The dynamics of information packaging*

Enric Vallduví
Centre for Cognitive Science & Human Communication Research Centre
University of Edinburgh
enric@cogsci.ed.ac.uk

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.

Contents

*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 and M. Poesio for granting me permission to use their EKN and DRT LaTeX macros, respectively. This work is carried out within Tasks 1.3 and 2.1 of the ESPRIT-funded DYANA-2 project (Basic Research Project 6852).

1 Introduction

1.1 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 or discourse. 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, a speaker may realize it by means of different sentential structures according to his/her beliefs about the hearer's knowledge and attentional state with respect to ϕ. The term 'packaging' is due to Chafe (1976):

> I have been using the term *packaging* to refer to the kind of phenomena at issue here, with the idea that they have to do primarily with how the message is sent and only secondarily with the message itself, just as the packaging of toothpaste can affect sales in partial independence of the quality of the toothpaste inside. (Chafe 1976:28)

As a first approximation to the effect of packaging in natural language interpretation, consider the examples in (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), which differ in their truth conditions (in fact, given a fixed indexical value for the pronoun *he*, they are mutually exclusive). Nevertheless, (b) and (c) 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 the proposal presented below, each one of the ways in which propositional content is packaged is referred to as an instruction. Each sentence encodes a logical formula ϕ and an information-packaging instruction. Thus, (a) and (b) encode ϕ and (c) encodes ψ. Simultaneously, (a) encodes an instruction of, say, type A, and (b) and (c) encode an instruction of type B.

In every language there is an array of sentences which, like (1a) and (1b), differ only in the way the say what they say about the world. However, these alternatives cannot be used interchangeably in context. Use of particular alternatives is constrained by the current 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 infelicity of (1b) is due to the fact that its packaging is not appropriate in this context.

The partition of sentences into focus-ground (also known as focus-topic, rheme-theme, new-given, etc.) has been argued to play a central role in information packaging (see Hockett 1958, Kuno 1972, Halliday 1985, Erteschik-Shir 1986, Prince 1986, Rochemont 1986, Ward 1988, among others). 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'. The point of such a partition is to optimize the communicative process, as suggested by Clark & Haviland 1977:

> To ensure reasonably efficient communication, [...] [t]he speaker tries, to the best of his ability, to make the structure of his utterances congruent with his knowledge of the listener's mental world. (Clark & Haviland 1977:5)

In Vallduví 1992 I developed 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. I 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. These four instruction-types, in fact, encompass systematically all the informational constructions previously described in the literature from the topic-comment, ground-focus, and other perspectives.

From this perspective, the two information-packaging notions of focus and ground were defined as follows. Let ϕ_{s} be the proposition conveyed by a sentence S and $\mathrm{K_h}$ (the relevant subset of) the hearer's model of the common ground at the time of utterance ($\mathrm{t_u}$):

- FOCUS: the part of the sentence that encodes *information* ($\mathrm{I_s}$), i.e. the only augmentation or modification to be made to the hearer's model of the common ground ($\mathrm{I_s} = \phi_{\mathrm{s}}$ - $\mathrm{K_h}$).

- GROUND: the part of the sentence that encodes what is already established and under discussion in $\mathrm{K_h}$ at $\mathrm{t_u}$; ushers $\mathrm{I_s}$ to the right location (from the speaker's viewpoint) in $\mathrm{K_h}$; further subdivided into LINK and TAIL.

Truth-conditionally equivalent sentences with different focus-ground partitions differ in their informativeness (in the sense of information just introduced).

If the ground is simply already established, noninformative data, a question that arises from a dynamic perspective 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* $\mathrm{I_s}$ goes and the tail indicates *how* it fits there. Sentences have a ground only if $\mathrm{I_s}$ requires ushering to a particular location in $\mathrm{K_h}$. If $\mathrm{I_s}$ can be appropriately added to $\mathrm{K_h}$ without ushering, i.e., if specification of a location in $\mathrm{K_h}$ 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 $\mathrm{K_h}$ in Section 2).[1]

Focus, link, and tail combine to yield the four distinct instruction-types: link-focus, link-focus-tail, all-focus, and focus-tail (no constituency implied by this ordering). Each instruction-type is associated with a packaging interpretation which results from the interpretation of each one of the informational primitives of which it is composed. Knowledge of how to use and interpret information-packaging instructions is part of any speaker's linguistic

[1]There is a number of motivations for the subdivision of the ground into two elements, both empirical and conceptual, but one straightforward advantage is that, with this move, the two traditional binomial informational articulations of ground-focus and topic-comment are conflated into a single trinomial hierarchical one, given that links are more or less akin to the notion of topic in the topic-comment framework (cf. e.g. Gundel 1974, Reinhardt 1982).

competence. The structural mechanisms exploited by information packaging vary from language to language. In English, the focus-ground partition is realized mostly via prosody. Foci, for instance, are associated with a H* pitch accent and links are associated with an L+H* pitch accent.[2] In other languages, the focus-ground partition has necessary syntactic effects (e.g. links are necessarily left-detached and foci are associated with a particular syntactic position too).[3] We return to the interpretation of instructions in Section 3.

1.2 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, '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 1993b, an information state s is a subset of the set of possible worlds, W, akin to a Stalnakerian context set (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. As communication proceeds, more and more worlds are eliminated with each update, thus reducing uncertainty. Eliminativity is an important property of the updating mechanism. Updating is eliminative in that interpretation of a given formula ϕ can only eliminate possibilities. The output of an update, s_2, must be a subset of the input information state s_1.

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, Update Semantics and its offsprings are entirely recipient oriented (see Dekker 1993b: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. Beaver (1993b:77), for instance, defines 'an information state at some point in a discourse [as the] hearer's instantaneous snapshot of the common ground of the conversational participants'. Intuitively, this is a desirable move. Talking about the information states of the updating agent makes more sense from a dynamic perspective. After all, in communication speakers do not engage in self-update, since as reflected in the Gricean maxim of quality, they communicate only what they already (pretend to) know or believe.

In relativizing information states to the hearer, Update Semantics sets foot in precisely 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 the time of utterance. In fact, the notion of information ($\mathrm{I_s}$) introduced in Section 1.1 is akin to

[2]This is an somewhat idealized picture. Pitch accents that occur in the prenuclear field (to the left of the most prominent H* accent) may undergo a process of eradication. Also, the need to express prosodically other aspects of interpretation may override the default prosodic realization of foci or links, e.g. the L*+H accent associated with speaker uncertainty (Ward & Hirschberg 1985). Finally, in speaking of prosodic association with foci and links, perhaps the relevant intonational unit is not the accent but rather the phrase. There would be no such thing as a link-associated accent but rather a link-associated tune.

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

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 ($= K_h$). Traditional pragmatic analyses of information packaging can benefit from the new insights gained by dynamic approaches to meaning. In particular, the notion of focus as what is 'new' or 'informative' or 'contrary to expectation' 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 or focus in any of these pragmatic approaches and the notion of update potential in Update Semantics is that I_s can be propositional or subpropositional (depending on whether S has a ground or not), whereas in Update Semantics 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.

2 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 1.1 above, does not make much 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, since information packaging is concerned not only with the identification of the information conveyed by a sentence, but also with how that information fits 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 Update Semantics.

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) and Situation Theory. Among these is Heim's (1982, 1983) File Change Semantics (FCS), which views them as file-like structures. In talking about Stalnaker's common ground, which is defined as a bare set of possible worlds, Heim points out:

> From my point of view, [Stalnaker] thus restricts himself to one aspect of the common ground [...] From his point of view, on the other hand, my 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. (Heim 1982:288)

We agree with Heim that there has to be some additional internal structure in the hearer's model of the common ground that plays an important role in natural language interpretation, even if this internal structure is of tangential relevance to truth-value computation. It is this internal structure of information states which is, in fact, crucially exploited by the different information-packaging strategies used by speakers in pursuing communicative efficiency.

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 (much like in DRT, where ϕ_s acts as a function from an input DRS to an output DRS). The appropriateness of the truth-value computation mechanisms in FCS need not concern us here and no reference to them will be made from now on. Files are collections of FILE CARDS. These file cards are akin to the notion of discourse referent in Karttunen (1976) or to DRT-style discourse markers and mediate between referring expressions and real-world entities. Each file card has a number of RECORDS—analogous to conditions in a DRS—written on it listing descriptions (attributes and relations) about the entity it denotes. A file, however, differs from a DRS in the way in which the relationship between records/conditions and file-cards/discourse-markers is represented. A file is dimensionally richer than a DRS. Each file card introduces its own representational space where all records concerning that file card are to be found. In contrast, conditions in a DRS are represented in a space shared by all the markers in the discourse universe.

File growth, i.e. the transition from an input file to an output file, may be seen as consisting of two simultaneous processes that are in principle distinct: file-card management and content update. File-card management is an entity-level process which concerns the novelty/familiarity status of file cards or discourse markers vis-à-vis the discourse universe. Issues of global and local focus in Sidner 1981, cognitive status in Gundel, Hedberg & Zacharski 1992, and ranking of utterance centers in Centering Theory (Grosz, Joshi & Weinstein 1983) are issues of file-card management (Heim 1983 aims to capture familiarity and novelty from a FCS perspective). File-card management is responsible for the creation of novel file cards and the activation of already familiar but dormant file cards. As is noted in Heim 1983, 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. These can be expressed by pronominal means.[4] Actual content update, on the other hand, is effected when the information conveyed by a given sentence is incorporated on these novel and familiar file cards in the form of a record or condition. As noted, the raison d'être of information packaging is to optimize this process of content update. 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 (i.e. cognitive or referential status), in contrast, reflects the speaker's assumptions about the hearer's knowledge and attentional state with respect to the set of discourse referents. The need to distinguish between these two levels is convincingly argued for, among others, by Reinhart 1982 and Prince 1992b.[5]

Consider these two distinct aspects of file growth in discourse (2). A speaker S is telling the story to a hearer H. F_1 is H's initial information state (which is nevertheless nonempty, because, as noted by Heim, there is no such thing as a clean-slate start for any given linguistic interaction) and F_5 is H's final information state. File growth proceeds as in (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, sentence (a)'s output file, contains one file card, for the weird story, which was not present in F_1. The use of the indefinite *a*

[4] See Prince 1992a, Gundel, Hedberg, and Zacharsky 1993 for more on the structural effects of file-card management. Salient entities are also said to be *in focus*, especially in the computational and psycholinguistics literature. This use of 'focus' must not be confused with the use of focus in the focus-ground partition.

[5] It is interesting to note that Update Semantics (Veltman 1990) is concerned mostly with content update, whereas Groenendijk & Stokhof's (1991) Dynamic Predicate Logic is concerned partially with what I called file-card management (assignment of values to variables). Dekker 1993a puts the two together in his update-inspired MDPL, a proposal that, as Dekker himself points out, is close in spirit to Heim's FCS.

weird story triggers its creation and addition to F_1. 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 in activation. In F_3 two new file cards have been created, corresponding to the two indefinite noun phrases *this man* and *a big banana*, and the activation of Pat's file card is maintained via the use of the pronoun *she*. We need to assume that a file card that has just been created remains in activation (is salient) for a while. This is why in F_4 the speaker can use the pronoun-like epithet *the guy* and the pronoun *it* to refer to the file cards created in the previous sentence. In the transitions to F_4 and F_5 there is no creation of novel file cards or activation of dormant file cards (if we take *an orangutan* in (d) to be predicative), although some 'local focus' mechanism is at work to determine the anaphoric resolution of the pronominals forms.

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 statement will be modified in Section 2.1). Similarly for the content update from F_3 to F_4 and from F_4 to F_5. F_5 results from both file-card management and content update for the entire discourse in (2). The relevant subset of F_5, i.e. of the hearer's model of the common ground after processing (2), could be represented as in Figure 1 (ignoring the first sentence and, therefore, the file cards for the speaker and the story):

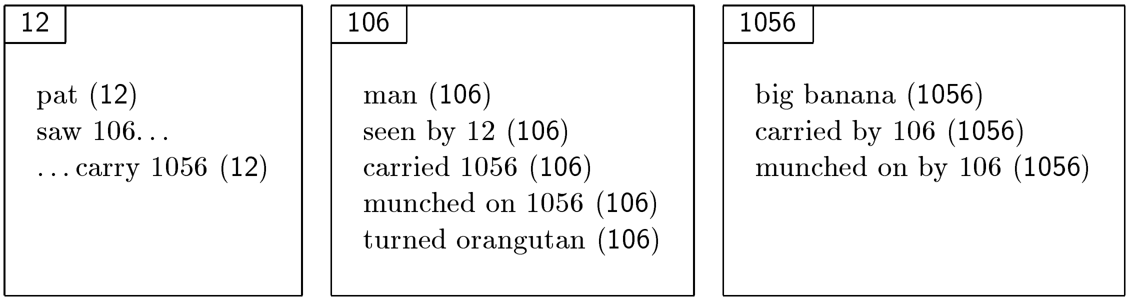

Figure 1: Partial view of F_5

File-card management guarantees that the correct file cards are available at each transition. Content update is responsible for the addition of records to these file cards.

2.1 Locus of update

Given this characterization of the structure of information states and the mechanics of file growth, how does information packaging contribute to the optimization of content update? The contribution of information packaging is actually twofold. Information-packaging instructions provide a means to (*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. This section discusses the first contribution and Section 3 discusses the second contribution.

It was just noted in the discussion around discourse (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 the subsequent discourse. For instance, after processing the third sentence in (2) we know about file card 106 that its denotation munched on the denotation of 1056 and about file card 1056 that its denotation was munched on by the denotation of 106. Fortunately, the same end-result can be attained in an alternative, much more efficient way. If we think about computer filing systems, it is easy to see that data that has been recorded on one single location in the database is, in fact, accessible from many different places provided they are connected hypercard-style to the single entry point by means of some symbolic-link mechanism. Consider discourse (2) again. F_5, H's final output information state, need not be represented as in Figure 1. Rather, it may also be optimally represented as in Figure 2. 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 $\rightsquigarrow$:

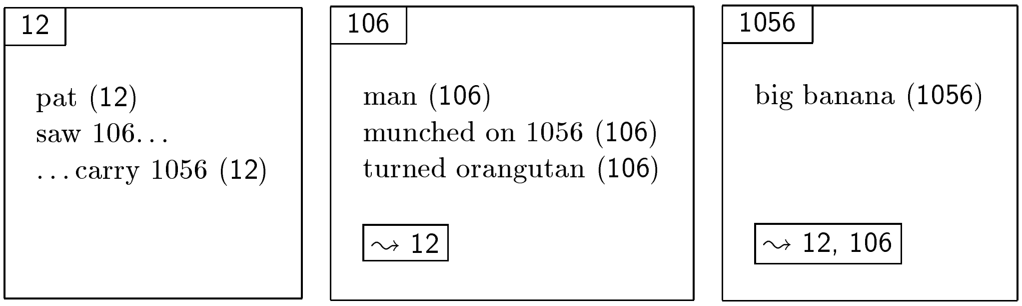

Figure 2: Hypercard-style partial view of F_5

Information-packaging instructions circumvent the redundancy that would arise from 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: I_s is 'ushered' to the file card in F_1 designated by the link as the locus of update, where, so to speak, it will be able to 'discharge' its update potential. In the case illustrated in Figure 2, it has been assumed, given the particular prosodic realization in these examples, that the link in (2a) is *Pat* and that *the guy* is the link in (2c), i.e., that file cards 12 and 106 have been designated as the loci of update (although in other contexts other file card could have also been chosen). A similar proposal is found in Reinhart (1982:24), where it is suggested that links (for her, topics in a topic-comment partition) are a signal to determine under which entries to classify a new proposition in a Stalnakerian context set. In fact, the notion of locus of update is also akin to the notion of SORT KEY in Kuno 1972: designating a given file card *fc* as a locus of update for I_s can indeed be thought of as establishing a key with respect to which I_s is sorted.

Even though designation of a locus of update involves, in some sense, an 'activation' of the designated file card, it is an integral part of content update rather than of file-card management. Designation as locus of update is not a reflection of familiarity/novelty. Although referential status and linkhood interact in an interesting way, loci of update can, in principle, have any referential status. Designation as locus of update must not be mistaken for activation in the file-card-management sense or for any other entity-level process. Rather, its role can only be conceived within the wider ushering role played by the ground in information-packaging instructions. The particular choice of locus of update is actually up to the speaker, although the choice is constrained by discourse structure, thematic cohesion, the familiarity status of a given file card, and rhetorical effect, among other things.

3 Information packaging and update potential

3.1 Focus as update potential

The second contribution of information packaging to the optimization of content update concerns the identification of what part of a sentence actually corresponds to I_s. As noted in Section 1.1, 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, unlike in Update Semantics, 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). Take S_0 to be a presidential aid, H_1 a newly-appointed White House butler, and H_2 the Foreign Office 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:[6]

125	4029
president (125) owns 4029 (125) hates 4029 (125)	Delft-china-set (4029) $\boxed{\leadsto 125}$

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, knows anything 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 knows or believes that the president has some attitude towards the

[6] 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.

Delft china set (perhaps without knowing which one it is) and is attending to this knowledge or belief. 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.[7]

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):

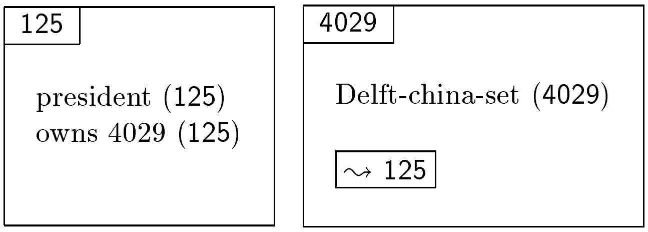

Figure 4: H_1's (partial) input file F_1 in (4)

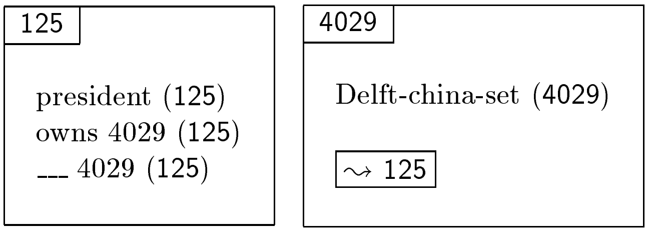

Figure 5: H_2's (partial) input file F_1 in (5)

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 information state.[8] 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, i.e., $\text{PHD}(F_1(H_1)) = (F_2(H_1))$ eliminates more possibilities than $\text{PHD}(F_1(H_2)) = (F_2(H_2))$.[9]

[7]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?*. The same points made with the contexts used in (4) and (5) can be made using these simpler context-setting questions, but the resulting dialogs do not feel natural (e.g. the most natural answer to *How does the president feel about the Delft china set?* is not a full-blown sentence, but rather a reduced *He hates it*).

[8]The appropriate formal characterization of an underspecified record will be left as an open issue. As early as Jackendoff (1972), grounds have been represented as lambda abstracts, so using a lambda abstract here is obviously a possibility. Alternatively, one could resort to a situation-theoretic object along the lines of questions (Ginzburg 1993).

[9]In a slightly different contextualization for (5), a second alternative interpretation is available to S_0. Instead

3.2 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. This difference in update potential between (4b) and (5b) is reflected in the different FOCUS-GROUND partitions they display. 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. The verb is focal and the remainder of the sentence is ground. In (4b), where no such assumption is warranted, the entire verb phrase is focal. The update potential of the sentence I_s is expressed in the focus. But information-packaging instructions also indicate how I_s fits in the (speaker's view) of the hearer's input file. In this connection, the role of the link within this ushering function was briefly described in Section 2.1. But in addition to the link, one needs to distinguish between two basic ways in which update of F_1 may proceed (this is argued for in Vallduví 1992). Let us call this two kinds of update 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.

To see what the two modes of update consist of, 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 1.1, links are associated with a particular structural realization, although the actual realization varies from language to language.

Tails, on the other hand, 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

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. We will discuss examples like this in Section 4 below.

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 in (c) should provide an intuitive idea of how the distinction between plain addition to a file card *fc* and modification of a record R on *fc* is implemented. In the case of (7a), this record R is underspecified, so UPDATE-REPLACE is in fact a completion of R with I_s. As noted above, however, R can be a fully specified record as well, in which case I_s alters it (see Section 4).

There are two other possible instruction-types. They are analogous to (6b) and (7b), in that they represent the UPDATE-ADD and the UPDATE-REPLACE modes of update, except they have no link (all-focus and focus-tail instructions). 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. The paraphrases are provided in (d):

(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 the current locus of update, add the 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 the 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. the UPDATE-ADD update is to effected on file card 125.[10]

The contrast between linkless and link-containing instructions is again clear when comparing (10) to (11):

(10) [$_L$ **Sue**] [$_F$ insulted CATHY] and then [$_F$ (she) HIT (her)].

(11) [$_L$ **Sue**] [$_F$ insulted CATHY] and then [$_L$ **she**] [$_F$ hit HER].

The first conjunct in (10) is S_{n-1}. It is a link-focus instruction, which contains the link *Sue*. The file card for Sue is designated as the current locus of update. The second conjunct is S_n. It contains no link, which indicates that content update is to be carried out on the file card for Sue, which was inherited from S_{n-1}. In (11) S_{n-1} again designates the file card for Sue as a locus of update. But here S_n, the second conjuct, is a link-containing instruction as well, as shown by the L+H* prosodic realization of the subject (strong) pronoun. Since establishment of a new locus of update is an inherent part of the meaning of links (if locus of update is inherited, no link appears), the locus of update designated by the link in S_n must be a file card other then Sue's. This, in effect, rules out an anaphoric relation between the link pronoun in the second clause and the link noun phrase in the first clause, thus creating the obligatory switch-reference effect (in (10) both co-reference and disjoint reference are possible).

In sum, a link-containing instruction need be used only when the locus of update is changed from one sentence to the next. If a speaker chooses a file card fc_1 as the locus of update for the information in *both* a sentence S_{n-1} and a sentence S_n, S_n need not contain a link. If, in contrast, file card fc_1 is chosen as a locus of update for S_{n-1} and file card fc_2 is chosen for S_n, then S_n must be realized as a link-containing instruction. The contrast between (10) and (11) regarding the absence or the presence of a link is perhaps more clearly illustrated with the Catalan equivalents of these sentences:

(10′) [$_L$ La Sió] [$_F$ va insultar la COIA]. [$_F$ *e* (li) va fotre una HÒSTIA].

(11′) [$_L$ La Sió] [$_F$ va insultar la COIA]. [$_L$ Ella] [$_F$ (li) va fotre una HÒSTIA].

In the Catalan equivalents of (10) and (11), the absence versus presence of a subject link is not correlated with a prosodic distinction between unaccented and L+H*-accented pronouns, but rather with a distinction between null and overt subject pronouns. If S_n inherits the locus of update from S_{n-1}, as in (10′), it contains no link and is realized with a null subject. If, as in (11′), S_n contains a subject link, i.e., if a new locus of update is designated, then an overt subject appears (ruling out an anaphoric relation between the subjects in S_{n-1} and S_n).

Finally, there is a type of all-focus instruction which deserves special mention. People have referred to these as neutral descriptions (Kuno 1972), thetic judgments (Kuroda 1972, Sasse 1987), news sentences (Schmerling 1976), and event-reporting sentences (Lambrecht 1987). The following are examples of thetic judgments of this sort:

[10] A little aside on the informational import of weak and strong pronominals is in order. In English, weak and strong pronouns are only distinguished by stress/accent. In many other languages, however, the two classes are not only prosodically but also morphologically distinct. Catalan is one such language. Catalan weak pronouns are either null or clitic, whereas strong pronouns are full-fledged, autonomous words. In (8) and (9) the subject pronouns appear within the focal bracketing. In Catalan, these subjects would be null. I take the presence of the pronouns in these examples to be due to independent requirements of English grammar, which does not allow for null argument slots. Weak pronouns are inert as far as information packaging is concerned (but crucial for file-card management) and in (8) and (9) remain vacuously within the focus as place-holders. Strong pronouns are full-fledged items and may take part in an instruction as foci, links, or tails.

(12) So, did anything happen while I was gone?
The PRESIDENT called.

(13) Why don't you go to the theater more often?
TICKETS are expensive

The intuitive informational characterization of this type of sentences is that they are not 'about' any particular referent (see Kuno 1972). They differ in this from link-focus instructions (prototypical topic-comment structures), which have been traditionally analyzed as being 'about' the denotation of the file card designated as locus of update. This can be captured by saying that no particular locus of update is designated for thetic judgments. Rather a salient general temporary situation file card, which is accessible with no ushering, is used to record I_s on.[11]

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):

- UPDATE-ADD(I_s)
- UPDATE-REPLACE(I_s,record(fc))
- GOTO(fc)(UPDATE-ADD(I_s))
- GOTO(fc)(UPDATE-REPLACE(I_s,record(fc)))

The instruction-types described in this section using the notions of mode of update (UPDATE-ADD, UPDATE-REPLACE) and locus of update are a partially formalized reformulation of several traditional proposals concerning the communicative dimension of language. The crucial insight that both these traditional analyses and the instruction-based approach described here share is that speakers operate on the assumption that hearers' models of the common ground, i.e. information states, are highly structured objects. Speakers do not simply present information to their hearers. They also exploit the internal structure of information states to provide hearers with detailed instructions about how to carry out information update optimally. In particular, they exploit the availability of file cards to designate a locus of update or sort key for I_s, and the availability of records on these file cards to express differences in update potential and in mode of update.[12]

[11] Gundel (1974:36) proposes that thetic judgments are, despite appearances, topic-comment structures as well: they are about 'the particular situation (time and place) about which it is asserted (questioned, etc.)'. A similar position is found in Erteschik-Shir 1992. Thus, thetic judgments could be analyzed as link-focus structures where the designated locus of update is the file card associated with a particular time-space slice, i.e. the spatiotemporal argument of the verb. The implications of this option will not be pursued.

[12] Claims about what the structure of the hearer's model of the common ground is or should be like are necessarily made with a speculative slant and must be eventually confirmed by experimental psycholinguistic testing. Another issue awaiting experimental work is whether the filing or sorting choices suggested in information-packaging instructions imply any long-term memory classification of information or are only a part of any working-memory manipulations. In this respect, exploring the connection between information packaging and work on memory by Stenning, Shepherd & Levy (1988) would appear to be worthwhile.

3.3 Without file cards

The present description of the mechanisms underlying information update relies heavily on the particular architecture provided by FCS. For instance, the task carried out by links, that of designating one file card among those expressed in the sentence as a locus of update, loses much of its purpose if we switch to a flatter architecture which does not contain file cards. In DRT, for instance, where discourse markers are dimensionally simpler than file cards, there is only one possible locus of update for I_s: the DRS itself (although independent factors will determine I_s's degree of embeddedness in the DRS). It could be argued that without a file-like architecture, a notion of a locus of update or a notion of sort key is not needed and that arguing that the structure of information states must include such a notion is begging the question. Languages, however, do have links. All languages, in one way or another, may realize a subsegment of the ground in some structurally distinct way. In English and other Germanic languages, there is a subsegment of the ground that is associated with a L+H* accent. In Catalan-type langauges this subsegment of the ground must be expressed in a left-hand clause-peripheral position. In Japanese it is associated with an affix *wa*. If we do away with the notions of locus of update or sort key, we must find another explanation for the consistent marking of links across languages.

Alternatively, one could try to express the notion of locus of update or sort key in a framework that lacks file cards, even if the need for such a notion, given a different architecture, is less clear. Take a simplified DRS like (14), which is the semantic representation of sentences (4) and (5) in Section 3:

$$(14)\quad \begin{array}{|l|}\hline x\ y \\ \hline \textsc{president}(x) \\ \text{D}\textsc{elft-china-set}(y) \\ \textsc{owns}(x,y) \\ \textsc{hates}(x,y) \\ \hline \end{array}$$

Both discourse markers in the universe of (14) have the same status. To reflect the fact that *the president* in (4) and (5) is associated with a link prosody, x must be designated as the sort key for these sentences. It must be thought of as having some special status compared to the other markers in the DRS universe. Whereas in the DRS the notion of locus of update, strictly speaking, has lost much of its raison d'être, the idea of having one of the discourse markers be assigned a special status for sorting purposes could be easily accommodated. Links, then, would be closer to Chafe's (1976) idea of topics as designating the 'frame within which the sentence holds.'

Robin Cooper has pointed out that, from a situation-theoretic viewpoint, linkhood could be thought of as an abstraction. We abstract over the object which is associated with a link realization, yielding a unary type. A unary proposition abstract of this kind is, for our purposes, analogous to a file card, since, in a sense, the object we abstract over has 'scope' over the subsequent information, just as the designated locus-of-update file card has scope over the information recorded on it. In this light, thetic judgments are sentences where no abstraction takes place and where, therefore, no particular object is understood as the 'frame within which the sentence holds'. Of course, these ideas need further elaboration and their applicability to DRSs needs to be ascertained, but they indicate that expressing linkhood in frameworks other than FCS is indeed a feasible task.

As for tails, abandoning a file-like architecture in favor of a DRT architecture does not affect the way in which their role is characterized above. Their role in determining the mode of update and in highlighting a particular record R in the input file can be equally expressed in the absence of file cards if we bear in mind that records are analogous to conditions in a DRS. Of course, DRT, as it stands, does not contemplate the possibility that conditions be modified during communication by means of the UPDATE-REPLACE mode. The same modifications made to the file architecture employed above should be made to DRSs.

4 Nonmonotonicity in information update

Currently, in Update Semantics, downdating with a formula ϕ, i.e. outputing an information state s_2 which is a superset of the input state s_1, is not contemplated as an option. And neither is revision, a situation in which the input s_1 and the output s_2 are disjoint. In fact, in Dekker 1993a, 1993b one of the restrictions imposed on any 'proper information exchange' is that of ϕ not being false in the hearer's input information state at the time of the exchange because '(the resolution of) conflicting information cannot properly be dealt with as long as we stick to the exchange of pure, factual information about the world and about the value of variables' (Dekker 1993b:39).

However, if the relativization of information states to the hearer is taken literally, the existence of downdating and revision must be not only tolerated but expected as well. At any given point in conversation speakers may realize that the content of the hearer's model of the common ground does not agree with their own model, which, of course, they assume to be correct. The speaker, then, may desire to set the hearer straight by indicating that s/he should modify her/his model of the common ground. This constitutes either a revision or a downdate, depending on the sort of modification it entails. If information states were not relativized to the hearer, revision and downdating would indeed be problematic, since by engaging in revision a speaker would be contradicting the common ground and, therefore, him or herself. But what is revised is *not* the common ground, but rather the hearer's model of it (in fact, the speaker's assumptions about the hearer's model of it). As will be shown below, revisions and downdates are in fact standard linguistic exchanges and are particularly associated with an UPDATE-REPLACE mode.

Recognition of the UPDATE-REPLACE mode of update opens the door to nonmonotonic information update. It was noted above that UPDATE-REPLACE carries out the update of information by completing or altering a record already present on a given file card. In the case of example (5) in Section 3, the use of a UPDATE-REPLACE mode of update did not pose any problems for eliminativity. In the contextualization provided, the UPDATE-REPLACE encoded by this sentence is still a reduction of the input information state (albeit a lesser reduction than the one effected with (4)). The record R specified by UPDATE-REPLACE was analyzed as being underspecified, since in this context the presidential aid assumes that no *particular* attitude that the president may have towards the Delft china set is attended to by the Foreign Office Secretary. I_s, in this cases, merely instantiates the blank in R. However, it was pointed out that in a slightly different contextualization, one in which the presidential aid assumes that the Foreign Office Secretary erroneously believes, and is attending to the belief, that the president likes the Delft china set, we would not be dealing with an underspecified R. Rather, R would be fully specified. I such cases, I_s does not instantiate a blank. It actually replaces an element in R.

Carrying out an information update by *altering* a record in the input file clearly leads to nonmonotonicity. Consider (15b), an UPDATE-ADD, and (16b), an UPDATE-REPLACE:

(15) a. S_1: What did you find out?

b. S_2: John is not dead.

(16) a. S_1: Since John is dead, we can now split his inheritance.

b. S_2: I hate to spoil the fun, but John is [$_F$ NOT] dead.

Updating s_1 with the interpretation of (15b) can indeed be viewed as an eliminative process. The output s_2 is the subset of s_1 that contains only the worlds in which John is not dead (all worlds in which John is dead have been eliminated). That is, s_1 contained worlds in which John is dead and worlds in which John is alive. But consider (16b). It is clear from the preceding context that at the time (16b) is uttered, the hearer's s_1 contains only worlds in which John is dead. The update carried out with (16b) yields an s_2 which contains only those worlds in which John is *not* dead. Clearly, s_1 and s_2 are disjoint.

Sentence (16b) is a link-focus-tail (GOTO,UPDATE-REPLACE) instruction where *John* is the link and where (narrow) focus is on the affirmative/negative polarity (*verum* focus; see Jacobs 1984, Prince 1986). This instruction can be represented as in (17) (let us assume 'John' is associated with file card 15):

(17) a. GOTO(15)(UPDATE-REPLACE($-$, { + | + be dead(15)}))

b. 'Go to file card 15 in your input file and update its content by replacing '$-$' for + in the preexisting record "+ be dead(125)".'

Given the context of utterance for (16b), it is obvious that the record R designated by the UPDATE-REPLACE instruction in (17) is not underspecified, but rather fully specified. It is clear S_1 believes John is dead. Hence the update with I_s is not a completion of R but rather an alteration of R. S_1's input F_1, i.e. S_1's information state before processing (16b), contains the datum that John is dead. In S_1's output F_2 that datum has been replaced by the datum that John is not dead. This is a clear case of revision.

In (16b), revision affects the polarity value of R, but other elements in R can be altered. Take a version of (5) where (the presidential aid assumes) the Foreign Office Secretary does believe that the president likes the Delft china set. Here the UPDATE-REPLACE instruction encoded by (5) effects a revision instead of an elimination of possibilities. The revision, however, does not affect polarity but the lexical content of the verb. In the Foreign Office Secretary's F_1 the president likes the Delft china set. In F_2, after (5) is processed, the president hates the Delft china set. Again, F_1 and F_2 are disjoint. The use of an UPDATE-REPLACE for this purpose is natural, given that this mode of update allows the speaker to highlight a particular record R in the hearer's file. R is precisely the premise that must be revised. The revision is carried out when I_s replaces an element in R.

Examples of downdating are viewed much in the same way. I_s replaces an element in the R highlighted by UPDATE-REPLACE. The replacement, however, does not result in disjoint F_1 and F_2 files. Rather, F_2 is a superset of F_1. Consider (18):

(18) a. S_1: I know Mary washed the dishes.

b. S_2: She MAY have washed the dishes, but you can't know for sure.

In (18) I_s is the modal (*may* is a narrow focus). Since (18) encodes an instruction cum tail, the hearer is instructed to find a record R (probably on the file card for *Mary*) describing

Mary having washed the dishes. An element in R (apparently, its epistemic value) is to be replaced by I_s. This is the only change to F_1 that (18) effects. In this case, however, the output s_2 is a superset of the input s_1.

Linguistic interactions of this nature are not felt to be extraordinary or deviant in any sense. They are actually commonplace in natural dialogue and, given the relativization to the hearer of information update, their existence is only to be expected. In fact, Stalnalker himself judged it appropriate to explicitly mention that context change is contingent on the willingness of the audience to accept the assertion (Stalnaker 1978:323). Uttering (16b), however, cannot be taken as S_2's rejection of a previous assertion about John being dead because this bit of information is backgrounded within the embedded clause in (16a)). Of course, S_2's model of the common ground is incompatible with John's being dead, but what S_2 is doing with (16b) is informing S_1 that his model of the common ground is (in S_2's view) contrary to fact and simultaneously providing the necessary revision.

Examples like (16b) are fundamentally different from examples like (19), which illustrate a self-revision carried out by the speaker:

(19) a. S_1: Since John is not here, we can now speak frankly about him.

b. Oh, wait a minute... What am I saying...? John [$_F$ IS] here.

These cases are not problematic for an eliminative update semantics, since they reflect a delusion on the part of the speaker. The putative update effected by S_1 by means of (19a) is a bona fide contribution to the hearer's model of the common ground. The subsequent retraction does have an air of deviation and is better analyzed as a plain and simple repair. In this respect, it is perhaps significant that removing the hedging from (19b) yields an infelicitous sequence. Doing exactly the same in the revision of the hearer's file in (16b), however, does *not* trigger infelicity.[13]

5 Presupposition and ground

Beaver (1993a, 1993b) has developed an account of presupposition within the framework of Update Semantics where presupposition is seen as a 'test' on input information states, i.e. a precondition for update. A sentence containing a given presupposition $\partial\phi$ can update only input states that accept ϕ (Beaver 1993a:43). The ground in UPDATE-REPLACE instructions, according to the anlysis provided in Section 3, also represents, in some sense, a precondition for update. A record R highlighted by the ground in an UPDATE-REPLACE instruction must generally be already present in the input file F_1(it is this preexisting R that I_s modifies). Otherwise use of an UPDATE-REPLACE instruction results in infelicity. This is analogous to saying, in update-semantic terms, that R must be accepted in s_1.

Of course, in the instruction-based approach defended here, the basic characteristic of grounds is the ushering role they perform, which is alien to presupposition. But if a different approach to focus-ground were taken, could grounds be analyzed as a test on input information states too? Can ground and presupposition be collapsed into one single notion? Judging by

[13]There is a way in which revision can be incorporated without abandoning the monotonicity of information growth. It entails introducing time indices in records. In the case of (16a) and (16b), for instance, S_1's F_2 could contain, among many other things, two records. One would specify the belief at time t_1 that John is dead; the other the belief that at time t_2 that John is not dead. Of course, the proposed semantics of UPDATE-REPLACE loses much of its purpose within such a model, since, strictly speaking, no records are actually modified.

the literature on the topic, the collapse of ground and presupposition into one single notion is not empirically warranted. The following is a brief overview of the relevant facts.

Chomsky 1971, Jackendoff 1972, and others after them use the term 'presupposition' to denote the ground in an informational focus-ground structure. In (20a), for instance, *to Joan* is identified as the focus and something like *Bev talked to someone at the party* is called 'presupposition':

(20) a. **Bev** talked [$_F$ to JOAN] at the party.

b. **Bev** talked [$_F$ to NO one] at the party.

Jackendoff (1972:246), however, carefully points out that 'presupposition' is not a presupposition of the traditional sort (e.g. constancy under negation). The ground ('presupposition') of both (20a) and (20b) is the same. Whereas (20a) may appear to presuppose that there is someone that Bev talked to at the party, this is certainly not the case in (20b): in (20b) the ground is not an existential presupposition.[14]

Clearly, we cannot enforce a constraint that the ground of a given sentence S be entailed or presupposed by ϕ_S. This has actually led many authors to abandon the term 'presupposition' in favor of terms like open-proposition, background, ground, or theme. Classic presupposition, in fact, appears to be orthogonal to the focus-ground partition. Erteschik-Shir (1986) discusses examples of the type of (21) with presuppositional predicates like *regret*, where the presupposition of a sentence is focal:

(21) What does John regret?
John regrets [$_F$ Peter lost his job].

Delin 1991, analyzing it-clefts, argues that assertion-presupposition and focus-ground are orthogonal dimensions as well. It-clefts are presuppositional structures where the nonclefted part is existentially presupposed. For instance, the cleft in (22a) presupposes (22b):

(22) a. It's Harry that taught me how to tango.

b. $\exists x$ x taught me how to tango

(23) Who taught you how to tango?
It's [$_F$ HARRY] who taught me how to tango.

While it is possible–indeed quite common—to find it-clefts where the presupposition is identical with the ground, as in (23), it is also possible to find clefts where the focus is precisely the presupposition (and where, therefore, the assertion is part of the ground). Sentence (24a) is synonymous, as far as focus-ground structure is concerned, with (24b), not with (24c):

(24) Why are you so fond of Harry?

a. It is Harry that taught me how to TANGO.

b. Harry [$_F$ taught me how to TANGO].

c. # [$_F$ HARRY] taught me how to tango.

[14] Given that information states are relativized to the hearer, it could be argued correctly that, while it is true that for the speaker of (20b) the existential presupposition that Bev talked to someone at the party does not hold, for a potential hearer, which is in an erroneous belief state (in the eyes of the speaker), it does hold. This move, however, effectively relativizes existential presupposition to the hearer, something that runs against the received view and which is problematic in many cases, including those we will see below.

In sum, when the full range of facts is considered, there is no empirical motivation to identify the ground with a presupposition, despite what may be suggested by some unfortunate terminological choices.

Sentences (21) and (24a) presuppose that Peter lost his job (call this presupposition ϕ) and that someone taught the speaker to tango (call it ψ), respectively. Following Beaver 1993a, ϕ and ψ must be accepted by the information states they update, otherwise there is no update possible. However, the contexts of both (21) and (24a) favor (or at least allow) a scenario where ϕ and ψ do not hold in the hearer's input state. The answer in (21) is felicitous even in a context where the hearer knows nothing about Peter's dismissal and (24a) is a felicitous addition to a context in which the hearer does not know the speaker knows how to tango. These sentences successfully update their hearer's input information states by identifying John's regret and by identifying the reason for the speaker's fondness for Harry, even though these are precisely the presuppositions of (21) and (24a).

Presuppositions ϕ and ψ, then, are not (or need not be) accepted in the hearer's input information state. In fact, ϕ and ψ are more amenable to an analysis in terms of speaker presupposition. In other words, ϕ and ψ are not (assumed by the speaker to be) things that hold in the *hearer's* model of the common ground, but rather things that the speaker takes for granted, i.e. things the speaker assumes are indeed part of the actual common ground. The distinction appears to be important. Presuppositions ϕ and ψ are taken to be part of (the speaker's model of) the common ground, but not necessarily to be part of (the speaker's view of) the hearer's model of the common ground. Contrast this with example (20b) above. In (20b) the ground (a record R on the file card for Bev in the hearer's F_1) is, roughly, that Bev talked to someone at the party. The speaker of (20b), clearly, does not take this ground for granted or as part of (her model of) the common ground, although she assumes is part of the *hearer's* model of the common ground. Grounds, therefore, hold in the hearer's model of the common ground (in the speaker's view), while presuppositions, at least ϕ and ψ, hold in (the speaker's model of) the actual common ground. Of course, this distinction surfaces mostly in sentences where assertion-presupposition and focus-ground are in tension. In examples like (23), where the focus is the assertion and the ground the presupposition the contrast vanishes. That someone taught the speaker to tango is both speaker-presupposed, as a presupposition, and a part of the hearer's model of the common ground, as a ground.[15]

So far, a substantial part of Beaver's analysis of presupposition, specifically his treatment of presuppositional accommodation, has been ignored. Beaver acknowledges that presuppositions can in fact be used to provide information, i.e., that presuppositions may not be accepted in the hearer's input infomation state. Following Lewis (1979) he identifies these cases as 'special'. They are cases of accommodation. Presumably, it could be argued that presuppositions ϕ and ψ in (21) and (24a), which are indeed informative, are, in fact, cases of accommodation, although in doing so we stretch the scope of accommodation considerably.

Beaver suggests that accommodation is a form of filtering operation on epistemic alternatives. Epistemic alternatives are alternative models of the common ground that the hearer may be entertaining at any given moment. The idea is that hearers are always uncertain about what the common ground contains and, therefore, do not have a unique model of the common ground, but rather juggle with a set of distinct models that are compatible with their

[15] A similar structure with multiple representations of the common ground for speaker and hearer is used by von Klopp 1993 to lay out a dynamic analysis of negation. She points out the widespread use of such multiple representations in artificial intelligence, e.g. Cohen & Perrault 1979.

knowledge about the common ground. The task of $\partial\phi$ would be to filter out any epistemic alternatives where ϕ is not accepted. Presuppositions ϕ and ψ in (21) and (24a) would not be acting as tests on the hearer's input information state, but rather as filters of unwarranted epistemic alternatives.

Given the need for the filtering of epistemic alternatives and in light of the remarks above about the comparison of ground and presupposition, an alternative to the dual nature of presupposition suggests itself: filtering is not just the function of accommodated presupposition, but rather the basic function of presupposition in general. Presuppositions are used by the speaker to make sure that the hearer's model of the common ground is appropriate. Upon encountering $\partial\phi$, a hearer knows that any epistemic alternatives that do not accept ϕ do not correspond to the (speaker's model of the) common ground and proceeds to filter them out. This is what is going on in (21) and (24a), which arguably are not cases of accommodation. Presupposition performs the same task in (23), which presupposes that someone taught the speaker how to tango, although here the filtering is vacuous, since each one of the hearer's alternative models of the common ground happens to accept this presupposition. From this perspective, then, $\partial\phi$ holds of (the speaker's model of) the common ground, while a ground in an UPDATE-REPLACE instruction holds of the hearer's model of the common ground. The function of the former is to filter out hearer's epistemic alternatives that conflict with ϕ, while the function of the latter is to highlight which record in the hearer's F_1 is affected by the completion or alteration effected by I_s.

Of course, this proposal, as it stands, is very rough and sketchy. Whether relativizing presupposition to the speaker as opposed to the hearer withstands more serious analysis is yet to be seen. The traditional insight that any account of assertion-presupposition and focus-ground must reflect, however, is that these distinctions are orthogonal to each other and that their functions in the dynamics of communication must be distinct.

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 has taken a look at the sorts of instructions we find in communication and has suggested 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 seems to favor a view of an information state as a collection 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, it should be mentioned that the approach is coached in a constructive view of information update, rather than an eliminative one, although the consequences of this 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.

References

Beaver, David I., 1993a. The kinematics of presupposition. In Hans Kamp, ed., *Presupposition (DYANA-2 Report R2.2.A, part II)*, 31-71. Institute for Language, Logic and Computation, University of Amsterdam.

Beaver, David I., 1993b. Kinematic Montague Grammar. In Hans Kamp, ed., *Presupposition (DYANA-2 Report R2.2.A, part II)*, 73-133. Institute for Language, Logic and Computation, University of Amsterdam.

Chafe, Wallace L., 1976. Givenness, contrastiveness, definiteness, subjects, topics, and point of view. In Charles N. Li, ed., *Subject and topic*, 25-55. New York: Associated Press.

Chomsky, Noam, 1971. Deep structure, surface structure, and semantic interpretation. In Daniel Steinberg & Leon Jacobovits, eds., *Semantics*, 183-216. Cambridge: Cambridge University Press.

Clark, Herbert H. & Susan E. Haviland, 1977. Comprehension and the given-new contract. In Roy O. Freedle, ed., *Discourse production and comprehension*, 1-40. Hillsdale, N.J.: Lawrence Erlbaum Associates.

Cohen, Philip R. & C. Raymond Perrault, 1979. Elements of a plan-based theory of speech acts. *Cognitive Science* 3.177-212.

Dekker, Paul, 1993a. *Transsentential meditations*. Universiteit van Amsterdam Ph.D. dissertation.

Dekker, Paul, 1993b. Updates in dynamic semantics. In Robin Cooper, ed., *Integrating semantic theories (DYANA-2 Report R2.1.A)*, 1-44. Institute for Language, Logic and Computation, University of Amsterdam.

Delin, Judy L., 1991. Presupposition and shared knowledge in *it*-clefts. Submitted to *Language and Cognitive Processes*.

Erteschik-Shir, Nomi, 1986. Wh-questions and focus. *Linguistics and Philosophy* 9.117-49.

Erteschik-Shir, Nomi, 1992. Focus structure and predication. *Belgian Journal of Linguistics* 7.35-51.

Ginzburg, Jonathan, 1993. Resolving questions. Research Paper HCRC/RP-46, Human Communication Research Centre, University of Edinburgh.

Groenendijk, Jeroen & Martin Stokhof, 1991. Dynamic predicate logic. *Linguistics and Philosophy* 14.39-100.

Grosz, Barbara, Aravind K. Joshi & Scott 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, Jeanette K., 1974. *The role of topic and comment in linguistic theory*. University of Texas Ph.D. dissertation (published by Garland, New York, 1988).

Gundel, Jeanette K., Nancy Hedberg & Ron Zacharski, 1993. Cognitive status and the form of referring expressions. *Language* 69.274-307.

Halliday, Michael A.K., 1985. *An introduction to functional grammar*. London: Arnold.

Heim, Irene, 1982. *The semantics of definite and indefinite noun phrases*. University of Massachusetts Ph.D. dissertation.

Heim, Irene, 1983. File Change Semantics and the familiarity theory of definiteness. In Reiner Bäuerle et al., eds., *Meaning, use, and interpretation of language*, 164-189. Berlin: de Gruyter.

Hockett, Charles F., 1958. *A course in modern linguistics*. New York: Macmillan.

Jackendoff, Ray, 1972. *Semantic interpretation in generative grammar*. Cambridge: MIT Press.

Jacobs, Joachim, 1984. Funktionale Satzperspektive und Illokutionssemantik. *Linguistische Berichte* 91.25-58.

Karttunen, Lauri, 1976. Discourse referents. In James D. McCawley, ed., *Notes from the linguistic underground* (*Syntax and Semantics* 7), 363-385. New York: Academic Press.

Klopp, Ana von, 1993. *Negation: Implications for theories of natural language*. University of Edinburgh Ph.D. dissertation.

Kuno, Susumu, 1972. Functional sentence perspective. *Linguistic Inquiry* 3.269-320.

Kuroda, Shige-Yuki, 1972. The categorical and the thetic judgment: Evidence from Japanese syntax. *Foundations of Language* 9.153-85.

Lambrecht, Knud, 1987. Sentence focus, information structure, and the thetic-categorical distinction. *Berkeley Linguistics Society* 13.366-382.

Lewis, David, 1979. Score-keeping in a language game. *Journal of Philosophical Logic* 8.339-359.

Prince, Ellen F., 1986. On the syntactic marking of presupposed open propositions. *Chicago Linguistic Society* (*Parasession*) 22.208-22.

Prince, Ellen F., 1992a. The ZPG letter: Subjects, definiteness, and information-status. In Sandra A. Thompson & William C. Mann, eds., *Discourse description: Diverse analyses of a fund-raising text*, 295-325. Amsterdam & Philadelphia: John Benjamins.

Prince, Ellen F., 1992b. Informational and rhetorical structure. In William Bright, ed., *International encyclopedia of linguistics*, Vol. 3, 399-401. Oxford: Oxford University Press.

Reinhart, Tanya, 1982. Pragmatics and linguistics: an analysis of sentence topics. *Philosophica* 27.53-94.

Rochemont, Michael S., 1986. *Focus in generative grammar*. Amsterdam/Philadelphia: John Benjamins.

Sasse, Hans-Jurgen, 1987. The Thetic/categorical distinction revisited. *Linguistics* 25.511-580.

Schmerling. Susan F., 1976. *Aspects of English sentence stress*. Austin: University of Texas Press.

Sidner, Candace L., 1981. Focusing for interpretation of pronouns. *American Journal for Computational Linguistics* 7.217-231.

Stalnaker, Robert, 1978. Assertion. In Peter Cole, ed., *Pragmatics* (*Syntax & Semantics* 9), 315-332. New York: Academic Press.

Stenning, Keith, Martin Shepherd & Joe Levy, 1988. On the construction of representations for individuals from descriptions in text. *Language and Cognitive Processes* 3.129-164.

Vallduví, Enric, 1992. *The informational component*. New York: Garland.

Veltman, Frank, 1990. Defaults in update semantics. In Hans Kamp, ed., *Conditionals, defaults, and belief revision* (*DYANA Report R2.5.A*), 28-63. Edinburgh: Centre for Cognitive Science, University of Edinburgh.

Ward, Gregory L., 1988. *The semantics and pragmatics of preposing*. New York: Garland.

Ward, Gregory L. & Julia Hirschberg, 1985. Implicating uncertainty. *Language* 61.747-776.

Enric Vallduví
Centre for Cognitive Science
University of Edinburgh
2 Buccleuch Place
Edinburgh EH8 9LW, Scotland

enric@cogsci.ed.ac.uk