Updates, Files, and Focus-Ground[1]

Enric Vallduví
University of Edinburgh

Abstract
Accounts of information packaging (focus-ground, theme-rheme, topic-comment) in the pragmatic tradition are inherently dynamic in that they are concerned with the linguistic processes by means of which information is presented to an updating agent in discourse. The main insight gained from this line of research is that, in discourse, speakers not only present information to their interlocutors, but also provide them with detailed 'instructions' on how to manipulate and integrate this information. The use of these instructions reveals that speakers treat information states as highly structured objects and exploit their structure to make information update more efficient for their hearers. The analysis of information packaging acts as a valuable probe into the nature of linguistic information update and the structure of information states.

1 INFORMATION PACKAGING

Information packaging (a.k.a. communicative dimension, psychological structure) is a structuring of sentences by syntactic, prosodic, or morphological means that arises from the need to meet the communicative demands of a particular context. In particular, information packaging indicates how information conveyed by linguistic means fits into a (hearer's mental model of the) context or discourse. When communicating a proposition ϕ, speakers may realize it by means of different sentential structures according to their beliefs about the hearer's knowledge and attentional state with respect to ϕ. As a first approximation to information packaging in natural language interpretation, consider (1):

(1) a. He hates chocolate.
b. Chocolate he hates.
c. Chocolate he loves.

In (1), (a) and (b) are truth-conditionally equivalent. They differ not in what they say about the world, but in how they say what they say about the world, i.e., they differ in the way they are packaged. Compare now (b) and (c): they differ in their truth conditions, but exude a certain interpretive equivalence, which is a result of the fact that they are packaged in the same way. In other words, (a) and (b) differ in what they say about the world, but not in how they say it. In every language there is an array of sentences which, like (1a) and (1b), differ only in the way they are packaged. However, these alternatives cannot be used interchangeably in context. This can be empirically confirmed in terms of discourse felicity. For instance, while (1a) is a felicitous answer to the question *What does he hate?*, (1b) is not.

The partition of sentences into a focus-ground (also known as focus-topic, rheme-theme, new-given) plays a central role in information packaging (see Kuno 1972, Halliday 1985, Prince 1986, Gundel 1988, Ward 1988, *inter alia*). There is a wealth of characterizations of focus-ground, but they all share one characteristic: focus-ground divides the sentence into a part that anchors the sentence to the previous discourse or the hearer's

'mental world' and an informative part that makes some contribution to the discourse or the hearer's 'mental world'. As suggested by Clark & Haviland (1977:5), the point of such a partition is to optimize the communicative process. Vallduví 1992 contains a partial formalization of information packaging which takes up on these traditional ideas on focus-ground with the aim of establishing a first step towards an implementation in a dialogue-modeling system. It is argued that, in order to optimize the communicative process, each sentence encodes an information-packaging INSTRUCTION. Each instruction-type—there are four of them—is designed to indicate what part of the sentence constitutes INFORMATION and where and how that information fits in the hearer's model of the common ground.

From this perspective, the two notions of focus and ground were defined as follows. Let ϕ_s be the proposition conveyed by a sentence S and K_h (the relevant subset of) the hearer's model of the common ground at the time of utterance (t_u):

- FOCUS: the part of S that encodes *information* (I_s), i.e. the only augmentation or modification to the hearer's model of the common ground ($I_s = \phi_s$ - K_h).
- GROUND: the part of S that encodes what is already established in K_h at t_u; ushers I_s to the right location (from the speaker's viewpoint) in K_h; further subdivided into LINK and TAIL.

If the ground is simply already established, a question that arises is why sentences have a ground at all. The answer lies precisely in the ushering role it performs. Within the general ushering role of the ground, link and tail perform different specific tasks: links indicate *where* I_s goes and the tail indicates *how* it fits there. Sentences have a ground only if I_s requires ushering to a particular location in K_h. If I_s can be appropriately added to K_h without ushering, i.e., if specification of a location in K_h is inherited from a previous utterance, a sentence may be groundless or it may lack one of the two elements of the ground. We return to the notion of location in K_h in Section 3.

Focus, link, and tail combine to yield the four instruction-types: link-focus, link-focus-tail, all-focus, and focus-tail (no constituency implied). The structural mechanisms exploited by information packaging vary from language to language. In English, the focus-ground partition is realized mostly via prosody. Roughly, foci are associated with a H* pitch accent and links are associated with an L+H* pitch accent. In other languages, the focus-ground partition has necessary syntactic or morphological effects. We return to instructions in Sections 4 and 5.

2 UPDATE SEMANTICS

Dynamic approaches to meaning capitalize on the idea that the semantic contribution of sentences lies in their potential to change the context. A sentence is seen as a function from an input information state to an output information state, i.e. as an UPDATE. In Veltman's Update Semantics US, 'you know the meaning of a sentence if you know the *change* it brings about in the information state of anyone who wants to incorporate the piece of news conveyed by it' (Veltman 1990:29).

Following Dekker 1993, an information state s is a subset of the set of possible worlds, W, akin to the context set in Stalnaker 1978. The minimal information state is W. If $s = W$, no worlds have been ruled out yet. Updating with any given sentence in discourse consists in eliminating from s all those worlds that are incompatible with the interpretation of that sentence. Even though Stalnaker 1978 defines an information state

(actually, his context set) as the set of possible worlds compatible with what is accepted to be true by *both* speaker and hearer at a given time-point, US and its offsprings are entirely recipient oriented (see Dekker 1993:38). What is updated is not an information state independent of or shared by both speaker and hearer, but rather an information state in the hearer's mental model of the common ground, as suggested in Veltman's quote above. Intuitively, this is a desirable move. Talking about the information states of the updating agent makes more sense from a Gricean perspective.

In relativizing information states to the hearer, US sets foot in the traditional domain of information packaging, which deals with how information may best be put forth given (the speaker's assumptions about) the hearer's information state at t_u. In fact, the notion of I_s introduced in Section 1 is akin to the notion of update potential: I_s is the only part of S that brings about a change in a hearer's input information state s_1 (= K_h). The notion of focus as what is 'new' or 'informative' or what 'pushes the communication forward' may be now thought of as an update function (what takes an input context to an output context). Of course, a difference between I_s/focus in any of these pragmatic approaches and the notion of update potential in US is that I_s can be propositional or subpropositional (depending on whether S has a ground or not), whereas in US the update potential of a sentence is its propositional content ϕ. Therefore, if integration is to be pursued, some compromises will have to be made.

3 THE STRUCTURE OF K_h

Talking about ushering I_s to a location in the hearer's model of the common ground K_h, as in Section 1, does not make sense unless one assumes some sort of rich internal structure for K_h. The assumption that K_h has such a rich internal structure is commonplace in the pragmatic literature concerned with how incoming information meshes with what the hearer already knows or believes and is attending to. In contrast, Update Semantics takes information states, which are also hearers' models of the common ground, to be sets of possible worlds without any further internal structure. This is another point of departure between information packaging and US.

There are other dynamic frameworks, however, that do provide a richer structure in their representation of hearers' models of the common ground or information states, e.g. Discourse Representation Theory (DRT), Situation Theory, and Heim's (1982) File Change Semantics (FCS), which views infromation states as file-like structures. In comparing her files to Stalnaker's common grounds, Heim (1982:288) points out that her files 'do not just represent common grounds, but add to them some kind of internal structure that plays no role in their evaluation w.r.t. truth and falsity' It is precisely this internal structure which is crucially exploited by the different information-packaging strategies used by speakers in communication.

Let us view information states or hearers' mental models of the common ground as a FILE, i.e., let us refer to K_h or to s_1 as F_1. In FCS, a proposition ϕ_s in discourse acts as a function from an input file F_1 to an output file F_2 and truth values are computed over entire files. Files are collections of FILE CARDS. Each file card has a number of RECORDS written on it listing attributes and relations about the entity it denotes.

File growth, i.e. the transition from an input file to an output file, may be seen as consisting of two distinct simultaneous processes: file-card management and content update. File-card management is an entity-level process which concerns the novelty/familiarity status of file cards vis-à-vis the discourse universe. File-card management is responsible

for the creation of novel file cards and the activation of already familiar but dormant file cards. There is a strong but imperfect correlation between familiarity and formal definiteness and between novelty and formal indefiniteness. A subset of the familiar file cards are maintained in activation—they are salient and can be expressed by pronominal means. Actual content update, on the other hand, is effected when the information conveyed by a given sentence is incorporated to these novel and familiar file cards in the form of a record or condition. Information packaging reflects the speaker's assumptions about the hearer's knowledge and attentional state with respect to content update, i.e. it is a propositional-level phenomenon. File-card management, in contrast, reflects the speaker's assumptions about the hearer's knowledge and attentional state with respect to the set of discourse referents.[2]

Consider these two distinct aspects of file growth in discourse (2) as told to a hearer H. F_1 is H's (nonempty) initial information state and F_5 is H's final information state.[3]

(2) a. **Pat** told me a weird story.
b. She saw this man carrying a big banana.
c. Well, **the guy** started munching on it,
d. and, lo and behold, he turned into an orangutan.

(3) $F_1 \overset{a}{\mapsto} F_2 \overset{b}{\mapsto} F_3 \overset{c}{\mapsto} F_4 \overset{d}{\mapsto} F_5$

Here file-card management involves the following steps. F_2, the first sentence's output file, contains one file card, for the weird story, which was not present in F_1. The use of the indefinite *a weird story* triggers its creation and addition to F_1, where it remains salient. In addition, the file card for Pat, which (according to speaker's assumptions) was present but dormant in F_1, has been activated by the proper name *Pat* and is now salient. In F_3 two new file cards have been created, corresponding to the two indefinite NPs *this man* and *a big banana*, and the activation of Pat's file card is maintained via the use of the pronoun *she*. In F_4 uses the pronoun-like epithet *the guy* and the pronoun *it* to refer to the salient file cards created in the previous sentence.

Content update is carried out when the information conveyed by a sentence is recorded on the file cards created or activated during file construction. For instance, content update in the change from F_2 to F_3 in (2) entails adding a record on the file cards for Pat, the man, and the big banana describing the information that Pat saw the man carrying the big banana (this will be modified in Section 3). F_5 results from both file-card management and content update after the hearer has processed the entire discourse in (2). The relevant subset of F_5 could be represented as in Figure 1 (ignoring (a) and, therefore, the speaker and the story):

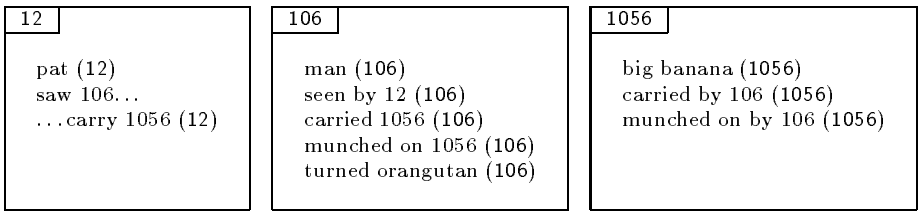

Figure 1: Partial view of F_5

How does information packaging contribute to the optimization of content update? The contribution is actually twofold. Instructions (a) avoid unnecessary multiple recording

of information by designating a unique location for content update and (*b*) correctly identify I_s and establish how it relates to data already present in K_h. Here we will discuss the first contribution and Section 4 will discuss the second contribution.

It was noted in the discussion around (2) that in content update information is recorded several times: once for each one of the file cards expressed in a given utterance. Multiple update, at first blush, appears to be necessary to guarantee that all data about a given file card is accessible when evoking it in subsequent discourse. Fortunately, the same end-result can be attained in an alternative, much more efficient way. Data that has been recorded on one single location in a database can be, in fact, accessed from many different places provided they are connected hypercard-style to the single entry location by means of some symbolic link. Consider (2) again. F_5, H's final output information state, may be represented as in Figure 2, rather than as in Figure 1. In Figure 2, content update has taken place only once per sentence on a designated file card chosen among the file cards expressed in the sentence. The other file cards are linked to this designated file card by $\rightsquigarrow$:

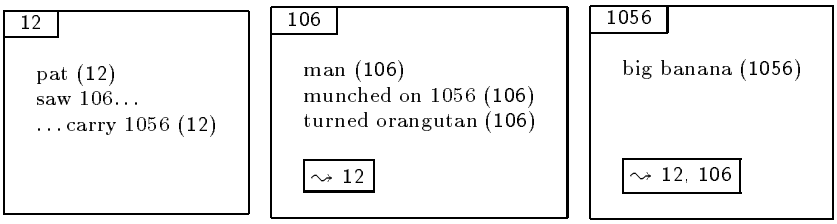

Figure 2: Hypercard-style partial view of F_5

Information-packaging instructions circumvent multiple update by designating one of the file cards expressed in a sentence as the LOCUS OF UPDATE. Within the instruction, this particular task is carried out by the link. The link ushers I_s to a file card *fc* in F_1, where, so to speak, it will be able to 'discharge' its update potential. In Figure 2, given the particular prosodic realization in these examples, *Pat* is taken to be the link in (2a) and *the guy* the link in (2c), i.e., file cards 12 and 106 have been designated loci of update (although in other contexts other file card could have also been chosen). Proposals similar to this are Reinhart (1982:24), where it is suggested that links (for her, topics) are a signal to determine under which entries to classify a new proposition in a Stalnakerian context set, and Kuno's (1972) notion of SORT KEY, since designating *fc* as a locus of update for I_s is indeed similar to establishing *fc* as a sorting key for I_s.

4 FOCUS AS UPDATE POTENTIAL

The second contribution of information packaging to the optimization of content update concerns the identification of I_s. As noted in Section 1, I_s is the only augmentation or modification ϕ_s makes to K_h at t_u. In update-semantic terms I_s is the update potential of ϕ_s. However, the update potential of ϕ_s is not necessarily ϕ_s: two propositionally equivalent sentences will differ in their update potential if the value of K_h varies, e.g. if they are communicated to hearers with different input information states.

Consider dialogues (4) and (5). S_0 is a presidential aid, H_1 a newly-appointed White House butler, and H_2 the Foreign Secretary after returning from a trip to Europe:

(4) a. H_1: I'm arranging things for the president's dinner. Anything I should know?

b. S_0: Yes. The **president** [$_F$ hates the Delft CHINA SET]. Don't use it.

(5) a. H_2: In the Netherlands I got the president a big Delft china tray that matches the set he has in the living room. Was that a good idea?

b. S_0: Nope. The **president** [$_F$ HATES] the Delft china set.

Both (4b) and (5b) express the same propositional content, namely that the president hates the Delft china set (abbreviated as PHD). This is reflected in the fact that (4b) and (5b) yield identical output files. In other words, both H_1 and H_2 are in the same output information state s_2. $F_2(H_1)$ and $F_2(H_2)$ exclude the same worlds: any worlds in which PHD is false. $F_2(H_1)$ and $F_2(H_2)$ can be both (partially) represented as in Figure 3:[4]

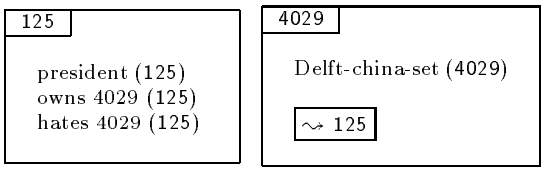

Figure 3: H_1's and H_2's (partial) output files after (4b) and (5b), respectively

Nevertheless, the update potential of (4b) and (5b) is *not* the same. This is because the input file F_1 (the value of K_h) is not the same in each scenario. From both contexts it can be inferred that both H_1 and H_2 know about the president, the Delft china set, and about the former owning the latter. However, in (4) S_0, the presidential aid, has no reason to assume that the hearer, H_1, has any beliefs about (and is attending to) the president's attitude towards the Delft china set. H_1's question, *Anything I should know?*, could have been given a number of equally relevant, felicitous answers: that the president does not like fish, that the president always eats at nine o'clock, that he has high cholesterol, that he eats in the Oval Office, and so on. In contrast, in (5) S_0 is warranted to assume, given what she has heard in the immediately previous dialogue, that H_2 believes (and is attending to it) that the president has some attitude towards the Delft china set (perhaps without knowing which one it is). Basically, and using Jackendoff's (1972) words, the president's having some attitude towards the Delft china set is 'under discussion' at t_u in context (5) but not in context (4). This means that $F_1(H_1)$ at the time (4b) is uttered contains less information than $F_1(H_2)$ at the time (5b) is uttered (in update-semantic terms, $F_1(H_2)$ is a proper subset of $F_1(H_1)$). The difference in informativeness or update potential between (4b) and (5b) is determined by the contents of the input files that they can felicitously augment.[5]

The different (partial) input files of the two hearers H_1 and H_2 before processing (4b) and (5b), respectively, can be represented as in Figures 4 and 5 (again, identical indices assigned in both files for the sake of exposition):

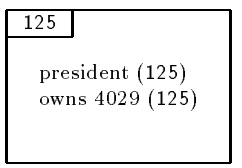

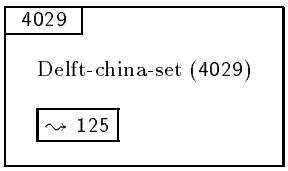

. Figure 4: H_1's (partial) F_1 in (4)

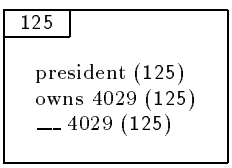

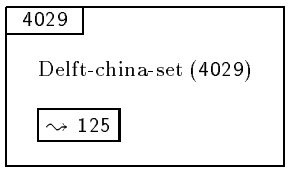

Figure 5: H_2's (partial) F_1 in (5)

The underspecified record on file card 125 in Figure 5 represents the fact that, from S_0's viewpoint, the president's having an attitude towards the Delft china set is believed and attended to by H_2. The blank (__) in that record notates the lack of information about the nature of this attitude. The underspecified record is absent from file card 125 in

Figure 4, since in that context (the speaker assumes) the president's having an attitude towards the Delft china set is not in H_1's input file. Comparing Figures 3, 4, and 5 it is easy to gauge that PHD is more informative for H_1 than it is for H_2.[6]

5 INSTRUCTIONS

The two contexts (4) and (5) push S_0 to makes different assumptions about the information states of H_1 and H_2. This motivates her to package her utterances differently so as to correctly identify, in each case, the update potential of ϕ_s and to indicate how it fits into the input file. In (5b), for instance, where the president's having some attitude towards the Delft china set is treated as being believed and attended to by the hearer, a structure is chosen that singles out the verb *hates* as the only informative part of the sentence. In (4b), where no such assumption is warranted, the entire verb phrase is focal. Also, the role of the link within this ushering function was briefly described in Section 3. But in addition to the link, one needs to distinguish between two basic ways in which update of F_1 may proceed. Let us call them MODES OF UPDATE. Information-packaging instructions, then, perform the following tasks: (*a*) they identify the update potential I_s, (*b*) they may designate a locus of update, and (*c*) they specify a mode of update.

Let us describe the exact interpretation of information-packaging instructions. Each information-packaging primitive (focus, link, and tail) contributes its constant informational interpretation to the combinatoric rules yielding the interpretation associated with each instruction-type. The focal segment is interpreted as being the argument of an UPDATE predicate. The core informational interpretation of any sentence is UPDATE(I_s), i.e. an instruction to update the input file with I_s. This bare all-focus structure is enriched if any ground elements are present. The link was described above as being the designator of a locus of update. When an expression denoting a file card *fc* is structurally encoded as a link it is informationally interpreted as GOTO(*fc*), i.e. it ushers I_s to a particular file card *fc*, where the update must be effected. If the locus of information update for S_n is inherited from S_{n-1}, no link is required. As noted in Section 1, links are associated with a particular structural realization in every language.

Tails designate a given record R already listed on *fc*. In tailless instructions, I_s updates F_1 by simply being added to it. But if a tail is present, I_s updates F_1 by completing or altering R. In fact, the presence of a tail triggers a switch in the mode of update. The UPDATE predicate has two distinct manifestations. One, the one-place UPDATE-ADD mode, acts as a default and, as said, is associated with tailless instructions (link-focus and all-focus). The other, the two-place UPDATE-REPLACE mode is triggered by the presence of a tail (link-focus-tail and focus-tail instructions). The UPDATE-REPLACE mode indicates that I_s is not merely added to a file card but rather must complete or alter a record R already present on that file card.

Sentence (4b) is an instance of the UPDATE-ADD mode (tailless instruction) and (5b) is an example of the UPDATE-REPLACE mode (cum-tail instruction). Both are link-containing instructions. The interpretation of the instruction-types they instantiate is as in (6a) for (4b) and (7a) for (5b) (the Delft china set is not represented as an index in the instruction for the sake of clarity):

(6) a. [$_L$ The **president**] [$_F$ hates the Delft CHINA SET]

b. GOTO(**125**)(UPDATE-ADD(hates the Delft-china-set(**125**))

c. 'Go to file card **125** in your input file and update its content by adding the record "hates the Delft-china-set(**125**)".'

(7) a. [L The **president**] [F HATES] the Delft china set.

b. GOTO(125)(UPDATE-REPLACE(hates , {__ | __ Delft-china-set(125)}))

c. 'Go to file card 125 in your input file and update its content by replacing 'hates' for __ in the preexisting record "__ the Delft-china-set(125)".'

The paraphrases provide an intuitive idea of how the distinction between plain addition to *fc* and modification of a record R on *fc* is implemented. In the case of (7a), the record R is underspecified, so UPDATE-REPLACE is in fact a completion of R with I_s. As noted, however, R can also be a fully specified record, in which case I_s alters it.

There are two other possible instruction-types. They are UPDATE-ADD and the UPDATE-REPLACE modes of update, analogous to (6b) and (7b), except they have no link. No locus of update need be designated because the current one is already established in and inherited from the previous discourse. Examples (8) and (9) correspond to a linkless all-focus instruction and a linkless focus-tail instruction, respectively:

(8) a. H_1: I'm arranging things for the president's dinner. Anything I should know?

b. S_0: Yes. The **president** always uses plastic dishes.
.[F (He) hates the Delft CHINA SET].

c. UPDATE-ADD(hates the Delft-china-set(125))

d. 'At current locus of update, add record "hates the Delft-china-set(125)".'

(9) a. H_2: In the Netherlands I got the president a big Delft china tray that matches the set he has in the living room. Will the president like it?

b. S_0: Nope. [F (He) HATES] the Delft china set.

c. UPDATE-REPLACE(hates , {__ | __ Delft-china-set(125)})

d. 'At current locus of update, replace 'hates' for __ in the preexisting record "__ the Delft-china-set(125)".'

The second sentence in (8b), let us call it S_n, is, as noted, an example of an all-focus instruction. The first sentence in (8b), S_{n-1}, is a link-focus instruction, which contains the link *the president* (realized with the L+H* prosodic structure characteristic of links). In S_{n-1} file card 125 is designated as the current locus of update. S_n contains no link (there is no phrase realized with link prosody), which indicates that content update is to be carried out on the file card inherited from S_{n-1}, i.e. file card 125.[7]

Summarizing, we have two modes of update, UPDATE-ADD and UPDATE-REPLACE, and a mechanism to designate a file card as the locus of update. Designation of a locus of update is not always necessary, since it may be inherited from previous discourse. This result in four possible instruction-types: the linkless all-focus (UPDATE-ADD) and focus-tail (UPDATE-REPLACE) and the link-contaning link-focus (GOTO,UPDATE-ADD) and link-focus-tail (GOTO,UPDATE-REPLACE).

6 CONCLUSION

Many analysts have noted that, in discourse, information is not presented in an unstructured way. Rather, linguistic structure provides hearers with detailed instructions about how to retrieve, sort, and file information. This, of course, makes sense only if it is assumed that information states are highly structured objects that allow—or even require—information to come with (un)packaging instructions. This paper discusses the sorts of instructions found in communication and suggests a particular internal structure

for information states that seems to accord with the nature of these instructions. The particular formulation of information-packaging endorsed in this paper favors a view of information states as collections of records or conditions that can be individually accessed or highlighted, rather than as being merely a set of possible worlds. It also favors a view that includes, as an integral part of information states, file-card-like constructs that act as a sort key for subsequent information in communication. Finally, let us notice that this approach is coached in a constructive view of information update, rather than an eliminative one, although any consequences have been left unexplored.

The main point of this paper is to argue that a proper understanding of information packaging, i.e. of the actual strategies used by human agents in effecting information update by linguistic means, will help us gain further insight into the structural properties of the cognitive states these dynamic strategies manipulate. The focus-ground articulation of the sentence is the linguistic means by which agents carry out these diverse updating strategies. As such, the notions of focus and ground (or theme and rheme, etc.) should play a central role in any dynamic approach to meaning. To conclude with a slogan: what you get matters, but how you get it matters too.

NOTES

[1]I am indebted to D. Beaver, R. Cooper, P. Dekker, E. Engdahl, J. Ginzburg, H. Hendriks, R. Ladd, D. Milward, M. Poesio, and M. Rooth, among others, for useful comments and suggestions. Thanks to R. Cooper for granting me permission to use his EKN LaTeX macros. This work is carried out within ESPRIT project DYANA-2 (Basic Research Project 6852).

[2]Issues of global and local focus in Sidner 1981 (a different use of the term focus), ranking of utterance centers in Centering Theory in Grosz, Joshi & Weinstein 1983, referentials status in Prince 1992, and cognitive status in Gundel, Hedberg & Zacharski 1992 are issues of file-card management. Heim 1983 aims to capture familiarity and novelty from a FCS perspective.

[3]Here and in examples below, where judged helpful, links are enclosed in L-labeled brackets and the L+H*-accented element within the link is in boldface. Foci are enclosed in F-labeled brackets and the H*-accented item within the focus appears in small caps.

[4]Identical indices are assigned to both $F_2(H_1)$ and $F_2(H_2)$ for the sake of exposition. Also, since *the president* is realized with a L+H* accent, the file card for the president is taken to be the designated locus of update.

[5]In the literature one often finds that the contexts used to illustrate the felicity of different information-packaging structures is less sophisticated than the ones used in (4) and (5). For instance, (4b) would be given as the answer to a question like *What can you tell me about the president (that is relevant for my stated goal)?* and (5b) as the answer to *How does the president feel about the Delft china set?*.

[6]The appropriate formal characterization of an underspecified record (e.g. lambda abstract, situation-theoretic question) will be left as an open issue. A further point is that in a slightly different contextualization for (5), a second alternative interpretation is available to S_0. Instead of assuming that H_2 has no belief about the president having an attitude towards the Delft china set, S_0 may assume that he does, but that his knowledge or belief is erroneous (i.e. love instead of hatred). In such a scenario, we would not be dealing with an underspecified record containing a blank, but with a fully specified record.

[7]Weak pronouns are taken to be inert as far as information packaging is concerned (but crucial for file-card management). In (8) and (9) they remain vacuously within the focus as place-holders (due to independent requirements of English grammar; they would be null in other languages). Strong pronouns are full-fledged items and may take part in an instruction as foci, links, or tails.

REFERENCES

Clark, H.H. & S.E. Haviland (1977), 'Comprehension and the given-new contract', in R.O. Freedle (ed.), *Discourse production and comprehension*, 1-40, Erlbaum, Hillsdale, NJ.

Dekker, P. (1993), 'Updates in dynamic semantics', in R. Cooper (ed.), *Integrating semantic theories* (*DYANA-2 Report R2.1.A*), 1-44, ILLC, Amsterdam.

Grosz, B., A.K. Joshi & S. Weinstein (1983), 'Providing a unified account of definite noun phrases in discourse', *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, **21**, 44-50.

Gundel, J.K. (1988), *The role of topic and comment in linguistic theory*, Garland, New York.

Gundel, J.K., N. Hedberg & R. Zacharski (1993), 'Cognitive status and the form of referring expressions', *Language*, **69**, 274-307.

Halliday, M.A.K. (1985), *An introduction to functional grammar*, Arnold, London.

Heim, I. (1982), 'The semantics of definite and indefinite noun phrases', Ph.D. thesis, University of Massachusetts.

Heim, I. (1983), 'File Change Semantics and the familiarity theory of definiteness', in R. Bäuerle et al. (eds.), *Meaning, use, and interpretation of language*, 164-189, de Gruyter, Berlin.

Jackendoff, R. (1972), *Semantic interpretation in generative grammar*, The MIT Press, Cambridge.

Kuno, S. (1972), 'Functional sentence perspective', *Linguistic Inquiry*, **3**, 269-320.

Prince, E.F. (1986), 'On the syntactic marking of presupposed open propositions', *Chicago Linguistic Society* (*Parasession*), **22**, 208-22.

Prince, E.F. (1992), 'The ZPG letter: Subjects, definiteness, and information-status', in S.A. Thompson & W.C. Mann (eds.), *Discourse description: Diverse analyses of a fund-raising text*, 295-325, John Benjamins, Amsterdam.

Reinhart, T. (1982), 'Pragmatics and linguistics: an analysis of sentence topics', *Philosophica*, **27**, 53-94.

Sidner, C.L. (1981), 'Focusing for interpretation of pronouns', *American Journal for Computational Linguistics*, **7**, 217-31.

Stalnaker, R. (1978), 'Assertion', in P. Cole (ed.), *Pragmatics* (*Syntax & Semantics* 9), 315-332, Academic Press, New York.

Vallduví, E. (1992), *The informational component*, Garland, New York.

Veltman, F. (1990), 'Defaults in update semantics', in H. Kamp (ed.), *Conditionals, defaults, and belief revision* (*DYANA-2 Report R2.5.A*), 28-63, CCS, Edinburgh.

Ward, G.L. (1988), *The semantics and pragmatics of preposing*, Garland, New York.

Enric Vallduví
Centre for Cognitive Science
University of Edinburgh
2 Buccleuch Place
Edinburgh EH8 9LW, Scotland
e-mail: enric@cogsci.ed.ac.uk