Decalog 2007: Proceedings of the 11th Workshop on the Semantics and Pragmatics of Dialogue, pages 65–72.
Trento, Italy, 30 May – 1 June 2007. Edited by Ron Artstein and Laure Vieu.

Clarification Requests: An Incremental Account

Ruth Kempson
King's College London
ruth.kempson@kcl.ac.uk

Andrew Gargett
King's College London
andrew.gargett@kcl.ac.uk

Eleni Gregoromichelaki
King's College London
eleni.greg@blueyonder.co.uk

Abstract

In dialogue models, fragment clarification requests (CRs) characteristically involve pre-processing, lifting the fragment to sentential level or coercing the context (Ginzburg and Cooper, 2004; Purver, 2004). This paper introduces an incremental account of CRs using the Dynamic Syntax (DS) dialogue model (Purver et al., 2006) that allows CRs as interruptions with no structure-specific stipulations. Generation of CRs is licensed at any point in the tree growth process which constitutes the build-up of utterance interpretation. The content attributed to CRs in context is one step more advanced than what has been achieved by the (interrupted) parse, either querying lexical content, checking/querying means of identification in context, or checking/querying resulting content (in the last of these, update may be trivial). Fragment responses (FRs) may reconstruct the apparent source of difficulty from the CR parse providing/confirming update from that reconstructed partial tree. However, the FR may constitute a trivial update of the clarifier's own context (the latter being the tree-representation of their initiating utterance), as the CR has been equally parsed via trivial context-update. All ambiguities arise from interaction of lexical specification, available partial structure as context, and available means of update: no irreducible ambiguity is required.

1 Introduction

Accounts of clarifications presume, following Ginzburg and Cooper (2004), that clarification-request fragments (CR) bifurcate according to whether what is queried concerns contentious content (the "clausal reading") or problematic identification of the meaning of the word used ("constituent reading"), the latter taken as a distinct "anaphoric utterance" use, with both being assigned a propositional-type construal. However, not only can it be shown that propositional-type analyses are not necessary in accounting for such ellipsis construals, as we shall see in due course, but it is also well-known that clarification requests and their fragment response can be made incrementally at a sub-sentential level:

(1) A. They X-rayed me, and took a urine sample, took a blood sample.
A: Er, the doctor
B: Chorlton?
A: Chorlton, mhm, he examined me, erm, he, he said now they were on about a slight [shadow] on my heart.

Furthermore, there is a broad range of readings associated with such fragments which do not seem to fall easily into two such clear-cut categories. To illustrate, we set out the following possible modes of clarifying the subject of a statement using the repeated fragment (CR) with its equally fragmentary reply (FR), and outline some of the different possible CR construals when the time-linear dimension of the parse is taken into account:

(2)
	A (female):	Bill left.
(i)	B (male):	Bill?
(ii)	B:	"Bill"?
	A:	Bill (Smith).

Case (i) of B's responses is a CR that can be paraphrased in terms of the whole of A's original utterance, in other words, as *Bill left?*. One might distinguish three reasons to justify the utterance of such a CR: (a) the entire utterance has been understood, but the CR conveys doubt of the involvement of the individual referred to; (b) although who is intended has in principle been identified, confirmation is still requested for certainty; (c) the meaning of the word is understood, hence the sentence successfully understood *qua* sentence, but the query is a request for provision of information to identify who is being referred to in the face of lacking this information. B's response (ii), as annotated, might seem to be construed as making a meta-linguistic response, and there are arguably three bases which suggest this form of construal: (a) the word *Bill* has been parsed, but uncertainty as to who A is talking about has led to B abandoning the parse at that juncture without establishing a full understanding of the sentence; (b) B fears he has misheard, and (on the basis of some word segment he has heard) is guessing what was said (e.g. here B might say *Bill* and be right, or *Jill* and be wrong), and (c) where B is explicitly asking for a repeat of the information provided by that word. There are thus a considerable number of different ways of grounding CR uses.

Three features of CRs provide clues as to how best to model them. First, they repeat specific material from the context. Unlike standard questions, this type of clarification is not about requesting new information from interlocutors (as with WH-questions), but focuses on repeating items from (the immediate) context. Second, their brevity opens up a range of possible interpretations, not always distinguishable. Third, they have a characteristic intonation, whose function is to indicate some non-canonical mode of interpretation in response to the immediate context (Rodriguez and Schlangen, 2004).

This paper presents the claim that the Dynamic Syntax (DS) model of dialogue (Purver et al., 2006) extends seamlessly to these phenomena. The account of clarificatory requests (CR) and fragment replies (FR) allows incremental request/provision of clarification at arbitrary points in the dialogue, while retaining a unitary characterisation of the lexical input. There is no need for coercion operations in order to resolve the fragment (Ginzburg and Cooper, 2004; Purver, 2004; Purver, 2006). We shall also argue that the distinction between clausal and constituent CRs emerges as a consequence of clarification being possible for all licensed tree transitions, including those involving the update provided by the word itself, so there is no recourse to stipulated input ambiguity between clausal and constituent CR's. The analysis of CR's and FR's furthermore fits directly within an overall account of ellipsis that construal of fragments is determined by structures/formulae/actions that context directly provides (Purver et al., 2006; Cann et al., 2007).

2 Previous Literature

As a form of nonsentential utterance (NSU), CRs have typically been modelled through preprocessing of some kind. Approaches adopt either a syntactic approach lifting them to sentence level (assuming missing information is "hidden"), or a semantic one, raising the information presented by some previous sentence so this can combine with the content of the fragment to yield back a propositional content (for representative papers see Elugardo and Stainton, 2005). A third approach associated with Ginzburg and colleagues (eg Ginzburg and Sag, 2000; Purver, 2004; Ginzburg and Cooper, 2004; Fernández-Rovira, 2006) both lifts the fragment to a clausal level and processes contextual information (which they term *context coercion*) (Purver, 2006).

This last approach has been described as incremental in involving phonological, syntactic, and semantic projection of subparts of complex signs in parallel as information becomes available (Ginzburg and Cooper, 2004). However, it is also desirable that computational accounts meet a notion of incrementality in which projection of structure/interpretation follows as closely as possible word-by-word processing with progressive interaction between linguistic and contextual information for which there is psycholinguistic evidence (see Pickering and Garrod among others); and the DS model of dialogue (Purver et al., 2006) purports to match this, as part of meeting the Pickering and Garrod challenge that formalisms for language modelling should be evaluated by how good a basis they provide for reflecting patterns that occur in conversational dialogue.

3 Dynamic Syntax: Background

Dynamic Syntax (DS) is a parsing-based approach to linguistic modelling in which syntax is defined as the progressive projection of semantic representations from the words taken in left-to-right sequence. Such representations take the form of decorated (linked) binary branching trees representing predicate-argument structures, with each node decorated with a sub-term of some propositional formula. The interpretation process is defined as goal-directed growth along various dimensions of tree decoration: type and formula decorations ($Ty(X)$, $Fo(Y)$), tree-node identification ($Tn(Z)$), and tree-relations (see below). Formula decorations are lambda terms of the epsilon calculus, with all quantified terms of type e, their restrictor being part of the term.[1]

The central tree-growth process of the model is defined in terms of the procedures whereby such structures are built up; taking the form of general structure-building principles (*computational actions*) and specific actions induced by parsing particular lexical items (*lexical actions*). The core of the formal language is the modal tree logic LOFT, which defines modal operators $\langle\downarrow\rangle$, $\langle\uparrow\rangle$, which are interpreted as indicating daughter and mother relations, respectively, $\langle\uparrow_*\rangle$, $\langle\downarrow_*\rangle$ operators characterizing *dominate* and *be dominated by*, and two additional operators $\langle L\rangle$, $\langle L^{-1}\rangle$ to license paired *linked* trees. Tree nodes can then be identified from the rootnode $Tn(0)$ in terms such as $\langle\uparrow\rangle Tn(0), \langle\uparrow_*\rangle Tn(0)$, etc. The actions defined using this language are transition functions between intermediate states, which monotonically extend tree structures and node decorations. The concept of *requirement* is central to this process, $?X$ representing the imposition of a goal to establish X, for any label X. Requirements may thus take the form $?Ty(t)$, $?Ty(e)$, $?Ty(e \rightarrow t)$, $?\langle\downarrow\rangle Ty(e \rightarrow t)$, $?\exists x Fo(x)$, $?\exists x Tn(x)$, etc.

All aspects of underspecification have an associated requirement for update. Pronouns illustrate formula underspecification, the pronoun *he* being assigned lexical actions from a trigger $?Ty(e)$ that projects a metavariable $Fo(\mathbf{U}_{Male(\mathbf{U})})$ of $Ty(e)$ with requirement $?\exists x Fo(x)$ (also a case requirement); and such metavariables are replaced by a Substitution process from a term available in context. We assume that the restriction $Male(\mathbf{U})$ would be specified as resulting from an action to construct a LINK transition to a tree of topnode to be decorated as $Male(\mathbf{U})$ as part of the actions encoded by the pronoun *he* (the mechanism of constructing a LINK relation being the means of constructing paired trees to be evaluated as compound forms of conjunction: Cann et al., 2005).

The process is thus essentially representational: the resolution of pronoun construal is established as part of the construction process. We propose that names too project a metavariable, eg *Bill* projecting a metavariable which we annotate as $Fo(\mathbf{U}_{Bill'(\mathbf{U})})$, with instruction to construct a LINK transition to a linked tree of topnode $Ty(t)$ decorated with the formula value $Bill'(\mathbf{U})$, characterising the predicate 'being named Bill', this constituting a constraint on the logical constant to be assigned as construal of the use of that name in the particular context.[2] We shall represent such logical constants, m_{21}, m_{22} etc, as having an attendant predicate attribute, eg $(m_{21,Bill'(m21)})$, but these are short-hand for the projection of such a pair of linked trees, one containing an argument node decorated with a formula (m_{21}) of type e, linked to a tree with topnode decorated with the formula $Bill'(m_{21})$.

The construction of structurally underspecified relations is also licensed (displayed in trees as a dashed line), with construction of nodes through an operation **Adjunction* licensing construction from a node $Tn(a)$ of a node described only as $\langle\uparrow_*\rangle Tn(a)$, an underspecification which is resolved, if introduced early on in a parse, only at a later point in the parse, when this characterisation can be satisfied by some introduced node of appropriate type. A variant, *Late*Adjunction*, applies to an initiating node of a given type to induce a dominated node requiring the same type, which with subsequent parse provides a basis for update to that initiating node, hence to some interim metavariable decorating it: Cann et al. (2005) analysed expletive pronouns in these terms.

Since, in any parse sequence, there may and characteristically will be more than one update possibility, a parse state P is defined as a set of triples $\langle T, W, A\rangle$, where: T is a (possibly partial)

[1] These take the form of variable-binder, variable of type e, and restrictor. Composite restrictors can be constructed through the building of linked trees, the resulting propositional content then by a step of LINK-evaluation, taken as an enrichment of the restrictor-specification (Kempson et al., 2001).

[2] Such an analysis suggests presuppositions in general involve constructing linked trees (Cann et al., 2005, ch.8).

tree; W is the associated sequence of words; A is the associated sequence of lexical and computational actions. Context is then defined in similar terms. At any point in the parsing process, the context $\mathcal{C}$ for a particular partial tree T in the set P can be taken to consist of: a set of triples $P' = \{\ldots, \langle T_i, W_i, A_i\rangle, \ldots\}$ resulting from the previous sentence(s); and the triple $\langle T, W, A\rangle$ itself.[3] Wellformedness is then definable as the availability of at least one sequence of transitions from some partial tree-specification as output to some complete tree with topnode decorated with a formula of type t having used all the words in sequence and with no outstanding requirements, a characterisation which Cann et al. (2007) extend to define a concept of context-dependent wellformedness.

In Purver et al. (2006) generation is defined to follow the parsing dynamics, this being the core mechanism, but it too is goal-directed: speakers have a goal tree representing what they wish to communicate, and each licensed step of the update transition defined by the core formalism constitutes the grounding for some possible generation step subject to a requirement of a subsumption relation between the constructed parse tree and the goal tree, in the sense of allowing a successful derivation from the parse tree as updated to the goal tree. Incremental (word-by-word) parsing, and lexicon search for words which provide appropriate tree-update relative to this goal tree enables speakers to produce the associated natural language string (see Purver et al., 2006). A generator state G is thus a pair (T_G, X) of a goal tree T_G and a set X of pairs (S, P), where: S is a candidate partial string; P is the associated parser state (a set of $\langle T, W, A\rangle$ triples). Search for appropriate words is said to be made from context wherever possible, reducing the production task.

Ellipsis, in both parsing and generation equally, involves use of context in a number of different ways. Strict readings of (VP) ellipsis involve taking some formula value as the value of the metavariable supplied at the ellipsis site. Sloppy readings of such fragments reuse sequences of actions stored in context, leading to different information given their re-application relative to the partial tree provided by the construal of the fragment itself. Answers to questions involve using some structure in context as their point of departure, the answer expression providing the update to that structure to yield some propositional formula. In the generation of such ellipses, the same parse actions are subject to the added restriction that the update to the partial tree under construction subsume the goal tree. What integrates these accounts of different elliptical forms is that each makes direct use of some attribute of context, without any coercion of the context prior to such use, thereby dramatically reducing the parsing/production task, as full lexicon search is side-stepped.

4 Towards an Incremental Account of CRs

In the general case, parsing and generation are presumed to start from the Axiom, the initial one-node tree $?Ty(t)$ and reach some goal tree $Ty(t), Fo(\alpha)$ via an accumulated sequence of transitions across partial trees, but this restriction is not essential: both parse and generation tasks may start from arbitrarily rich partial trees and end at any richer partial tree (see Purver et al., 2006 for an account of split utterances that depends on this). It is these partial tree inputs and outputs which constitute the core of the CR account.

The general schema is as follows. We take questions overall to be an encoding of a explicit request for coordination with some other party with respect to input provided by the question form. There are two core cases: those where some particular (*wh*-marked) constituent is signalled as being the information to be provided by the answer;[4] and those where the request concerns some whole propositional unit (polar interrogatives), which may be marked by word order or often merely by intonation alone. However, there is also a whole range of cases where individual words, their intrinsic system-based meaning, or their particular context-established construal may constitute the request for explicit coordination. These are the CR cases – a fragment associated with an explicit coordination request. Given the DS account of dialogue, all such fragments are taken to be both understood or have their production licensed relative to whatever structure is provided in context, whether a partial tree representation, with pointer indicating where the emergent growth of some tree

[3] For simplicity, we shall generally take this to comprise the triple P' resulting from A's initial utterance, and any partial trees established in subsequent parsed fragments associated with clarification of aspects of P'.

[4] See Kempson et al., 2001, where it is argued that *wh* expressions encode specialised meta-variables.

structure has got to, or a completed propositional tree representation. The encoding that this is a coordination request we take, at this juncture, simply to be a primitive Q annotation, to be removed by rule in the process of response.[5]

Informally, then, the dynamics of how such CR's or FR's come to be uttered is as follows. The idea is that the formal account should model very directly the dynamics of information processing – A, the initiator starts to say something, and clarification can be requested and provided at any point whenever B "gets stuck". What this will mean in DS terms is the construction of some partial pointed tree which will then constitute the context for B's interruption: the goal of the CR is then the request for provision of a value at exactly that point, with intonation playing the critical role of indicating where that is. The goal tree for such an interruption is characteristically just one point of update from the partial tree established at that achieved parse state, or may even be identical to it. A then may reply with the needed clarification, often also a fragment (FR), both being able to rely on re-use of actions from context to round out the interpretation intended. With uncertainty in the parse process in principle possible at any point in the parse sequence, requests for clarification may occur at any point in the parse-update process.

In the case of (2), this yields at least the following possibilities, each tree displaying the construction step immediately upon uttering the word *Bill*.

I: B may have failed to fully parse the first word but makes a stab at what it was, his goal tree being a structure constituting a request for confirmation of the provided name-based update (Q is taken to decorate the node indicated by intonation):

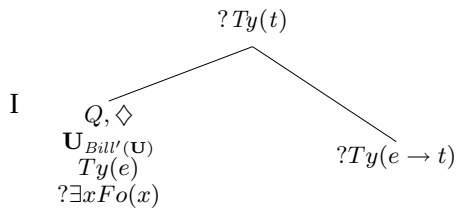

[5]This is clearly only an intermediate formulation, but the critical aspect is that it not be presented as itself in predicate-argument form in the representation, unless this is explicitly made clear through words whose content is to present such a request.

In this case, the goal tree set up contains a decoration as provided by the name but no identification of the individual in question. It is notable that if a word is even to be guessed at, it will induce tentative update of the partial tree, hence characterising even clarifications at the word-level as just one among a whole set of possible bases for clarification, without need of any concept of anaphoric utterance (contra Ginzburg and Cooper, 2004)

II: B may have successfully processed the first word but, not knowing who is being talked about, wishes to establish this before proceeding to the development of the predicate node (the analogue of incremental anaphora resolution). Such a request can be achieved by repeating the name, because a licensed parse step is to build an unfixed node (by *Late*Adjunction*) from the node already decorated with $\mathbf{U}_{\mathbf{Bill'(U)}}$ thereby providing an additional node decorated with $?Ty(e)$ which will license the update induced by the parse of the repeated name and lead via unification of the two nodes back to the parse tree which constituted the source of his request for clarification:

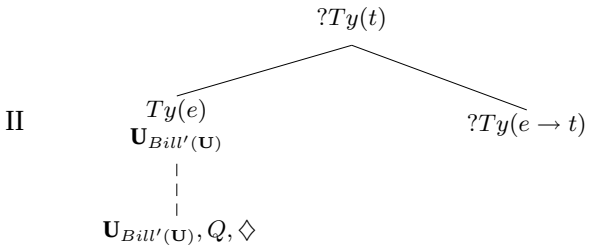

In other words, the assumption of a goal tree and a (distinct) parse tree can be retained even in cases where some understood word is nevertheless being repeated.

III: B may have understood the first word, using it to establish a type value and a place-holding metavariable, without having been able to identify who is being talked about. Nonetheless, because the word itself is processed the parse can continue. The predicate value can then be established (which may help to identify the individual in question). Yet, in coming to build up a whole propositional content, B may still yet fail to identify who is being talked about and so need to repeat the word as before. This would be the analogue of expletive pronouns, for which a type value is assigned to the node in question early on in the interpretation process, but the formula value is established at a relatively late point:

III

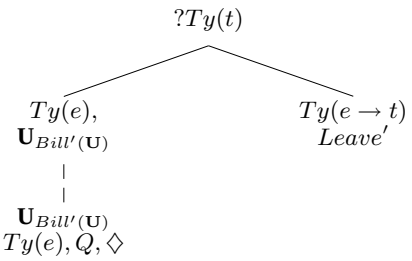

IV: B may have fully understood what A has said but might wish to have this confirmed, and thus decides on a goal tree hopefully identical with that which he has just retrieved in parsing A's utterance. *Late*Adjunction* can be applied relative to this structure also (with pointer return to the subject node as in parsing the answer to questions), again enabling the parse actions associated with *Bill* to be added to the introduced node, albeit one which will turn out to be a trivial update:

IV

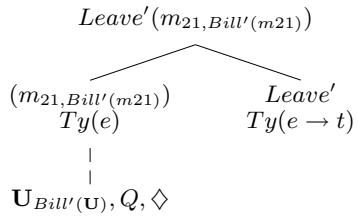

There are more possible interim goal trees on the basis of which B might seek clarification. In each case, the effect of adopting a goal tree which is some minimal (possibly null) enrichment from the partial tree established in the parse process (taken as context for the generation task) ensures that clarificatory fragments are assigned a context, parse tree and goal tree which are (all but) identical. This may seem to render the account trivial, with so little differentiation between context input to the generation task, goal output of the task, and parse-tree indicating the current update; but the near-identity is definitive of the clarification task. The goal tree in such cases is not some radical enrichment of the input parse tree: to the contrary, what is requested is either suggestions for or confirmation of some current putative tree-growth step, in order to proceed. With all these signals, in conjunction with the indicative intonation across the fragment in question, it follows that as long as A successfully parses B's fragment, she will have a basis for identifying the source requiring clarification, and hence the basis for a reply.

4.1 Detailing a CR/FR Exchange

In each such clarification exchange, there are altogether three turns, with each turn having a goal and parse tree for the speaker, and a parse tree for the hearer. Figures 1-2 detail trees for a simple request for clarification, where both *Bill* and *left* have been parsed, upon the assumption that B has parsed A's input and recovered a decoration for the subject-node but without identifying which Bill is being talked about. Figure 1 schematically represents the generation of B's CF. With the current parse tree as context and input to the generation task, and the goal of querying the update to that subject decoration, B can make use of *Late*Adjunction*, licensing the construction of a node decorated with $?Ty(e)$ in order to provide a vehicle for licensing the lexical actions of the word *Bill*, i.e. the update with metavariable $\mathbf{U}_{Bill'(\mathbf{U})}$ which would then license unification with the already decorated subject node, yielding back the partial tree which was his parse tree as context differing from it only in the decoration Q which constitutes the request. The focus here is on modelling CR as an interactive strategy for repairing potential misalignment (eg Pickering and Garrod, 2004). For interactive repair of the misalignment to occur, A and B must agree on the node for which clarification is requested. The question is: how does B signal to A where to start? Here is where repetition and intonation jointly determine (re-)positioning of the pointer for both parties.

Figure 2 displays the update involved in A's fragment reply by licensing empty modification of her own initially established tree. On the tree under construction, the Q feature remains at the point of retrieval of the word *Bill*, but will be removed with identification of m_{21} as the value, hence falling within the subsumption constraint as defined. B, then, given the update provided by parsing A's FR, this time applies *Substitution* using the context provided (possibly by a more explicit utterance on A's part), and recovers $Leave'(m_{21,Bill'(m21)})$. The result, if so, is that A and B have re-aligned, and whatever failure in communication there had been in A's first utterance is successfully re-aligned.

On this account, we would expect that FR's can be made on the basis of a number of different assumptions. A may merely repeat the word used relative to her own context as in Fig.2. She may, however repeat the word *Bill* relative to a re-start,

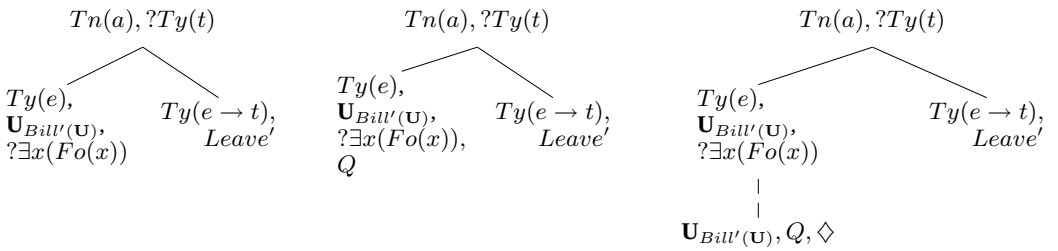

(a) **B speaking**; context tree (left), goal tree (centre), tree under construction (right)

Figure 1: Clarification Query: Result of B's CR

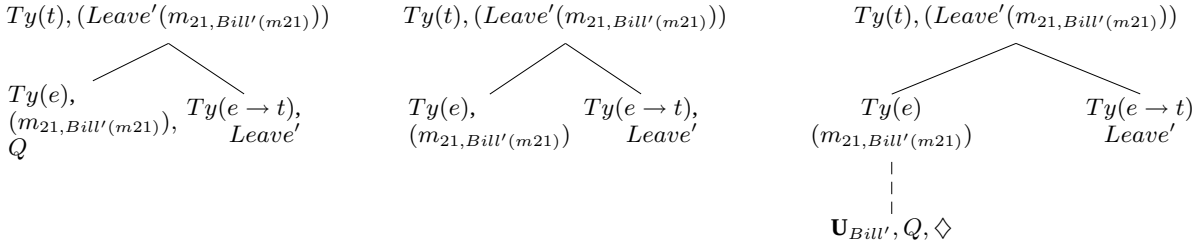

(a) **A speaking**; context tree (left), goal tree (centre), tree under construction (right)

Figure 2: Clarification Response: Result of A's FR

introducing an unfixed node, then re-using actions from her own context as appropriate to yield a possibly different tree. This is in any case needed in cases of mistaken construal. In order to understand such a case, in which B utters say "Jill", A will have to parse that name as providing an utterance whose interpretation has to be constructed independently: to the contrary, merely to add decorations to her own tree as context would lead to it being discarded as inconsistent, thus preventing her from assigning B's fragment an interpretation. But with *Adjunction available, A can build an interpretation for Bill's utterance from a new goal of $?Ty(t)$ straightforwardly, taking it to provide a metavariable decorating an unfixed node, and from there A can nonetheless select a subset of actions to yield an understanding of Bill's clarification request based on the context provided by her own utterance. Her own reply might well thus also involve such a re-start introducing an unfixed node by *Adjunction following exactly the same pattern of actions as established by the immediately previous parse sequence used in processing the utterance of *Jill*. In such a case, with her utterance of *No* indicating her rejection of that established proposition as part of her own context, re-start is indeed a putative option, since she can use it nevertheless to build an unfixed node but also thereafter to recover the very same actions used in the processing of his utterance. However, given her rejection of the tree constructed from the parse of B's CR, as indicated by her utterance of *No*, she might also simply revert to using her own utterance as context with trivial update as in Fig.2. Either option is possible, clearly licensed by the DS perspective of setting out alternative strategies.

5 Discussion

As these displays have indicated, CR and FR generation can be made relative to the immediate parse context, which may be any partial tree along the transition from initial so-called Axiom state to some completed tree. Furthermore, the assumption, as here, that generation of FRs can (but need not) be on the basis of trivial modifications of some complete tree provides a basis for explaining why even young children can answer clarificatory questions without making any hypothesis as to the basis for clarification other than identifying the node in question.[6]

The added significance of this incremental approach to CR, is that no difference in principle needs to be stipulated to distinguish constituent and clausal types of CR/FR. Even the type of which Purver et al call *a reprise gap* falls into the same type of explanation, and is, on this ac-

[6] In principle the account extends to predicate words, if we make assumptions analogous to those made here for linguistic names, but this assumption needs extended justification to be developed elsewhere.

count, no more than the mechanism one would need in cases where the individual speaker repeats the word as in A's third utterance in (3) (Healey et al., 2003):

(3) A: Could I have some toast please?
B: some....?
A Toast.

All that needs to be assumed is that in order for B to utter "Some....", B will have already had to have parsed A's previous utterance via the construction of some metavariable of cn type as formula value. On this scenario, B will not however succeed in fully understanding what A has said without establishing the value for that metavariable. One way of getting help with this is to initiate the term-construction process again, harmlessly over-riding the earlier specification of $\lambda P.\epsilon.P$, but then signalling the leaving of the pointer at the $?Ty(cn)$ node, which A then provides. All that is needed to model this process is the assumption of a meta-variable for any type licensed, and the assumption that repeat actions may trivially update the determiner node (see Cann et al., 2005).

The analysis of CR's and FR's is thus general: for all apparently distinct subtypes, there is simply a cline of possible partial trees from outset parse state to completed tree, any one of which can constitute a point for clarification by generation of the appropriate word, with the goal of providing some minimal update to that interrupted parse sequence in order, once clarification is provided, to be able to proceed. This account has three advantages. First, the characterisation of the lexical content of the fragment remains constant across all construals of its uses, both fragmentary and non-fragmentary. Second, the phenomenon is explained in an integrated way across both CR and FR fragments. But, more generally than this, the mechanisms posited for this account of CR/FR fragments are none other than those posited for the account of ellipsis in general. Fragments in language are those cases in which their construal can be provided directly from the context, whether by taking whatever partial structure that context provides and building on it, or by taking formulae established in context, or by taking a sequence of actions recorded in context. Clarificatory fragments are those where both input and output to the local parsing/production process may be a partial structure. The only constraint put on such a process is that use of context in language construal has to be licensed by the form of input: and in the case of clarificatory fragments, it is precisely such a license, which intonation provides, indicating both the need to use context for construal and the fact that such construal will be essentially local, particular to the sequence of expressions so picked out.

References

R Cann, R Kempson, and L Marten. 2005. *The Dynamics of Language: An Introduction*. Elsevier.

R Cann, R Kempson, and M Purver. 2007. Context and wellformedness: the dynamics of ellipsis. *Research on Language and Computation*, 5.

R Elugardo and R Stainton, editors. 2005. *Ellipsis and Non-Sentential Speech*. Springer.

R Fernández-Rovira. 2006. *Non-Sentential Utterances in Dialogue: Classification, Resolution and Use*. Ph.d., Kings College London.

J Ginzburg and R Cooper. 2004. Clarification, ellipsis, and the nature of contextual updates. *Linguistics and Philosophy*, 27:297–365.

J Ginzburg and I Sag. 2000. *Interrogative Investigations*. CSLI Publications.

P Healey, M Purver, J King, J Ginzburg, and G Mills. 2003. Experimenting with clarification in dialogue. In *Proceedings of the 25th Annual Meeting of the Cognitive Science Society (CogSci 2003)*.

R Kempson, W Meyer-Viol, and D Gabbay. 2001. *Dynamic Syntax: The Flow of Language Understanding*. Blackwell.

M J Pickering and S Garrod. 2004. Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27:169–226.

M Purver, R Cann, and R Kempson. 2006. Grammars as parsers: Meeting the dialogue challenge. *Research on Language and Computation*, 4(2–3):289–326.

M Purver. 2004. *The Theory and Use of Clarification Requests in Dialogue*. Ph.d., Kings College London.

M Purver. 2006. Clarie: Handling clarification requests in a dialogue system. *Research on Language and Computation*, 4(2–3):259–288.

K J Rodriguez and D Schlangen. 2004. Form, intonation and function of clarification requests in german task oriented spoken dialogues. In *Proceedings of Catalog '04 (The 8th Workshop on the Semantics and Pragmatics of Dialogue, SemDial04)*.