{content}(e_{\pi_1}, {}^{\wedge}K_{\pi_2})]\!]$ (to be revised)

Note that this semantics presupposes that π_1 introduces a main eventuality (e_{π_1}) that can plausibly be said to have a propositional content, such as an utterance event, an occurrent thought, an attitudinal state, or a perceptual state/event. This requirement should be included in the antecedent of a defeasible Glue Logic axiom for inferring
10 Attribution, though we'll leave it open here what other requirements might further inform this inference:

(11) $\text{contentful.eventuality}(e_{\pi_1}) \wedge \ldots \leadsto \text{Attribution}(\pi_1, \pi_2)$

In any case, in our example this requirement is clearly met: π_1 introduces a saying event with Mia as agent, and saying is the prototypical contentful event. (10) then tells us that
15 the content of that event must be the proposition expressed by π_2, which corresponds exactly to what we get in the compositional semantics in (8-b).

3.3 Reports beyond the clause

In the case of (8-a) the discourse-level Attribution analysis above gives us the exact same semantic predictions as the classic sentence-level compositional account. The
20 crucial difference lies in the architectural fact that we're now assuming that the attitude verb plus clausal complement construction no longer encodes the event–content relation as a matter of compositional grammar, but instead is treated as a cue that informs the pragmasemantic Glue Logic of SDRS construction to infer the discourse relation of Attribution between two discourse units. The situation is similar to that with various
25 connectives: In traditional grammar, 'because' is just a lexical item interpreted as some kind of operator in the grammar, but in SDRT it is a marker guiding the inference of an Explanation relation between two discourse units. The powerful added machinery of the discourse-level approach is warranted by cases that are not overtly marked as

reports, but nonetheless interpreted as such. The most salient example of this is probably free indirect discourse, to be discussed in section 5. Below we first discuss another case that has received far less attention: report continuations beyond the overtly embedded complement clause.

Consider the following extended dream report: 5

(12)   Dan went to bed early. He dreamed that he was a frog. He jumped around a bit and then he was eaten by a stork.

On our discourse-level approach we parse this discourse as consisting of 5 elementary discourse units.

(13) π_1 : Dan went to bed early   10
         π_2 : He dreamed
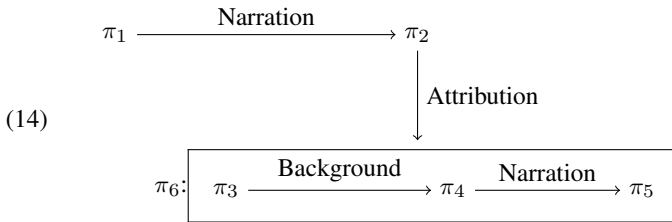 π_3 : (that) he was a frog
         π_4 : He jumped around a bit
         π_5 : (and then) he was eaten by a stork

These units are straightforwardly assigned simple DRS representations and then con- 15
nected by discourse relations to create an interpretable discourse structure graph. In the example, two discourse relations are arguably encoded grammatically: the complement construction in *dreamed that* encodes Attribution and *and then* encodes Narration. The rest of the relations can be defeasibly inferred by existing Glue Logic axioms, such as the sequence of events in π_1 and π_2 giving rise to a likely Narration inference, and the 20
sequence of state and event in π_3 and π_4 giving rise to a likely Background inference. Without going into Glue Logic details we're going to assume that under certain conditions, a coherent cluster of connected elementary units gets subsumed under a complex unit. Here, the most likely interpretation involves the construction of such a complex unit, π_6, out of the three horizontally connected elementary units describing the dream, 25
π_3 to π_5 and taking that complex unit (rather than just π_3) as the second argument of the Attribution.[4]

---

4 An interesting puzzle, beyond the scope of this paper, arises when we continue the discourse in (12) with 'He woke up screaming'. Evidently, the waking up is (probably) not part of the dream and hence should attach via Narration to π_2, but then we miss the fact that the waking up and screaming is (probably) causally related to the event of being eaten (or the dreaming of that event).

(14)

$$\pi_1 \xrightarrow{\text{Narration}} \pi_2$$

$$\downarrow \text{Attribution}$$

$$\pi_6: \quad \pi_3 \xrightarrow{\text{Background}} \pi_4 \xrightarrow{\text{Narration}} \pi_5$$

Crucially, this graph straightforwardly captures the 'modal subordination' (Roberts 1989) reading, where π_4 and π_5 are interpreted as describing the content of the dream, despite being syntactically outside the scope of the attitude verb. By contrast, the only way for a traditional sentence-level report semantics to deal with this would be to assume a silent dream operator in front of every proposition interpreted as a dream description. Note that such a sequence of hidden operators would still fail to capture the obvious discourse structural, temporal, and anaphoric relations between these segments.

Note that similar unmarked continuations of reports occur with other attitude and speech reports. In some languages, such syntactically unembedded continuations of speech reports can be marked with a reportative subjunctive mood on the verb:

(15)   Sie sagte sie habe   keine Zeit. Sie müsse   noch 86 Prüfungen
       She said she have-SUBJ *no   time. She must*-SUBJ *still   86 exams*
       bewerten.
       grade
       'She said she has no time. She still has 86 exams to grade (she said)' (German, Bary and Maier 2020)

In such constructions, the traditional, compositional approach would take the subjunctive morpheme as a semantic report operator (which causes significant complications for dealing with the overtly embedded subjunctive in the first sentence in (15), see Fabricius-Hansen and S{\textbackslash}a ebø 2004). On the current approach, we take the subjunctive merely as a grammatical cue that constrains the Glue Logic to block attachment of the current unit to a top-level unit, i.e. forcing it to attach to a unit under an Attribution.

In English, where we have no subjunctive inflection to mark something as reported content, we do occasionally find unmarked free standing clauses that are interpreted as speech report continuations:

(16)   Trump says he'll cut inflation in half. He'll also create record numbers of jobs and beat COVID before Christmas.

As in (14), by connecting the propositions about inflation, record job numbers, and COVID together into a complex unit (using coordinating, veridical relations like List

or Continuation between them), we automatically get the most likely reading where all
three are semantically interpreted as describing what Trump said, without relying on any
covert operators in the syntax.

4 Quotation as form attribution

The above event-based implementation of Hunter's discourse-structural approach to   5
Attribution applies to speech and attitude reports in the indirect discourse mode, where
we are reporting the content of another person's speech or attitudinal state in our own
words. I propose to generalize the semantics of the Attribution relation in order to
cover also quotation and free indirect discourse reports, which seem to exhibit similar
sensitivity to discourse structure, like allowing complex report continuations far beyond   10
the sentence level.

4.1 Direct discourse as pure quotation

We start with a simple, clausal, direct quotation. On an event-based account we can treat
direct and indirect speech uniformly as event modifiers, one that characterizes a speech
event by its propositional content, and one that characterizes it by its linguistic form   15
(Maier 2017):

(17)   a.   Mia said, "Don is a phony"
           b. $\exists e[\text{say}(e) \wedge \text{agent}(e, \text{mia}) \wedge \text{form}(e, \text{'Don is a phony'})]$

As with indirect reports I now propose a discourse-level alternative to this type of (near-
)compositional account that retains the idea of treating quotation as event modification.   20
We parse the quotation and the frame as distinct discourse units, connected by a discourse
relation of Attribution.

(18)

π_1:Mia said

$\quad\quad\quad$ | Attribution

π_2:"Don is a phony"
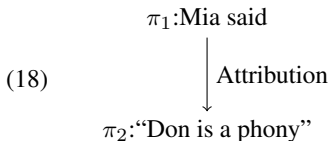
Now, to get the right truth conditions we could technically admit two distinct Attributions:
one defined as in (10), contributing $[\![\text{content}(e_{\pi_1}, {}^{\wedge}K_{\pi_2})]\!]$, and one, say QAttribution,   25
contributing instead something like $[\![\text{form}(e_{\pi_1}, \sigma_{\pi_2})]\!]$ (with σ_{π_2} denoting the linguis-
tic/graphemic/phonological surface form of speech act π_2). However, this move will lead

us down a path of multiplying coherence relations for each type of reporting, including mixed quotation, free indirect discourse, speech balloons, etc. In this paper I explore an alternative route where we stick with a single coherence relation of Attribution. To make this work we have to generalize its semantic contribution so that it subsumes both
5 form- and content-based reporting.

4.2 Attribution as underspecified event characterization

I propose to replace our original definition of the semantics of content-based Attribution in (10) with (19), which uses a more underspecified relation of 'characterization'. In this definition, $\text{Char}(\mathcal{F}(\pi), e)$, roughly '$\pi$ characterizes event e', I use the official SDRT
10 notation $\mathcal{F}$ for the function that maps labels in an SDRS to the SDRS constituents that they label – so far, $\mathcal{F}(\pi)$ is a notational variant for K_π but we'll introduce an extra feature in the SDRT syntax shortly that requires this slightly more general and official notation.
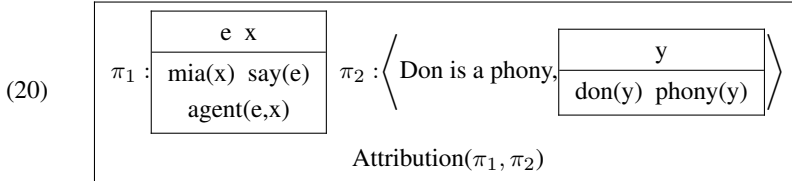
(19) $[\![\text{Attribution}(\pi_1, \pi_2)]\!] = [\![K_{\pi_1}]\!] \circ [\![\text{Char}(\mathcal{F}(\pi_2), e_{\pi_1})]\!]$

15 The idea behind (19) is that languages may allow different ways of characterizing what someone said, thought, or dreamed. We can characterize what someone said by reproducing its propositional content in our own words. That is what happens in indirect discourse reports, and it is exactly this type of 'loose' characterizing that is formalized explicitly in our original formulation of the semantics of Attribution in (10). But we can
20 also characterize what someone said at a more fine-grained level, by reproducing the exact words uttered. This is what happens in direct discourse.[5]

The proposed general approach to Attribution leaves us with the question of what to do with the actual quotation marks. Are they merely a cue to enforce the inference of an underspecified Attribution relation – the way we suggested treating the reportative
25 subjunctive mood in (15) above –, or are they a genuine semantic quotation operator applied to the second Attribution argument? The first option would mean that at the level of semantic representation, quoted sentences are treated just like any other discourse unit, i.e. parsed and assigned a DRS representation. But for reports with quotation marks we need more than just the semantic representation of the complement, we need access
30 to the actual form of the words used to express it. I propose that's what quotation marks

---

5 Below we'll encounter some other forms of characterization, such as simultaneous form and content characterization, and diagonal characterization. Ultimately I would propose to extend the same notion to capture the use of a picture characterizing what someone sees or thinks in a 'free perception sequence' or 'pictorial thought bubble' in comics and film narratives (Abusch and Rooth 2017; Maier and Bimpikou 2019), but that is beyond the scope of the current paper.

do: the tell the DRS construction algorithm to introduce a surface form into the semantic representation. For reasons to be discussed below we'll assume that we also construct the regular DRS representation of the quoted material, where possible. Hence, in the full SDRS representation of (18), the frame unit π_1 is represented as just a content DRS, while the quoted unit π_2 is represented as a form–content pair, consisting of a copy of the quoted surface form along with a DRS representation of its content.

(20)

$$\pi_1 : \boxed{\begin{array}{c} e \quad x \\ \hline mia(x) \quad say(e) \\ agent(e,x) \end{array}} \quad \pi_2 : \left\langle \text{Don is a phony,} \boxed{\begin{array}{c} y \\ \hline don(y) \quad phony(y) \end{array}} \right\rangle$$

$$Attribution(\pi_1, \pi_2)$$

We can now be more precise about the two most salient types of characterization that figure in the semantic definition of Attribution. First, propositional characterization: A DRS K propositionally characterizes a contentful eventuality e if the proposition expressed by K matches the propositional content of e.[6] Second, formal characterization: a form–content pair formally characterizes a speech or thought event e if the form component matches the linguistic form of the reported speech event.[7] We can rephrase this more formally as in (21), using the following notational conventions: $[\![\varphi]\!]_w^{f,c}$ is the (static)[8] semantic interpretation of an atomic DRS condition φ, i.e. its truth value relative to an assignment f, a Kaplanian context c and a possible world index w; $[\![\varphi]\!]^{f,c} = \lambda w[[\![\varphi]\!]_w^{f,c}]$, i.e. the proposition expressed by φ; and $Content$ and $Form$ are the by now familiar functions mapping certain events to their propositional contents and surface forms, respectively.

(21)   a. $[\![Char(K,e)]\!]_w^{f,c}$ is defined iff $f(e)$ is a contentful eventuality (speech event, belief state, etc.). If defined, $[\![Char(K,e)]\!]_w^{f,c} = 1$ iff $Content(f(e)) = [\![K_\pi]\!]^{f,c}$

---

6 I'm assuming here that propositional matching means identity between sets of possible worlds, but more realistically the original speech act may have been stronger than reported (I can report that Mary said that she's coming if she literally said something more specific, like "I'll be at the party between 9 and 10PM" (von Stechow and Zimmermann 2005; Abreu Zavaleta 2019).

7 Again, for simplicity I'll assume matching means identity between strings of letters or phonemes, though to model judgments regarding natural language quotation more realistically we have to make room for cleaning up false starts and filled pauses and allow literal translations, at the very least.

8 In DRT we typically use essentially static truth definitions for conditions as part of a definition of dynamic context change potentials for DRSs. See Kamp, van Genabith and Reyle (2003) for details.

b. $[\![\mathrm{Char}(\langle \sigma, K\rangle, e)]\!]_w^{f,c}$ is defined iff $f(e)$ is a linguistic speech act or language-like occurrent thought. If defined, $[\![\mathrm{Char}(\langle \sigma, K\rangle, e)]\!]_w^{f,c} = 1$ iff $Form(f(e)) = \sigma$ (to be revised)

The definition of characterization in (21) together with the general definition of Attribu-
5 tion from (19) allows us to model direct and indirect discourse uniformly. It effectively
recreates the truth-conditional predictions of a traditional account of direct discourse as
pure quotation, and a traditional account of indirect discourse as an intensional operator
(or contentful event) (Kaplan 1989; Brasoveanu and Farkas 2007; Maier 2017). A salient
feature of this account of quotation is that direct discourse reports can be interpreted
10 even if the quoted words are not interpretable or even parsable themselves:

(22) Mia stammered, "Wh..., wh-wh?...wheresil-uh-uh-egdesh?"

The reason we are able to correctly interpret gibberish quotes like (22) is that (21-b)
effectively ignores the second component of the form–content pair, i.e. the DRS parse
representing the content of the quoted words is just a dummy. The flip side of always
15 ignoring the semantic content of quoted passages in this way is that it suggests that
the quoted words are never really interpreted at all, they just contribute their form, i.e.
their 'shape' (D. Davidson 1979), to the eventual interpretation. This would be fine if all
we're interested in are the kinds of simple pure and direct quotations discussed in the
philosophical literature, like ''Boston' is a six letter word' and 'Otto said "I'm a fool"'.
20 When we're interested in more global discourse structures in actual text, this will prove
unsatisfactory.[9]

4.3 Complex quotations and simultaneous use and mention

Take a, still very simple, quotation like (23).

25 (23) "Oh, we'll be cutting," Trump told the audience. "But we're also going to have
tremendous growth."

---

9 Partee (1973) and others have already provided well-known arguments against the pure quotation
approach to direct discourse on the basis of anaphora and ellipsis dependencies between quotation and
surrounding discourse, as in:

(i) "Don't worry, my boss likes me! He'll give me a raise" said Mary, but given the economic
climate I doubt that he can. (Maier 2015)

On the analysis sketched above we could derive and interpret a graph with two separate Attributions, connecting the two quoted units to the intervening frame unit:

(24)

$$\pi_2\text{:Trump told audience}$$

Attribution

Attribution

$$\pi_1\text{:We'll be cutting} \qquad \pi_3\text{:We're going to have tremendous growth}$$
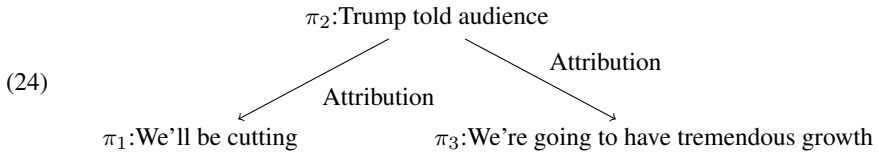
This graph would be semantically interpreted as conveying (i) that Trump spoke to some contextually salient audience (K_{π_2}); (ii) that the form of (part of) the speech act was 'we'll be cutting' (Attribution(π_2, π_1)); and (iii) that the form of (part of) the speech act was 'But we're also going to have tremen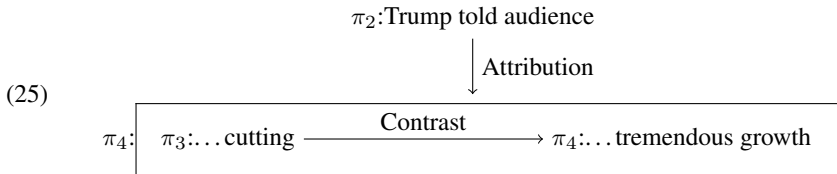dous growth' (Attribution(π_2, π_3)). The discourse in (23) clearly entails these three conditions, but by not actually interpreting the quoted material (other than as having a certain shape or form) we seem to be missing something important: a reader will naturally interpret the two quotation units and then on that basis infer some coherence relation between them. In this case presumably a relation of Contrast (marked with the overt 'But'), yielding a discourse graph with a complex unit:

(25)

$$\pi_2\text{:Trump told audience}$$

Attribution

$$\pi_4: \quad \pi_3\text{:}\dots\text{cutting} \xrightarrow{\quad\text{Contrast}\quad} \pi_4\text{:}\dots\text{tremendous growth}$$

In order to correctly infer coherence (and anaphoric) connections between multi-sentence quotations and derive graph structures like (25), the Glue Logic needs to have some access to the semantic content of quoted discourse units. This is the reason why I originally chose to semantically represent a quoted unit not as a mere surface form but as a pair, see (20) above. However, the semantic definition of formal characterization in (21) ignores the DRS component of the pair, making it semantically inert after all. Moreover, spelling out the full SDRS representation corresponding to (25) would gives us a complex unit, π_4, as the second component of our Attribution, which as it stands is not a form–content pair, and hence will not even trigger a quotational interpretation in the first place.

To remedy the last problem we extend the SDRT mechanism for forming complex units to take into account form–content pairs. In words, if we attach a form–content pair to another form–content pair we form a complex form–content pair by concatenating (notation: $\cap$) the two forms and combining the contents into a complex discourse unit, schematically:

(26) $\pi_1{:}\langle\sigma_{\pi_1}, K_{\pi_1}\rangle + \pi_2{:}\langle\sigma_{\pi_2}, K_{\pi_2}\rangle \;=\; \pi_3{:}\Big\langle\sigma^{\cap}_{\pi_1}\sigma_{\pi_2}, \boxed{\;\pi_1 : K_{\pi_1} \quad \pi_2 : K_{\pi_2}\;}\Big\rangle$

In other words, form components "project" from the elementary discourse units to the complex discourse unit containing them. Applied to our example we get the following full SDRS representation for (25):

5   (27)

$$\pi_4{:}\left\langle \begin{array}{l} \text{Oh, we'll be cutting.} \\ \text{But we're also going to} \\ \text{have tremendous growth} \end{array} \right., \begin{array}{l} \pi_2{:}\boxed{\begin{array}{c} e_2 \\ \hline \text{tell}(e_2) \end{array}} \\[4pt] \boxed{\pi_1{:}\boxed{\begin{array}{c} e_1 \\ \hline \text{cut}(e_1) \end{array}} \;\; \pi_3{:}\boxed{\begin{array}{c} e_3 \\ \hline \text{have.growth}(e_3) \end{array}} \\ \text{Contrast}(\pi_1, \pi_3) } \\[4pt] \text{Attribution}(\pi_2, \pi_4) \end{array} \right\rangle$$
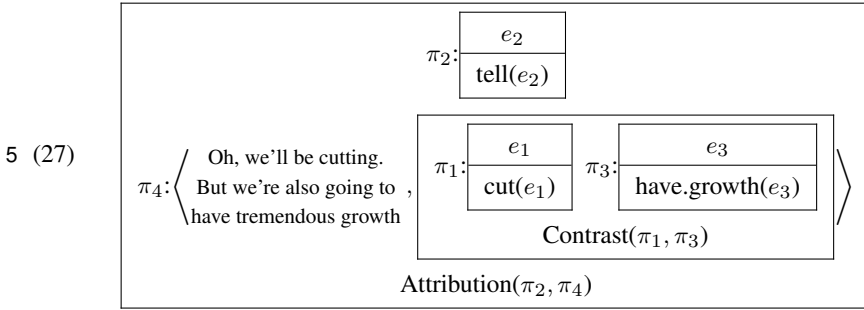
      The straightforward form-projection mechanism thus puts complex Attribution cases like (23) in the right format to feed into our semantics, as laid out in (19) and (21). But that semantics still ignores the semantic component of any quotations. To fix this we'll build on the idea that direct quotation tends to be interpreted as simultaneous use

10   and mention (D. Davidson 1979; Cappelen and Lepore 1997). I'll take here the most straightforward implementation of this idea, based on the two-dimensional account of mixed and direct quotation of Potts (2007):[10] a form–content pair $\langle\sigma, K\rangle$ characterizes a speech or thought event e if the first component σ formally characterizes e and the second component K propositionally characterizes e. If we stick with a single, classical

15   notion of propositional content here we have to follow Potts and assume a context shift: propositional characterization in the case of direct discourse must compare the content of e to the content of K relative to the shifted, reported context of utterance, not the actual, reporting context of utterance as in regular indirect discourse. This is necessary in order to get the reference of indexicals right: in direct discourse, all indexicals are

20   systematically shifted. I'll assume a function $Context$ mapping a speech/thought event to the context in which it occurs (Eckardt 2015).[11] In sum, we replace the second clause,

---

10 Maier (2015) accounts for this simultaneous use/mention intuition differently, viz. by treating direct quotation as essentially mixed quotation, and replacing the semantic interpretation of quoted material with a appropriately typed variable, leaving it for Gricean pragmatic/post-semantic reasoning to interpret the quoted words whenever that's called for.

11 $Context(e) = \langle w, t, x\rangle$ iff e occurs in w at time t and the agent of e is x. This is assuming events are world-bound. If we instead assume that a single event can occur in different possible worlds we would have to add the world as an extra parameter, i.e. $Context(e, w)$.

(21-b), in our general definition of characterization with a stricter definition that demands matching of form and content simultaneously, like this:

(28) $[\![\text{Char}(\langle \sigma, K \rangle, e)]\!]_w^{f,c}$ is defined iff $f(e)$ is a linguistic speech act or language-like occurrent thought. If defined, $[\![\text{Char}(\langle \sigma, K \rangle, e)]\!]_w^{f,c} = 1$ iff $Form(f(e)) = \sigma$ and $Content(f(e)) = [\![K]\!]^{f,Context(f(e))}$

5

Much more can be said about characterization than this. For instance, we could allow the content compartment to be empty and in such cases disregard it semantically, thereby subsuming the account of quoting gibberish as pure form quotation that we considered above. We could allow existentially quantified or presupposed variables of type t (or more complex types) in the content compartment to implement mixed quota- 10 tion (Geurts and Maier 2005; Maier 2014). We could also define an intermediate mode of characterization, somewhere in between formal and propositional characterization, viz. characterization at the level of Kaplanian character or its diagonal (Kaplan 1989; Stalnaker 1978; Zimmermann 1991). This would be useful for capturing monstrous or *de se* reports, but we could also use it to formalize the content matching requirement of 15 simultaneous form–content characterization for direct discourse, without the need of shifting the context of utterance.[12] Finally, and more tentatively, we might extend characterization beyond linguistic/contentful events to model demonstration more generally. For instance, the semantics of 'Mary ate like <gobbling gesture>' (K. Davidson 2015) would involve an event of eating being iconically characterized by a gobbling gesture. 20

Incorporating all these extensions and comparing various implementations is beyond the scope of this paper, which focuses on the general account of reporting as a discourse structural phenomenon. In the next section we will however incorporate a mechanism of unquotation, in order to deal with free indirect discourse.

---

12 To define this concisely we could use a well-known technical simplification, assuming that contexts ($c \in C$) and indices ($w \in W$) are tuples of the same type, i.e. indices are contexts with unused coordinates for agent, addressee, location etc., so that $C \subseteq W$ (von Stechow and Zimmermann 2005). Then we can easily define diagonal content as a a set of contexts:

(i)   diagonal DRS content: $\backslash K \backslash^f = \lambda c.\, [\![K]\!]_c^{f,c}$

We could now say that a discourse unit π with a DRS component K_π *diagonally characterizes* a contentful eventuality e if the '*de se* content' of e (the set of contexts 'compatible with e', Lewis 1979; Schlenker 2003) corresponds to the diagonal content of K_π.

5 A discourse-structural account of free indirect discourse

Free indirect discourse is a form of reporting speech or thought that shows characteristics of both direct and indirect discourse (Banfield 1982). Take (29).

5   (29)   Sue stared at the calendar. Oh no, she had to hand in that damn paper today! She'd never make it. . .

The first sentence is just a description of what's going on in the story world, but the next two seem to describe what's going on inside Sue's head. The way this 'perspective shift' is marked linguistically is often subtle but it involves a combination of the use of
10 expressive and indexical elements ('oh no', 'damn', 'today', '!') directly representing the protagonist Sue's point of view (i.e. as in direct speech), and the regular narrative past tense and third person pronouns ('she had to', 'she'd') representing the thinking protagonist from the narrator's 'third person' perspective (i.e. as in indirect speech).[13]
     Linguists have examined the semantic properties of free indirect discourse in some
15 detail, and have proposed various competing semantic analyses, e.g. in terms of monstrous indirect discourse (Sharvit 2008), the addition of an extra context parameter (Schlenker 2004; Eckardt 2014), and quotation plus unquotation (Maier 2015, 2017).
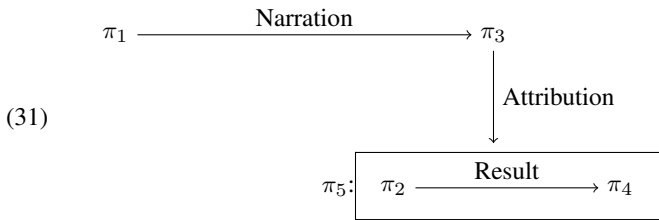   Some salient features of free indirect discourse that are often overlooked by semanticists are (i) that these types of reports tend to span several sentences or even entire
20 paragraphs, and (ii) that it may require intricate textual analysis to pinpoint exactly where such a report starts or ends. These neglected features however are more or less what we would expect on a discourse-structural approach. On our Attribution-based approach, once we have established that there's an Attribution, we get for each new incoming discourse unit a choice: do we attach it to the complex unit underneath that
25 Attribution (i.e., treat it as a continuation of the report), or to the main story line above it (i.e., treat is as a narrative description of the story world)? This choice is guided by often subtle considerations of global discourse coherence, i.e., which attachment generates a more coherent overall output SDRS (Asher and Lascarides 2003). Combined with the lack of clear, overt cues like quotation or (in English) subjunctive mood marking,
30 this explains the observed difficulty of determining the exact boundaries of free indirect discourse passages.
     Let me now flesh out the proposed discourse-structural Attribution account of free indirect discourse by applying it to the example in (29). Attuned to the grammatical cues

---

13 See Abrusán (2020a) for discussion of a more comprehensive algorithm for detecting 'perspective shift' based on grammatical, lexical and discourse-level cues.

for free indirect discourse detection, sketched above, we can recognize three discourse units, of which two form a complex node that is connected to the previous discourse via Attribution. But strictly speaking, Attribution can't attach to the first unit, because staring is not in any way a contentful or linguistic event that can sensibly be characterized by a form or a content. Following recent discourse-structural analyses of free indirect discourse (Abrusán 2020b; Bimpikou, Maier and Hendriks 2021; Altshuler and Maier 2020) I propose that we may in such cases accommodate a simple discourse unit, π_3, to introduce the required thought event.

(30) π_1 : Sue stared at the calendar.
          π_2 : Oh no, she had to hand in that damn paper today!
          π_3 : (she thought.)
          π_4 :She'd never make it. . .

(31)

$$\pi_1 \xrightarrow{\text{Narration}} \pi_3$$

$$\downarrow \text{Attribution}$$

$$\pi_5: \quad \boxed{\pi_2 \xrightarrow{\text{Result}} \pi_4}$$

Due to the inherent underspecification in the semantics of Attribution, this graph is in principle compatible with the various competing semantic analyses of the interpretation of free indirect discourse constructions. All that (31) tells us about the reports is that π_2 and π_4 together characterize the (accommodated) thought event in π_3. In its abstract graph form it doesn't specify what kind of characterization this is – simultaneous use/mention quotation, indirect discourse, or something else. But if we want to spell out the full SDRS box corresponding to the graph, and its interpretation, we'll eventually have to settle on a specific semantic theory. I'll explore here my own quotation-plus-unquotation approach.[14]

    Let's assume, following the argumentation in Maier 2015, that the DRS construction algorithm treats a free indirect discourse segment – recognized as such – as essentially quoted. This means that we introduce corresponding form layers for π_2 and π_4. But then, still following Maier 2015, pronouns and tenses are to be treated as 'unquoted'.[15]

---

14 A monstrous account à la Sharvit 2008 would involve defining a mode of characterization that preserves the character or diagonal for most of the report, but preserves only content for pronouns and tenses, presumably relying on some feature deletion mechanism already at the syntax/semantic level of DRS construction.
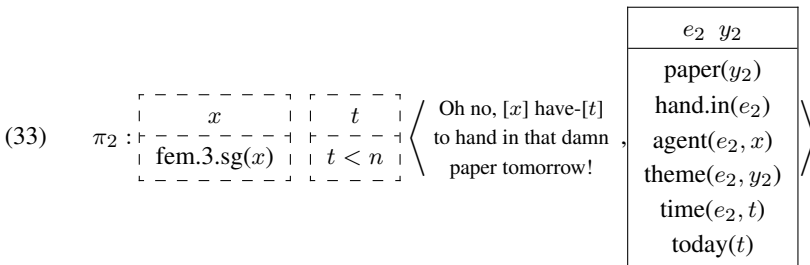
15 Maier 2017 seeks to derive the unquote-pronouns-and-tenses assumption from general pragmatic interpretation and production principles.

Technically, that means these pronouns and tenses are 'moved' out of the reports and interpreted separately, leaving (metalinguistic) traces (Maier 2014).
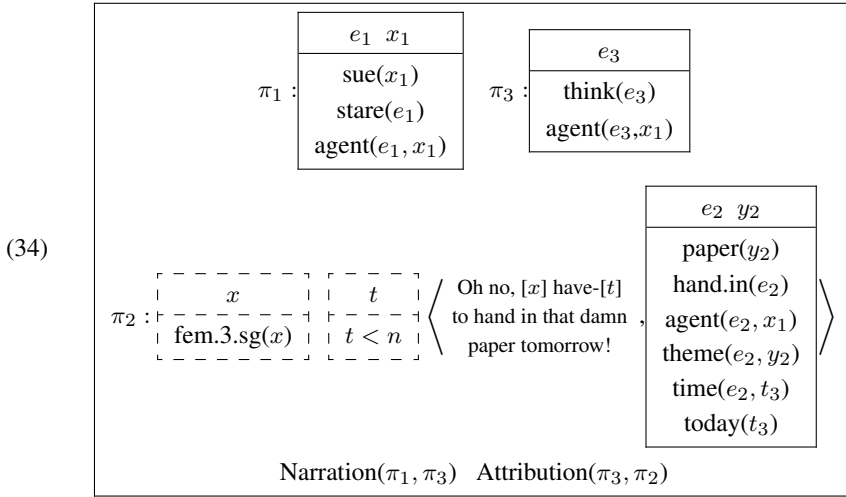
Let's go through the steps of the DRS construction algorithm for the first part of our example free indirect report. First, we assume a (usually covert) quotation with (covert) unquotation of all pronouns and tenses, (32-b). To interpret this we move the unquoted elements out of the quotation, (32-c).

(32) a. Oh no, she had to hand in that damn paper today!
      b. "Oh no, [she] have-[past] to hand in that damn paper today!"
      c. she_x $past_t$ "Oh no, [x] have-[t] to hand in that damn paper today!"

Now we apply the standard DRS construction algorithm to the expressions in (32-c). The two extraposed elements 'she_x' and '$past_t$' are anaphoric in nature and hence give rise to presupposition triggers (van der Sandt 1992), the quotation will give rise to a labeled form–content pair, consisting of the surface form (with two indexed holes) and a DRS box. The only new feature we have to add to the construction algorithm is how to deal with indexed holes in a surface form. Since the movement indexing ties each hole to a corresponding presupposition trigger, we can simply represent the contributions of the holes as the corresponding presupposed discourse referents, i.e. x and t, respectively.

(33) π_2 :

$$
\begin{array}{|c|}\hline x \\\hline \mathrm{fem.3.sg}(x) \\\hline \end{array}
\begin{array}{|c|}\hline t \\\hline t < n \\\hline \end{array}
\left\langle
\begin{array}{c} \text{Oh no, } [x] \text{ have-}[t] \\ \text{to hand in that damn} \\ \text{paper tomorrow!} \end{array}
,
\begin{array}{|c|}\hline e_2\ \ y_2 \\\hline \mathrm{paper}(y_2) \\ \mathrm{hand.in}(e_2) \\ \mathrm{agent}(e_2, x) \\ \mathrm{theme}(e_2, y_2) \\ \mathrm{time}(e_2, t) \\ \mathrm{today}(t) \\\hline \end{array}
\right\rangle
$$

We can now add (33) to the SDRS under construction by connecting its discourse label to a suitable existing label (e.g. to a thought event, through Attribution, or to another quoted or otherwise reported event under an Attribution, through Narration). Following the earlier graph structure in (31), we don't have an earlier Attribution so we have to accommodate a thought event unit π_3 and attach (33) to that with an Attribution, at which point we get the following SDRS:

(34)

$$\pi_1 : \boxed{\begin{array}{c} e_1 \ \ x_1 \\ \hline \text{sue}(x_1) \\ \text{stare}(e_1) \\ \text{agent}(e_1, x_1) \end{array}} \qquad \pi_3 : \boxed{\begin{array}{c} e_3 \\ \hline \text{think}(e_3) \\ \text{agent}(e_3, x_1) \end{array}}$$

$$\pi_2 : \boxed{\begin{array}{c} x \\ \hline \text{fem.3.sg}(x) \end{array}} \ \boxed{\begin{array}{c} t \\ \hline t < n \end{array}} \left\langle \begin{array}{c} \text{Oh no, } [x] \text{ have-}[t] \\ \text{to hand in that damn} \\ \text{paper tomorrow!} \end{array} \ , \ \boxed{\begin{array}{c} e_2 \ \ y_2 \\ \hline \text{paper}(y_2) \\ \text{hand.in}(e_2) \\ \text{agent}(e_2, x_1) \\ \text{theme}(e_2, y_2) \\ \text{time}(e_2, t_3) \\ \text{today}(t_3) \end{array}} \right\rangle$$

$$\text{Narration}(\pi_1, \pi_3) \quad \text{Attribution}(\pi_3, \pi_2)$$

Now we can resolve the presuppositions: x ('she') binds to x_1, the only salient female third person, and t binds to the time of the thinking (e_3). Now we add the final unit, π_4. We'll assume this is fed to the construction algorithm as a free indirect discourse, i.e. with quotation marks and unquotation holes, yielding a presuppositional form–content pair like (33). We attach this π_4 to the existing form–content pair, π_2, under the existing Attribution; create a complex discourse unit; project and concatenate the form components following (26); and bind π_4's unquoted tense and pronoun presuppositions. This gives the final output SDRS in (35), ascribing to Sue a complex thought whose form and content is characterized by two coherently connected discourse units.

(35)

$$\pi_1 : \boxed{\begin{array}{c} e_1 \ \ x_1 \\ \hline \text{sue}(x_1) \ \ \text{stare}(e_1) \\ \text{agent}(e_1, x_1) \end{array}} \quad \pi_3 : \boxed{\begin{array}{c} e_3 \ \ t_3 \\ \hline \text{think}(e_3) \\ \text{time}(e_3, t_3) \\ \text{agent}(e_3, x_1) \end{array}}$$

$$\pi_5 : \left\langle \begin{array}{c} \text{Oh no, } [x_1] \text{ have-}[t_3] \\ \text{to hand in that damn} \\ \text{paper tomorrow!} \\ [x_1] \text{ will-}[t_3] \text{ never} \\ \text{make it in time} \end{array} \right. , \quad \boxed{\begin{array}{l} \boxed{\begin{array}{c} e_2 \ \ y_2 \\ \hline \text{paper}(y) \\ \text{hand.in}(e_2) \\ \pi_2 : \ \text{agent}(e_2, x_1) \\ \text{theme}(e_2, y_2) \\ \text{time}(e_2, t_3) \\ \text{today}(t_3) \end{array}} \quad \pi_4 : \boxed{\begin{array}{c} e_4 \\ \hline \text{make.it}(e_4) \\ \text{agent}(e_4, x_1) \\ \text{time}(e_4, t_3) \\ \text{in.time}(t_3) \end{array}} \\ \text{Result}(\pi_2, \pi_4) \end{array}} \left. \vphantom{\begin{array}{c} a \\ b \\ c \\ d \\ e \end{array}} \right\rangle$$

Narration(π_1, π_3) Attribution(π_3, π_5)

6 Conclusion

I have proposed abandoning attempts to model reporting constructions in terms of various clausal operators integrated in a compositional semantics. Instead, we should
5 model them at the level of discourse structure. More specifically, I have proposed a discourse=structural account of all reporting in terms of a discourse relation of Attribution connecting two distinct discourse units: a frame unit ('she said', 'he dreamed') and a, possibly complex, report unit ('that he was unhappy', '"I'll beat COVID. But not global warming. That's still a hoax"'). I have proposed an underspecified semantics for the dis-
10 course relation of Attribution that involves a flexible notion of a speech/though/attitude eventuality being 'characterized' by a surface form or a propositional content, or both.

The proposed discourse-structural account is embedded in the general discourse semantics framework of SDRT. Clausal complements are simply analyzed as separate discourse units, represented by a labeled DRS in the discourse-level 'logical form'
15 (the SDRS). Quotation marks serve to introduce a surface form layer on top of the DRS representation of a quoted unit. These straightforward assumptions allow us to implement simultaneous use and mention for direct quotation, which I motivate with cases where multiple quoted sentences together form a complex discourse unit describing an internally coherent multi-sentence quotation.
20 More generally, it is such cases of extended direct, indirect, and free indirect reports, beyond the single reported clause, that have been the blind spots of traditional semantic

accounts of attitude reports and quotation and that motivate the proposed shift from syntax/semantics interface, to the level of discourse structure.

References

Abreu Zavaleta, Martín. 2019. Weak speech reports. *Philosophical Studies* 176(8). 2139–2166. https://doi.org/10.1007/s11098-018-1119-2.   5

Abrusán, Márta. 2020a. Computing perspective shift in narratives. In Emar Maier & Andreas Stokke (eds.), *The language of fiction*, forthcoming. Oxford: OUP.

Abrusán, Márta. 2020b. The spectrum of perspective shift: protagonist projection versus free indirect discourse. *Linguistics and Philosophy*. https://doi.org/10.1007/s10988-020-09300-z. https://doi.org/10.1007/s10988-020-09300-z.   10

Abusch, Dorit & Mats Rooth. 2017. The formal semantics of free perception in pictorial narratives. *Proceeding of the 21st Amsterdam Colloquium*.

Altshuler, Daniel & Emar Maier. 2020. Coping with imaginative resistance. Ms. https://philpapers.org/rec/ALTCWI.

Asher, Nicholas & Alex Lascarides. 2003. *Logics of conversation*. Cambridge University   15
Press. 556 pp.

Banfield, Ann. 1982. *Unspeakable sentences: narration and representation in the language of fiction*. London: Routledge & Kegan Paul.

Bary, Corien & Emar Maier. 2020. The landscape of speech reporting. Ms. Nijmegen/Groningen. https://ling.auf.net/lingbuzz/003526 (Accessed 27 September, 2018).   20

Bimpikou, Sofia, Emar Maier & Petra Hendriks. 2021. The discourse structure of free indirect discourse. In Mark Dingemanse & Elena Tribushinina (eds.), *LInguistics in the netherlands*. Amsterdam: John Benjamins.

Brasoveanu, Adrian & Donka Farkas. 2007. Say reports, assertion events and meaning dimensions. *Pitar Mos: A Building with a View. Papers in Honour of Alexandra Cornilescu*. 25
175–196.

Cappelen, Herman & Ernest Lepore. 1997. Varieties of quotation. *Mind* 106(423). 429–450. https://doi.org/10.1093/mind/106.423.429.

Cumming, Sam. 2020. Narrative and point-of-view. In Emar Maier & Andreas Stokke (eds.), *The language of fiction*, forthcoming. Oxford: OUP.   30

Davidson, Donald. 1979. Quotation. *Theory and Decision* 11(1). 27–40. https://doi.org/10.1007/BF00126690.

Davidson, Kathryn. 2015. Quotation, demonstration, and iconicity. *Linguistics and Philosophy* 38(6). 477–520. https://doi.org/10.1007/s10988-015-9180-1.

Eckardt, Regine. 2014. *The semantics of free indirect speech. how texts let you read minds* 35 *and eavesdrop*. Leiden: Brill.

Eckardt, Regine. 2015. Utterance events and indirect speech. *Grazer Linguistische Studien* 83. 27–46.

Fabricius-Hansen, Cathrine & Kjell Johan S{\textbackslash}a ebø. 2004. In a mediative mood: the semantics of the german reportive subjunctive. *Natural Language Semantics* 40 12(3). 213–257. http://link.springer.com/article/10.1023/B:NALS.0000034514.27887.d9 (Accessed 27 June, 2015).

Frege, Gottlob. 1892. Uber Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik* 100(1). 25–50.

Geurts, Bart & Emar Maier. 2005. Quotation in context. *Belgian Journal of Linguistics* 17(1). 109–128. https://doi.org/10.1075/bjl.17.07geu.

Groenendijk, Jeroen & Martin Stokhof. 1991. Dynamic predicate logic. *Linguistics and Philosophy* 14(1). 39–100. https://doi.org/10.1007/BF00628304.

Hacquard, Valentine. 2010. On the event relativity of modal auxiliaries. *Natural Language Semantics* 18(1). 79–114. https://doi.org/10.1007/s11050-010-9056-4. (Accessed 13 April, 2016).

Hintikka, Jaakko. 1969. Semantics for propositional attitudes. In *Models for modalities*, 87–112. Dordrecht: Reidel.

Hobbs, Jerry R. 1979. Coherence and coreference. *Cognitive Science* 3(1). 67–90. https://doi.org/10.1207/s15516709cog0301_4. (Accessed 13 January, 2020).

Hunter, Julie. 2016. Reports in discourse. *Dialogue & Discourse* 7(4). http://dad.uni-bielefeld.de/index.php/dad/article/view/3695 (Accessed 13 April, 2017).

Kamp, Hans, Josef van Genabith & Uwe Reyle. 2003. Discourse representation theory. In Dov Gabbay & Franz Guenthner (eds.), *Handbook of philosophical logic*, vol. 10, 125–394. Heidelberg: Springer.

Kamp, Hans & Uwe Reyle. 1993. *From discourse to logic: an introduction to modeltheoretic semantics in natural language, formal logic and discourse representation theory*. Vol. 1. Dordrecht: Kluwer.

Kaplan, David. 1989. Demonstratives. In Joseph Almog, John Perry & Howard Wettstein (eds.), *Themes from kaplan*, 481–614. New York: Oxford University Press.

Kratzer, Angelika. 2006. Decomposing attitude verbs. Handout. Honoring Anita Mittwoch on her 80th birthday. The Hebrew University of Jerusalem.

Lewis, David. 1979. Attitudes de dicto and de se. *The Philosophical Review* 88(4). 513–543. http://www.jstor.org/stable/2184843.

Maier, Emar. 2014. Mixed quotation: the grammar of apparently transparent opacity. *Semantics and Pragmatics* 7(7). 1–67. https://doi.org/10.3765/sp.7.7.

Maier, Emar. 2015. Quotation and unquotation in free indirect discourse. *Mind and Language* 30(3). 235–273. https://doi.org/10.1111/mila.12083.

Maier, Emar. 2017. The pragmatics of attraction: explaining unquotation in direct and free indirect discourse. In Paul Saka & Michael Johnson (eds.), *The semantics and pragmatics of quotation*. Berlin: Springer. http://ling.auf.net/lingbuzz/002966 (Accessed 3 May, 2016).

Maier, Emar & Sofia Bimpikou. 2019. Shifting perspectives in pictorial narratives. *Sinn und Bedeutung* 23.

Partee, Barbara. 1973. The syntax and semantics of quotation. In S. Anderson & Paul Kiparsky (eds.), *A festschrift for Morris Halle*, 410–418. New York: Holt, Rinehart and Winston.

Potts, Christopher. 2007. The dimensions of quotation. In Chris Barker & Pauline Jacobson (eds.), *Direct compositionality*, 405–431. New York: Oxford University Press.

Roberts, Craige. 1989. Modal subordination and pronominal anaphora in discourse. *Linguistics and Philosophy* 12(6). 683–721. https://doi.org/10.1007/BF00632602. (Accessed 25 October, 2016).

van der Sandt, Rob. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* 9(4). 333–377. https://doi.org/10.1093/jos/9.4.333.

Schlenker, Philippe. 2003. A plea for monsters. *Linguistics and Philosophy* 26(1). 29–120. https://doi.org/10.1023/A:1022225203544.

Schlenker, Philippe. 2004. Context of thought and context of utterance: a note on free indirect discourse and the historical present. *Mind and Language* 19(3). 279–304. https://doi.org/10.1111/j.1468-0017.2004.00259.x.

Sharvit, Yael. 2008. The puzzle of free indirect discourse. *Linguistics and Philosophy* 31(3). 353–395. https://doi.org/10.1007/s10988-008-9039-9.

Stalnaker, Robert. 1978. Assertion. In Peter Cole (ed.), *Syntax and semantics 9: pragmatics*, 315–332. New York: Academic Press.

von Stechow, Arnim & Ede Zimmermann. 2005. A problem for a compositional treatment of de re attitudes. In Greg Carlson & Francis Pelletier (eds.), *Reference and quantification: the partee effect*, 207–228. Stanford: CSLI.

Zimmermann, Thomas Ede. 1991. Kontextabhängigkeit. In Arnim von Stechow & Dieter Wunderlich (eds.), *Semantik: ein internationales handbuch der zeitgenössischen forschung*, 156–229. Berlin/New York: Walter de Gruyter.