Non-Canonical Applications of Topic Continuity: Restatement and Elaboration

Ekaterina Jasinskaja

1 Continuity in discourse

Givón (1983, pp. 7–8) points out three major aspects of discourse continuity which characterise a thematic paragraph: (a) thematic continuity; (b) action continuity; and (c) topics/participants continuity. The first one, thematic continuity, is the most difficult to operationalise and reflects the idea that the sentences within a paragraph are "about the same theme." Action continuity pertains to temporal sequentiality of the described events, which in other frameworks is known as a discourse relation of *Narration* or *Sequence*. Topic continuity in Givón's original sense characterises sequences of sentences that have a prominent common participant or participants, i.e. refer to the same individuals which can be viewed as the *leitmotif* or the protagonist of the discourse unit. The focus of this paper will be on the following cross-linguistic generalisation: "The more disruptive, surprising, discontinuous or hard to process a topic is, the more *coding material* must be assigned to it" (Givón, 1983, p. 18). This is equivalent to saying that the less coding material is present, the more continuous the topic must be, or in a binary, non-scalar fashion: by default (unless signalled otherwise) the topic does not change:

(1) *The Principle of Topic Continuity:*
By default, the discourse topic does not change.

This generalisation has already found its way into Optimality Theoretic pragmatics as a violable constraint on the interpretation of anaphoric expressions (cf. de Hoop, 2003; Beaver, 2004). Furthermore, Zeevat (2005) has recently suggested extending it to thematic continuity where the latter is understood as inertia in the maintenance of the question under discussion (QUD). I will follow Zeevat's terminology in which QUD is one of the possible interpretations of the notion of discourse topic, which promotes topic continuity to an overarching principle that comprises both thematic and participants continuity.

In this paper I explore some consequences of this principle in an approach to discourse interpretation similar to Zeevat's. More specifically, I investigate the relationship between topic continuity as a constraint on QUD management on the one hand and action continuity as well as participants continuity on the other. It will be shown that on the assumption that sentences are interpreted *exhaustively* with respect to their QUD's, (1) imposes strong requirements on the coreference of entities mentioned in the sentences. This does not only concern the coreference of nominal expressions referring to individuals, but also coreference between the described eventualities, which is characteristic of such discourse relations as *Restatement* and *Elaboration*. This leads to the unexpected result that the primary effect of topic continuity is establishing *Restatement* or *Elaboration*, rather than *Narration* (action continuity), which is traditionally viewed as *the* default discourse relation. Arguments against the default status of *Narration* have already been raised in the literature (e.g. by Asher and Lascarides, 2003). The present non-canonical understanding of the role of topic continuity realises the same idea. At the same time, the present approach does not exclude that topic continuity also plays a role in the inference of *Narration*, however, it will shown that *Narration* must involve a local violation of (1) whereas the topic may be maintained at a more global level.

This paper is structured as follows. Section 2 gives a rather dry outline of the formal framework, largely based on the approach to exhaustive interpretation proposed by van Rooy (2003); van Rooij (2004) and van Rooij and Schulz (2004) with a number of modifications which are (partly) explained and motivated in the subsequent sections, when applying the proposal to concrete examples of *Restatement*, *Elaboration*, and *Narration*. Section 3 demonstrates how the maintenance of the QUD at a local level leads to the inference of *Restatement* and *Elaboration*. Section 4 shows how a local change of topic makes the inference of *Narration* possible. Finally, general discussion and conclusions follow in Section 5.

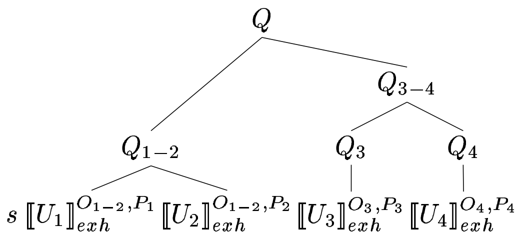

Figure 1: The interpretation of a sequence of declarative utterances $\langle U_1, ..., U_4\rangle$ and the assigned schematic QUD structure.

2 Outline of the framework

Our formal framework is constituted by a combination of discourse structure based on the notion of QUD along the lines of van Kuppevelt (1995) and update semantics (e.g. MDPL, Dekker, 1993) enriched with a notion of dynamic exhaustive interpretation (e.g. van Rooij and Schulz, 2004) and optimal interpretation function in the sense of van Rooij (2004). The interpretation of a monological discourse in the information state s is a sequence of optimal exhaustive updates $[\![\cdot]\!]^{O,P}_{exh}$ of s with the meanings of individual utterances, cf. Figure 1. This update function depends on two parameters—the optimal interpretation function O and the background predicate P. The background predicate is constrained by the information structure of the sentence. It is derived in a standard way by abstracting over the focused constituents. In the case of broad focus in utterances presenting events, the background predicate is assumed to be *happen(ed)*, reflecting the idea that the utterance answers a question like *What happened?* The optimal interpretation function O is an object of the same type as the conventional interpretation function F mapping expressions of the language (e.g. predicate symbols) to objects in the model (e.g. functions from worlds to sets of tuples of individuals); however O may deviate from F on demands of the current question under discussion reflecting the *sense* in which the expressions are used in the current utterance. The relationship between O and the QUD is explicated below in a number of steps.

The QUD: To begin with, the QUD is a partition of the set of possible worlds whose cells are associated with distinct most useful actions the agent is considering to take, which means that knowing the answer to the question determines the decision of the agent to act in a particular way (cf. van Rooy, 2003).[1] In this sense the QUD reflects the domain-level goals of the communication participants and is *a priori* independent from the literal semantics of the interrogative sentences they utter.

Relevance orderings: Relevance to the QUD is defined as a series of partial order relations on objects of different types. Relevance of propositions (sets of possible worlds) is defined in (2): roughly one proposition is more relevant than the other if it excludes (is inconsistent with) more cells of the QUD partition than the other.

(2) $p <_{rel} p'$ iff $\{q \in Q \mid p \subseteq \overline{q}\} \subset \{q \in Q \mid p' \subseteq \overline{q}\}$

In order to define the optimal interpretation of a predicate R in world w, we will need a notion of relative relevance of sets—possible extensions of R in w. Set Y is considered more relevant than X if among the *inertia worlds* of w, the set of worlds where Y is included in the conventional extension $F(R)$ constitutes a more relevant proposition than the set of worlds where X is included in $F(R)$, cf. (3). The inertia set of w, $I_R(w)$, is the set of worlds that are "similar" to w, i.e. differ from it at most in the extension of the predicate R and the predicates that depend on R syntactically or semantically, cf. (4).

(3) $X <_{R,w} Y$ iff $\{v \in I_R(w) \mid X \subseteq F(R)(v)\} <_{rel} \{v \in I_R(w) \mid Y \subseteq F(R)(v)\}$

(4) $I_R(w) = \{v \mid \forall P \in \mathcal{X}[F(P)(v) = F(P)(w)]\}$,
where $\mathcal{X}$ is a set of predicate symbols independent from R

[1]Strictly speaking, the QUD need not be a partition and may contain overlapping sets of worlds if the most useful action in some worlds in not uniquely determined. This gives rise to mention-some questions discussed by van Rooy (2003). Mention-some questions can be integrated in the present model, but will be ignored in this paper for space reasons.

Optimal interpretation: The optimal interpretation function O is derived from F, roughly, by picking the smallest ($min_{\subset}$) among the most relevant ($max_{<_{R,w}}$) interpretations of a predicate R from interpretations that are at least as specific as $F(R)$, cf. (5). $\mathcal{P}_R$ is the power set restricted by the meaning postulates that apply to R, e.g. if R is distributive, $\mathcal{P}_R(F(R)(w))$ is the set of all distributive subsets of $F(R)(w)$. Due to this restriction we only consider alternative extensions that are actually possible for R.

(5) $O(R)(w) = min_{\subset}(max_{<_{R,w}}(\mathcal{P}_R(F(R)(w))))$

Update: The optimal non-exhaustive update with respect to O is defined just like a standard update function (e.g. Dekker, 1993), except that the rule for atomic formulas uses the optimal interpretation function O instead of the conventional interpretation function F:

(6) $s[\![R(x)]\!]^O = \{\langle w, g\rangle \in s \mid g(x) \in O(R)(w)\}$

The exhaustive update of s with proposition ϕ with respect to an optimal interpretation function O and the background predicate P (contained in ϕ), $s[\![\phi]\!]^{O,P}_{exh}$, is an intersection of the optimal non-exhaustive update of the current information state s, $s[\![\phi]\!]^O$, and the optimal update of the *initial* information state s_0 minimised wrt. to a further relevance order relation $\prec_P$, cf. (7). In a semantics where information states are simply sets of worlds and no track is kept of the introduced discourse referents (e.g. Veltman, 1996), s_0 is just a set of all possible worlds that conform to the meaning postulates, i.e. s_0 contains no information except that about the lexical properties of the expressions of the language. If referents are taken into account, s_0 preserves this information (the assignment functions in s_0 are the same as in s) but loses the information about the world accumulated in s.

(7) $s[\![\phi]\!]^{O,P}_{exh} = s[\![\phi]\!]^O \cap min_{\prec_P}(s_0[\![\phi]\!]^O)$

Finally the $\prec_P$ relation on worlds (8) and world-assignment pairs (9) orders them along the proper inclusion relation between the optimal extentions of P in those worlds. This ensures the standard "only" effect of exhaustification: the worlds that are minimal wrt. $\prec_P$ are those where the optimal extension of P is smallest as long as this is consistent with ϕ, i.e. after minimisation only the individuals explicitly mentioned in ϕ have property P in the sense of O.

(8) $w_1 \prec_P w_2$ iff
a. $O(P)(w_1) \subset O(P)(w_2)$, and
b. $w_1 \in I_P(w_2)$

(9) $\langle w_1, g_1\rangle \prec_P \langle w_2, g_2\rangle$ iff $w_1 \prec_P w_2$ and $g_1 = g_2$

In sum, the QUD affects discourse interpretation by determining the relevance conditions and the optimal interpretation functions. It is assumed that legitimate sequences of QUDs in turn are subject to standard constraints, e.g. subquestion relationships between QUDs that stand in a dominance relation. Topic continuity (1) is a further constraint in this category. By requiring that adjacent utterances address the same QUD, it also ensures that they are interpreted wrt. the same optimal interpretation function. The consequences that the constancy of optimal interpretation has for the relations between exhaustively interpreted utterances are studied in the next section.

3 Local continuity: Restatement and Elaboration

The discourse in (10) is an instance of *Restatement* as the sentences describe the same event, i.e. the main eventualities of the sentences corefer.[2]

(10) Alena broke her skis. She lost her main transportation means.

In the framework outlined above, (10) is analysed as shown in (11). The sentences are interpreted with respect to the *same* optimal interpretation function O since, in accordance with the default assumption of topic continuity, they are dominated by the same QUD (like U_1 and U_2 dominated by Q_{1-2} in Figure 1). Assuming that both sentences have broad focus, the background predicate is $happen$ in both cases.

[2] The present notion of *Restatement* is broader than Mann and Thompson's (1988) as it includes all kinds of redescriptions of the same state of affairs, which need not be informationally equivalent.

(11) $s[\![\exists e_1[\text{ A. broke skis } (e_1) \wedge happen(e_1)]]\!]^{O,happen}_{exh}[\![\exists e_2[\text{ A. lost tr. means } (e_2) \wedge happen(e_2)]]\!]^{O,happen}_{exh}$

Informally, the exhaustive interpretation of the first sentence with respect to the background predicate $happen$ in the sense of O says that Alena breaking her skis is the only relevant event that happened. Similarly for the second sentence, Alena losing her main transportation means is the only relevant event that happened. Hence they must be the same event. Thus the identity of optimal interpretations of the background predicate leads to an event coreference, i.e. a *Restatement* reading for (10).

Technically, the derivation goes as follows. Suppose for simplicity that the optimal interpretation function O coincides with the conventional interpretation F, which would be the case if all events were relevant. The non-exhaustive update of s with the first proposition consists of world-assignment pairs $\langle w, g\rangle$ where the assignment function g maps e_1 to an eventuality in the intersection of the extensions of the predicates $\lambda e[\text{Alena broke her skis}(e)]$ and $happen$ in w. Apart from the eventuality that e_1 is mapped to, that intersection may or may not contain other elements, i.e. there may or may not be other events that happened. However, any inertia world w' of w where the extension of $happen$ contains some elements in addition to the referent of e_1 is greater than w with respect to the relevance order relation $\prec_{happen}$, hence $\langle w, g\rangle \prec_{happen} \langle w', g\rangle$. World-assignment pairs like $\langle w', g\rangle$ are eliminated from the information state as non-minimal wrt. $\prec_{happen}$ at the stage of exhaustification, cf. (7). Therefore, after exhaustive update of s with the sentence *Alena broke her skis*, the information state will only contain world-assignment pairs where the assignment maps e_1 to an event that belongs to Alena's ski-breakings and is the only event that happened in the given world. Similarly, the subsequent exhaustive update with the second sentence only contains world-assignment pairs $\langle w, g\rangle$ where g maps e_2 to an event of Alena losing her main transportation means which is the only event that happened in w. Obviously, only those world-assignment pairs survive both updates where the assignment function maps both e_1 and e_2 to the only event that happened in the respective world, which is only possible if e_1 and e_2 are mapped to the same event.

Elaborations that involve proper mereological part relationships between events are inferred via establishing a *Restatement* relation between descriptions of complex events that have simpler proper parts. E.g. in (12) coreference is established between the event e of John opening the door and the complex event described in the second sentence—the sum of the unlocking and the pushing events $e' + e''$, cf. (13).[3] The unlocking and the pushing in turn constitute a proper part of the opening.

(12) John opened the door. He unlocked it and pushed it.

(13) $s[\![\exists e[open(e) \wedge happen(e)]]\!]^{O,happen}_{exh}[\![\exists e'\exists e''[unlock(e') \wedge push(e'') \wedge happen(e' + e'')]]\!]^{O,happen}_{exh}$

Obviously, this derivation cannot proceed simply via inferring that opening and unlocking+pushing constitute the *only* relevant event that happened as in the previous example, since their proper parts (the unlocking and the pushing) happened as well. The solution to this problem is the same as that discussed by van Rooij and Schulz (2004) for the exhaustification of plural individuals. First, we assume that the domain of events is a complete join-semilattice closed under $+$.[4] For instance, if the domain of events contains simple events e_1, e_2 and e_3, it also contains their sums $e_1 + e_2 + e_3$, $e_1 + e_2$, and so on. Second, verbs like $happen$ have the property to distribute over subevents, so if a complex event happened, all the simpler events that it consisted of happened, too, and *vice versa*: $\forall e_1 \forall e_2[happen(e_1) \wedge happen(e_2) \leftrightarrow happen(e_1{+}e_2)]$. This is a meaning postulate that serves as a constraint on the model and the initial information state s_0, i.e. all worlds in s_0 conform to it. Hence $happen$ can have the extensions $\{e_1, e_2, e_1{+}e_2\}$, $\{e_1\}$, or $\{e_2\}$, but a world where the extension of $happen$ is e.g. $\{e_1{+}e_2\}$ is simply not available. Since exhaustification only selects among the worlds that actually exist, it will never contain worlds where only the complex event happened but not its proper parts. Thus ultimately, exhaustification gives us the *maximal* rather than the *only* event that happened in the relevant sense. Therefore event coreference is inferred via uniqueness of maximal events in the case of *Elaboration*.

Table 1 illustrates how this is achieved technically. It presents a number of "interesting" combinations of extensions of $open$, $unlock$, and $push$ (rows) with some of the possible extensions of $happen$ (columns). The first row contains worlds w where $F(open)(w) = \{e_1{+}e_2{+}e_3\}$, $F(unlock)(w) = \{e_1\}$, and $F(push)(w) = \{e_2\}$. This means that there is a complex opening event $e_1 + e_2 + e_3$ that involves unlocking (e_1) and pushing (e_2) as a proper part, but it also involves some other action e_3, which could

[3] The predicates $open$, $unlock$ and $push$ in (13) are abbreviations for $\lambda e[\text{John opened the door}(e)]$, etc.

[4] The $+$ operation is perhaps best interpreted as Link's (1983) material sum. That is, $e' + e''$ is a "complex" but singular event, rather than a plurality. However, in all currently relevant senses complex events behave like pluralities.

$happen \mapsto$	$\left\{\begin{matrix} e_1, e_2, e_3, \ldots \\ e_1+e_2+e_3 \end{matrix}\right\}$	$\left\{\begin{matrix} e_1, e_2, \\ e_1+e_2 \end{matrix}\right\}$	$\left\{\begin{matrix} e_1, e_3, \\ e_1+e_3 \end{matrix}\right\}$	$\left\{\begin{matrix} e_2, e_3, \\ e_2+e_3 \end{matrix}\right\}$	$\{e_1\}$	...
$open \mapsto \{e_1+e_2+e_3\}$ $unlock \mapsto \{e_1\}$ $push \mapsto \{e_2\}$	w_{11}	w_{12}	w_{13}	w_{14}	w_{15}	...
$open \mapsto \{e_1+e_2\}$ $unlock \mapsto \{e_1\}$ $push \mapsto \{e_2\}$	w_{21}	w_{22}	w_{23}	w_{24}	w_{25}	...
$open \mapsto \{e_1+e_2\}$ $unlock \mapsto \{e_1\}$ $push \mapsto \{e_3\}$	w_{31}	w_{32}	w_{33}	w_{34}	w_{35}	...
$open \mapsto \{e_1\}$ $unlock \mapsto \{e_1\}$ $push \mapsto \{e_2\}$	w_{41}	w_{42}	w_{43}	w_{44}	w_{45}	...
$open \mapsto \{e_1\}$ $unlock \mapsto \{e_2\}$ $push \mapsto \{e_3\}$	w_{51}	w_{52}	w_{53}	w_{54}	w_{55}	...
...	...	...	...	...	...	...

Table 1: The worlds in the body of the table represent a subset of the initial information state s_0. For each world in the given subset, the conventional interpretation $F(happen)$ in that world is given in the heading of the corresponding column, the values of $F(open)$, $F(unlock)$, and $F(push)$ in the corresponding row. On the assumption that the optimal interpretation coincides with the conventional interpretation, the marked worlds represent the following propositions: black ovals (dashed or solid line): $s_0[\![\exists e[open(e) \land happen(e)]]\!]^O$; black solid line ovals: $min_{\prec_{happen}}(s_0[\![\exists e[open(e) \land happen(e)]]\!]^O)$; grey ovals (dashed or solid line): $s_0[\![\exists e'\exists e''[unlock(e') \land push(e'') \land happen(e'+e'')]]\!]^O$; grey solid line ovals: $min_{\prec_{happen}}(s_0[\![\exists e'\exists e''[unlock(e') \land push(e'') \land happen(e'+e'')]]\!]^O)$.

be unbarring the door, or turning the door handle. The second row presents a case where unlocking and pushing is *all* the door opening is about. The complex opening event $e_1 + e_2$ contains an unlocking e_1 and a pushing e_2 as a proper part, but the unlocking and the pushing taken together exhaust $e_1 + e_2$. The third row corresponds to the possibility that the opening and the unlocking+pushing intersect but neither of them is part of the other. This could be the case, for instance, if John opened the door by unlocking and pulling it, and apart from that, pushed it (e.g. to close it again). In the forth row, opening is a proper part of unlocking+pushing. John opens the door by just unlocking it, and pushes it independently. Finally, the last row represents the possibility that opening, unlocking and pushing are three distinct events that have nothing to do with each other. What is interesting about these cases is that in all of them the literal conjunction of the clauses in (12) is satisfied since we indeed find an opening, an unlocking and a pushing event in all these worlds. The differences lie in the relations between these events, so next let's see how exhaustification constrains them.

Suppose once again for simplicity that all events are relevant, the optimal interpretation of the predicates coincides with their conventional interpretation, and the relevance ordering $\prec_{happen}$ on worlds reflects the set inclusion relations between the extensions of *happen* in those worlds. On these assumptions, the worlds marked black (dashed or solid line) in Table 1 have a non-empty intersection of *open* and *happen* and thus belong to the non-exhaustive update of s_0 with the first sentence *John opened the door*,[5] whereas the $\prec_{happen}$-minimal ones among them—$min_{\prec_{happen}}(s_0[\![\exists e[open(e) \land happen(e)]]\!]^O)$, cf. (7)—are marked with the solid line black ovals. Due to the distributivity of *happen*, the opening event is not necessarily the only event that happened in all these worlds, but it is the maximal such event, i.e. whatever else happened is a mereological part of it. Given that the second sentence *He unlocked it*

[5]Assignment functions are ignored henceforth.

and pushed it by topic continuity is interpreted with respect to the same optimal interpretation function, worlds marked with grey ovals in Table 1 belong to its non-exhaustive interpretation in s_0, solid line marking distinguishing the subset that corresponds to $min_{\prec_{happen}}(s_0[\![\exists e'\exists e''[unlock(e') \wedge push(e'') \wedge happen(e' + e'')]]\!]^O)$. As Table 1 shows, the only world that survives both minimisations is w_{22} in the second row of the table, where both unlocking and pushing are part of opening and opening has no other parts, i.e. that's all the opening is about. Thus it is predicted that the relation between the utterance *John opened the door* and each of the conjoined clauses in *John unlocked it and pushed it* is *Elaboration*, whereas the second utterance as a whole is a *Restatement* of the first.

All other relation possibilities mentioned above are ruled out by exhaustification. In the possibilities represented by the third, forth and fifth row of Table 1, the second sentence always introduces an event that is not part of the opening event mentioned in the first sentence, but this is already excluded by the exhaustive interpretation of the first sentence which says that whatever happened must be part of that opening event. This still allows for the possibility in the first row, where both unlocking and pushing form a proper part of opening but the latter also has other parts; however, the exhaustive interpretation of the second sentence only contains worlds where there are no happenings beyond unlocking and pushing, and that contradicts the existence of other components of the opening event.[6]

These examples demonstrate that applying topic continuity on the assumption that the utterances are interpreted exhaustively with respect to their topics (QUDs) has very strong consequences regarding the relations between the events described by those utterances. We always infer coreference between them (*Restatement*), whereas depending on the internal structure of those events, coreference may come along with a number of proper part-whole relations between the maximal events and their subparts, in which case *Elaboration* can be established. Thus the primary effect of topic continuity is the inference of a *Restatement* and *Elaboration* relations, rather than *Narration* or action continuity in Givón's sense. The latter in turn requires a violation of the continuity principle, at least a local one, as described in the next section.

4 Global continuity: Narration

In the above examples the derivation went through because the event descriptions presented by the sentences were compatible, i.e. applicable to the same event. However, this is not the case in (14). One of the assumed conditions on the identity of e_1 and e_2 is the identity of time points or intervals at which the events occur. But the world knowledge tells us that jumping on the horse and riding it cannot occur simultaneously, hence they cannot be the same event. Thus if an information state s containing this world knowledge is exhaustively updated with the sentences in (14) using the same optimal interpretation function, we arrive at an empty information state, i.e. the discourse is predicted to be inconsistent.

(14) The lone ranger jumped on his horse. He rode into the sunset.

The derived inconsistency can be viewed as a case in which the topic continuity default should be overridden at the local level. Once the interpretation with respect to the same QUD has failed, the next option to try is interpreting the sentences with respect to distinct but closely related QUDs, e.g. *What happened first?* and *What happened then?*, which are both dominated by a QUD like *What happened?* at the second level of QUD structure (cf. the relationship between Q_3, Q_4 and Q_{3-4} in Figure 1). These questions give rise to distinct $\prec_{happen}$ and $<_{happen,w}$ orderings, as well as distinct optimal interpretation functions O_1 and O_2, cf. (15).[7] Roughly, O_1 assigns the predicate symbol $happen$ an interpretation equivalent to *happen at* t_1, whereas in O_2, *happen* is understood as *happen at* t_2, where $t_1 < t_2$. If the predicate $happen$ is understood in these utterances in two different senses, no event coreference effect follows: from the fact that jumping on the horse is the only relevant event that happened at t_1 and riding into the sunset is the only relevant event that happened at t_2 it does not follow that jumping and riding are the same event.

[6]The latter inference is only valid if we minimise the update of the initial information state s_0, which contains no information beyond the meaning postulates, $min_{\prec_P}(s_0[\![\phi]\!]^O)$, and only after that intersect it with the non-exhaustive update of the current information state $s[\![\phi]\!]^O$, cf. (7). If $s[\![\phi]\!]^O$ is minimised directly, as in van Rooij and Schulz' (2004) original definition, then w_{11} rather than w_{12} (which is $happen$-smaller than w_{11}) would end up in $s[\![\exists e'\exists e''[unlock(e') \wedge push(e'') \wedge happen(e' + e'')]]\!]^{O,happen}_{exh}$ since w_{12} is not contained in s in the first place as a result of the previous update with *John opened the door*. See Jasinskaja (2006) for detailed argumentation for this move.

[7]Again, $jump$ and $ride$ abbreviate λe[Ranger jumped on horse(e)] and λe[Ranger rode into sunset(e)], respectively.

$happen \overset{F}{\mapsto}$	$\left\{\begin{matrix} e_1, e_2, \\ e_1 + e_2 \end{matrix}\right\}$	$\{e_1\}$	$\{e_2\}$	$\emptyset$
$happen \overset{O_1}{\mapsto}$	$\{e_1\}$	$\{e_1\}$	$\emptyset$	$\emptyset$
$happen \overset{O_2}{\mapsto}$	$\{e_2\}$	$\emptyset$	$\{e_2\}$	$\emptyset$
$jump \mapsto \{e_1\}$ $ride \mapsto \{e_2\}$ $at\ t_1 \mapsto \{e_1\}$ $at\ t_2 \mapsto \{e_2\}$	w_{11}	w_{12}	w_{13}	w_{14}
$happen \overset{O_1}{\mapsto}$	$\{e_2\}$	$\emptyset$	$\{e_2\}$	$\emptyset$
$happen \overset{O_2}{\mapsto}$	$\{e_1\}$	$\{e_1\}$	$\emptyset$	$\emptyset$
$jump \mapsto \{e_1\}$ $ride \mapsto \{e_2\}$ $at\ t_1 \mapsto \{e_2\}$ $at\ t_2 \mapsto \{e_1\}$	w_{21}	w_{22}	w_{23}	w_{24}

Table 2: For each world in the given subset, the conventional interpretation $F(happen)$ and two optimal interpretations $O_1(happen)$ and $O_2(happen)$ are given in the respective heading of the column, the values of $F(jump)$, $F(ride)$, $F(at\ t_1)$ and $F(at\ t_2)$ in the corresponding row. Worlds marked with black ovals belong to: $min_{\prec_{happen}}(s_0[\![\exists e[jump(e) \wedge happen(e)]]\!]^{O_1})$; grey ovals: $min_{\prec_{happen}}(s_0[\![\exists e'[ride(e') \wedge happen(e')]]\!]^{O_2})$.

(15) $s[\![\exists e[jump(e) \wedge happen(e)]]\!]^{O_1,happen}_{exh}[\![\exists e'[ride(e') \wedge happen(e')]]\!]^{O_2,happen}_{exh}$

Table 2 gives a brief illustration of how the change of QUD affects the optimal interpretation functions and ultimately prevents the event coreference reading for (14). The worlds represented in the table fall into three partition cells. For the first QUD *What happened at t_1?* these are $\{w_{11}, w_{12}\}$ "the lone ranger jumped on the horse at t_1," $\{w_{21}, w_{23}\}$ "the lone ranger rode into the sunset at t_1," and $\{w_{13}, w_{14}, w_{22}, w_{24}\}$ "nothing happened at t_1" (the intersection of $F(happen)$ and $F(at\ t_1)$ is empty). Obviously, w_{11} and w_{12} belong to the same cell of the partition since they only differ in whether or not e_2 happened, but e_2 is not located at t_1 and is therefore irrelevant. This makes these worlds and their extensions of *happen* incomparable wrt. the $\prec_{happen}$ and $<_{happen,w_{11}/w_{12}}$ relations, respectively. For instance, the inertia set of w_{11}, $I_{happen}(w_{11})$ is $\{w_{11}, w_{12}, w_{13}, w_{14}\}$—all the worlds that only differ in the extension of *happen*. Among these worlds, only in w_{11} is it the case that $\{e_1, e_2, e_1 + e_2\}$ is a subset of the $F(happen)$, and only in w_{11} and w_{12} is $\{e_1\}$. By definition of $<_{rel}$ (2), the propositions $\{w_{11}\}$ and $\{w_{11}, w_{12}\}$ are equally relevant since they exclude the same cells $\{w_{21}, w_{23}\}$ and $\{w_{13}, w_{14}, w_{22}, w_{24}\}$ of the QUD partition, therefore the sets $\{e_1, e_2, e_1 + e_2\}$ and $\{e_1\}$ are also equally relevant wrt. $<_{happen,w_{11}}$ by (3). By contrast, $\{e_1\}$ and $\{e_2\}$ are not equally relevant, $\{e_2\} <_{happen,w_{11}} \{e_1\}$, since $\{w_{11}, w_{13}\}$—the set of worlds where $\{e_2\}$ is a subset of $F(happen)$—only excludes one partition cell $\{w_{13}, w_{14}, w_{22}, w_{24}\}$, whereas $\{w_{11}, w_{12}\}$ excludes two as mentioned above. In this way we obtain that $\{e_2\}$ (as well as $\emptyset$) is smaller wrt. $<_{happen,w_{11}}$ than $\{e_1, e_2, e_1 + e_2\}$ and $\{e_1\}$. The latter two sets are therefore maximal among the distributive subsets of $F(happen)(w_{11})$, so $max_{<_{happen,w_{11}}}(\mathcal{P}_R(\{e_1, e_2, e_1 + e_2\})) = \{\{e_1, e_2, e_1 + e_2\}, \{e_1\}\}$. Minimising wrt. proper set inclusion gives us $\{e_1\}$ as the optimal interpretation of *happen* in w_{11}: $O_1(happen)(w_{11}) = \{e_1\}$. By repeating the same procedure for other worlds we get the O_1 values represented in Table 2. Basically, what we get is that $O_1(happen)(w)$ is reduced to the intersection of $F(happen)(w)$ and $F(at\ t_1)(w)$, which is what we want—*happen* is understood in the sense of *happen at t_1*.[8] The second QUD *What happened at t_2?* in turn has the cells $\{w_{11}, w_{13}\}$ "the lone ranger jumped on the horse at t_2," $\{w_{21}, w_{22}\}$ "the lone ranger rode into the sunset at t_2," and $\{w_{12}, w_{14}, w_{23}, w_{24}\}$ "nothing happened at t_2." Once again, the optimal interpretation function O_2 reduces *happen* to the intersection of the conventional interpretation of *happen* and *at t_2* in the corresponding worlds, cf. Table 2.

[8]Since the extension of *at t_1* (or any other implicit restriction) may change from world to world and we want O_1 to map *happen* to the events that happened at t_1 in *that* world, it is essential to constrain the compared propositions in the definition of $<_{R,w}$ (3) to the inertia worlds of w.

Given two different interpretations for $happen$, the rest of the inference proceeds as if the propositions in (15) were exhaustivized with respect to two different predicates. Thus the non-exhaustive update of s_0 with *The lone ranger jumped on the horse* only contains worlds like w_{11} and w_{12} (black ovals in Table 2) where the intersection of the *optimal* interpretations of $happen$ and $jump$ is non-empty.[9] Although the conventional extension of $happen$ in w_{11} is bigger than in w_{12}, the optimal extensions are equal, therefore w_{11} is not comparable to w_{12} by $\prec_{happen}$, and is not excluded by minimisation in the exhaustive update. Similarly, both w_{11} and w_{13} will survive in the non-exhaustive update of s_0 with the second utterance *He rode into the sunset* as well as its minimisation wrt. to $\prec_{happen}$ (this time based on O_2). These propositions are compatible as they both contain worlds like w_{11} where the lone ranger jumped on the horse at t_1 and rode into the sunset at t_2 (and no other events happened except the sum of these two, which is required by the distributivity of $happen$). Thus such worlds will constitute the consecutive exhaustive update of both utterances (15).

The above derivation shows how a local change of QUD and optimal interpretation function relaxes the inferred relations between the described eventualities, making it possible to establish relations other than *Restatement* and *Elaboration*. Splitting the main QUD into subquestions that address successive time points gives us *Narration*, whereas other splittings can lead to other discourse relations. It goes beyond the scope of this paper to motivate the particular choice of subquestions that violate topic continuity. What is important is that the inference of a *Narration* relation *requires* such a violation, assuming that utterances are interpreted exhaustively with respect to their QUDs. At first glance, this seems to contradict the common view (see e.g. Asher, 2004) that utterances connected by *Narration* are also united by a common topic. However, there is no contradiction given that the topic may continue at the second more global structural level, whereas the local topics are distinct, as in the case of Q_3, Q_4 and Q_{3-4} in Figure 1. Thus I propose that the role played by topic continuity in establishing "action continuity" is restricted to this more global level.

5 Discussion

The present proposal assigns a more prominent role to the topic continuity constraint in the inference of discourse relations than assumed so far, although this is associated with a non-standard ranking of defaults among discourse relations. On the assumption that utterances are interpreted exhaustively with respect to their topics, the default discourse relations turn out to be *Restatement* and *Elaboration* since they constitute the primary effect of topic continuity. Relations like *Narration* in turn require a local violation of this principle. If we assume that the derived inconsistency is enough to trigger such a violation, the proposed reranking of defaults is probably harmless since the effects of topic continuity are so strong that it does not take much to contradict them. This would be consistent with the intuition that *Restatement* is "rare," whereas *Narration* is one of the most common relations between utterances about events.

However, it was independently pointed out that sometimes the inference of *Narration* fails in the absence of explicit cues like *and then*, which can give rise to mildly suboptimal discourses like (17a), although *Narration* would be consistent with the semantics of the sentences.

(16) What did Kim do today?

(17) a. ? Kim watched TV. She studied.
b. Kim watched TV. Then she studied.

Cases like this made Asher and Lascarides (2003) argue against the traditional assumption that *Narration* is the default discourse relation. The proposed *Restatement* default could be employed to capture the infelicity of (17a) on the assumption that the violation of topic continuity does not come for free (even when the world knowledge contradicts event coreference) and indeed, as suggested by Givón's (1983) generalisation, requires the presence of appropriate linguistic cues to license a violation. It remains to be clarified how exactly topic continuity should interact with consistency and other pragmatic constraints in that case. It is clear, however, that it has different and much stronger effects on the inference of discourse relations than previously believed, which seems to be good rather than bad news since it means that more discourse relations can be derived from independently motivated general pragmatic principles.

[9] The optimal interpretations of other predicates contained in the sentences, e.g. $jump$ and $ride$, are affected by the implicit restriction $at\ t_1$ or $at\ t_2$ in the same way as in the case of $happen$. This effect can be ignored in the analysis of the current example, but it can be essential in other cases (see Jasinskaja, 2006, for detailed discussion).

References

Asher, N. (2004). Discourse topic. *Theoretical Linguistics*, 30:163–201.

Asher, N. and Lascarides, A. (2003). *Logics of Conversation*. Studies in Natural Language Processing. Cambridge University Press.

Beaver, D. I. (2004). The optimization of discourse anaphora. *Linguistics and Philosophy*, 27(1):3–56.

de Hoop, H. (2003). On the interpretation of stressed pronouns. In Blutner, R. and Zeevat, H., editors, *Pragmatics in Optimality Theory*, pages 25–41. Palgrave Macmillan.

Dekker, P. (1993). *Transsentential Meditations*. PhD thesis, ILLC-Department of Philosophy, University of Amsterdam.

Givón, T. (1983). Topic continuity in discourse: An introduction. In Givón, T., editor, *Topic Continuity in Discourse*, pages 1–42. John Benjamins, Amsterdam.

Jasinskaja, E. (2006). *Pragmatics and Prosody of Implicit Discourse Relations: The Case of Restatement*. PhD thesis, submitted to the University of Tübingen.

Link, G. (1983). The logical analysis of plurals and mass terms: A lattice-theoretical approach. In Bäuerle, C. S. R., Schwarze, C., and von Stechow, A., editors, *Meaning, Use and the Interpretation of Language*, pages 303–323. de Gruyter, Berlin.

Mann, W. C. and Thompson, S. (1988). Rhetorical Structure Theory: Toward a functional theory of text organization. *Text*, 8(3):243–281.

van Kuppevelt, J. (1995). Discourse structure, topicality and questioning. *Journal of Linguistics*, 31:109–147.

van Rooij, R. (2004). Pragmatic meaning and exhaustification. Handout for the presentation at the ESSLLI Workshop on Implicature and Conversational Meaning organized by Bart Geurts and Rob van der Sandt. Nancy, France.

van Rooij, R. and Schulz, K. (2004). Pragmatic meaning and non-monotonic reasoning: The case of exhaustive interpretation. © Kluwer Academic Publishers.

van Rooy, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy*, 26:727–763.

Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, 25:221–261.

Zeevat, H. (2005). Optimality theoretic pragmatics and rhetorical structure. Handout of the talk at the Sixth International Tbilisi Symposium on Language, Logic and Computation. September 12–16, 2005, Batumi.