Questions and Exhaustivity in Update Semantics

Henk Zeevat
University of Amsterdam

Abstract

This paper presents an exhaustivity operator suitable for update semantics and discusses one application of this operator: questions. The exhaustivity operator takes an open formula and assigns (if this is possible) the maximal values to the free variables such that the formula is true as a result and entails any sentence that can be obtained from the formula by assigning other values to the variables to obtain a true sentence.

Update semantics offers a simple and straightforward interpretation of the semantics of dialogue. Taking each conversational step to result in a change to the common ground (the information state manipulated by the update semantics) we obtain a semantics that is closely connected to the natural environment of natural language use: conversation.

Update semantics seems to offer the possibility of integrating certain pragmatical phenomena within the semantics. Stalnaker's approach has led to elegant treatments of presupposition and a partial treatment of implicatures. Elsewhere (Zeevat & Scha (1992)), we have defended the view that update semantics is particularly suited for developing pragmatics and semantics within a single theory. However, any theory of dialogue that can be taken seriously needs a theory of questions. This is not a trivial matter, as most of the available theories of questions seem to be in conflict with the fundamental assumption of update semantics that meaning is characterised as an update. Available theories treat questions as sets of answers, sets of true answers and as propositional concepts of true answers. None of these can themselves be seen as giving information. At most, the combination of the question with the answer can be seen as an information carrier.

On the view defended in this paper a question is a dynamical proposition. It expresses the statement that certain discourse markers (corresponding to the Wh-elements in the linguistic form) have an exhaustive value. Moreover, it makes these variables into discourse referents referring to their exhaustive values.

A positive answer to the question identifies the discourse markers of certain of its subexpressions with the Wh-discourse markers from the question. This results in exhaustive interpretations for the answers. Negative answers result in the negation of the question.

Integrating negative answers and answers that decline to answer the question makes it necessary to see the update associated with the question as a provisional one. The refusal to answer brings us back to the original information state[1].

The remainder of the paper consists of an analysis of exhaustivity in update semantics and a comparison of the current theory with the theory of Groenendijk & Stokhof (1984). I also refer to this work for a defense of the view that exhaustivity must be part of the semantics of questions.

0.1 An Eliminative Update Semantics

In an eliminative update semantics, we increase the information in an information state by eliminating information carriers: those where the new information does not hold. The system considered here is essentially the same as the system in Kamp (1981) but for the fact that it is formulated as an eliminative update semantics. Kamp's system can be seen as a syntactic update semantics where information growth is expressed by the addition of new conditions and discourse markers.

Information carriers will be models of sublanguages $L_0 \subset L$, where L is a set of sorted individual constants, signed predicate sysmbols and signed function symbols. In particular, we have constants for natural numbers and for sets of objects (without the empty set and the number zero).

Our language is that of variable-free first order logic without quantifiers, with the addition of individual constants appearing as atomic formulas. The function of such appearances of constants as formulas is similar to that of the discourse markers of Kamp (1981). Discourse markers can easily be defined by recursion over the formulas, but a semantic definition is preferred here.

Information states are sets of information carriers. The update of an information state σ by a formula φ is defined in (1).

(1)
1. $\sigma[c] = \{i \in \sigma : ic \text{ defined}\}$
2. $\sigma[Pt_1, \ldots, t_n] = \{i \in \sigma : \neg < it_1, \ldots it_n > \notin iP\}$
3. $\sigma[t_1 = t_2] = \{i \in \sigma : \neg(it_1 = u, it_2 = v \wedge u \neq v)\}$
4. $\sigma[\varphi \wedge \psi] = \sigma[\varphi][\psi]$
5. $\sigma[\neg\varphi] = \mathbf{neg}(\sigma, \sigma[\varphi])$
6. $\sigma[\varphi \rightarrow \psi] = \sigma[\neg(\varphi \wedge \neg\psi)]$

For the negation we need the definition (2),

(2) $dm(\sigma, \tau) = \{c \in L : \exists i \in \sigma\ ic \text{ is undefined} \wedge \forall i \in \tau\ ic \text{ is defined}\}$

[1] We do not treat here the epistemic states of the conversation participants. Otherwise, the answer *I do not know.* would have to be integrated in the information state.

In terms of this operation, we can define **neg** in (3).

(3) $\mathbf{neg}(\sigma, \tau) = \sigma - \{i : \exists j \in \tau \ i =_{dm(\sigma,\tau)} j\}$

The first three clauses are set up in such a way that there is a distinction between an atomic formula eliminating information carriers and updating the *conjunction* of the constants occurring within the atomic formula with the atomic formula: only in the latter case, it is guaranteed that each of the variables will be defined throughout the information state. The atomic formulas only eliminate those carriers that overtly contradict them. This distinction is necessary for the semantic treatment of discourse markers.

1 Exhaustivity

What is the exhaustive interpretation of a free variable in a formula? Intuitively, it is a value for the variable which makes the formula true and which makes it entail all true formulas that can be obtained by assigning another value to the variable. I.e. the exhaustive interpretation is the strongest interpretation that the open formula allows. Of course, there need not exist an exhaustive interpretation for a formula. This is indeed a common situation. Suppose, e.g., that five boys are asleep. It is then impossible to have an exhaustive reading for sentences like (4a/b) .

(4) a. One boy sleeps.
$boy(x) \wedge \#x = 1 \wedge sleep(x)$
b. Less than three boys sleep.
$x \wedge boy(x) \wedge \#x < 3 \wedge sleep(x)$

Exhaustification is thereby a combination of the statement that exhaustive readings are possible together with assigning the exhaustive value to the free variable. When exhaustification is possible, it gives minimal or maximal elements with respect to some order, e.g. set inclusion, smaller than on numbers.

The concept of interpreted open formulas entailing one another is not a standard one and the following remarks are intended to make it precise. The problem is that we need define a notion of entailment over interpreted formulas. Interpretation normally involves a single model and does not make sense over arbitrary classes of models (e.g. the value may not be available in another model, or it may play a completely different role). Entailment however essentially involves a quantification over models.

For this reason we will limit our models to classes K of models M which are expansions to a language L of a given model M_0 for a language $L_0 \subseteq L$. The given model fixes the domain and some privileged functions and relations. For the examples considered in this paper, it suffices to take the basic model to be the power set of some given set (without $\emptyset$) together with the set of natural numbers (without 0). The privileged functions and relations include *inclusion* between the sets, *smaller than* between numbers and *cardinality* relating sets and numbers. Object variables will range over sets of objects, number variables over numbers. (Of course, other relations and sorts can be considered.)

In addition, we assume a set MP of postulates about the non-privileged relations. With entailment, we will mean K-entailment from now on.

K can be written as $K = \{M : M \models MP \ \& \ M|L_0 = M_0\}$ where $|$ is the restriction operator.

Let φ be a formula with some free variable x and K a class of models M as described above. The exhaustification of φ in K with respect to the variable x is that object u in the domain U_M of M such that (5).

(5) (1) $M \models \varphi < u >$ and
(2) $\forall v \in U_M \forall M_1 \in K(M \models \varphi < v >$ and $M_1 \models \varphi < u > \Rightarrow M_1 \models \varphi < v >)$.

Example 1.

Let K be as described above. Let MP contain the formula: $Px \wedge y \subseteq x \rightarrow Py$.(gloss: If John has sheep x then John has sheep y for $y \subset x$.)

Let $\varphi = Px$. Then an exhaustive value for x in the model M is the set of all P in M. (gloss: John's sheep.)

Example 2.

$MP = \{Pn \wedge m > n \rightarrow Pm\}$. (gloss: If John runs the mile in n minutes then John runs it in m minutes if $m > n$)

Let $\varphi = Pn$. The exhaustive value is the smallest number such that P in M. (John's time for the mile.)

1.1 Exhaustivity in Update Semantics

Exhaustive updates are updates where the discourse markers of an update are exhaustified. We eliminate the information carriers where the formula does not hold and those in which the carrier does not give an exhaustive value to the variable. The first elimination is conventional, for the second we need to import exhaustification.

Information carriers are models. Quantification is dealt with by considering other information carriers in the information state which are almost exactly the same except for the value assigned to the discourse markers. This relation is standard. For exhaustiveness, we will introduce another relation of controlled variation. The two information carriers must have the same domain, the same interpretation of all individual constants and must coincide with respect to a set of privileged relations. (the analogue of the L_0 relations considered above).

They can vary with respect to the interpretation of non-privileged predicates and functions. We call this relation object-identity: the two information carriers agree with respect to the objects and basic ontological relationships. They may disagree about everything else.

As before, we expect the information state (i.e. every carrier in it) to satisfy a set of meaning postulates. Parts of the information state of the form $\{j \in \sigma : j$ is object-identical with $i\}$ have the same structure as the classes K consisting of expansions of a given model. So we obtain, for each of the information-carrier i in the information state σ, a set of models K_i that can support the entailment relation needed for judging whether i assigns an exhaustive value to the variable. Because σ contains the conceptual information in P, so does each of the K_i.

The following diagram shows the demand of exhaustiveness on a constant c with respect to the information state $\sigma[\varphi]$.

$$\begin{tikzcd}
\mathtt{i} \arrow[r, "\texttt{object-identical}"] \arrow[d, "\texttt{c-variant}"'] & \mathtt{j} \arrow[d, "\texttt{c-variant}"] \\
\mathtt{k} \arrow[r, "\texttt{object-identical}"'] & \mathtt{l}
\end{tikzcd}$$

$i : sleepers = \{j, b, m\}$
$c = \{j, b\}$

$j : sleepers = \{j, b\}$
$c = \{j, b\}$ (obj.id. i)

$k : sleepers = \{j, b, m\}$ (c-var. i)
$c = \{j, b, m\}$

$l : sleepers = \{j, b\}$ (c-var. j)
$c = \{j, b, m\}$ (obj.id. k)

So, l does not satisfy $sleep(c)$

Exhaustivity Diagram

i is exhaustive for c iff whenever i has an c-variant i_c and an object-identical variant i_o in $\sigma[\varphi]$ then i_o has an c-variant $i_{oc} \in \sigma[\varphi]$ that is object-identical to i_c.

To see that this is correct, consider what we mean by exhaustive values. φ according to i should entail all of the φ-meanings in c-variants i_c of i. It does not do so, if the carrier i has an c-variant i_c, but there is also a world w, where φ is the same as in i, but which lacks the corresponding c-variant. The value i assigns to the discourse markers is then compatible with the absence of the values i_c assigns to them and so φ with i's values does not entail φ with i_c's values.

Object-identity here guarantees two things: it guarantees that the same value is given to c and that c is not just formally the same: it plays the same role in the ontology of the other world. So w must be object-identical to i. The corresponding c-variant must similarly be object-identical to i_c. A counterexample to i being exhaustive for c and φ with respect to some σ is therefore an c-variant i_c and an object-identical i_o, both in σ, which lack an element i_{oc} that is object-identical to i_c and c-variant to i_o.

To go back to our earlier example:

We need the meaning postulate:$z \wedge y \wedge Pz \wedge y \subseteq z \rightarrow Py$, i.e. we assume that (6)

(6) $\sigma[z \wedge y \wedge Pz \wedge y \subseteq z \rightarrow Py] = \sigma$

and $\varphi = Pc$, where c is new to the information state. (6) forces the interpretation of P in an information carrier i to be of the form $pow(A)$ for some set A. We will show that i is not exhaustive for φ if i assigns to c a proper subset of A.

Suppose i assigns $pow(A)$ to P, and $B \subset A$ to c. Take i_c such that i_c assigns A to c. $i_c \in \sigma[Pc]$ since c is new and by the assumption.

Consider i_o such that i_o assigns $pow(B)$ to P. $i_o \in \sigma[Pc]$ since $i_o c = B$.

Then there is no i_{oc} such that $i_{oc} \in \sigma[Pc]$, i_{oc} is an c-variant of i_o and i_{oc} is object-identical to i_c, since (by object-identity) i_{oc} assigns A to c and (by c-variance) i_{oc} assigns $pow(B)$ to P. But then $i_{oc} \notin \sigma[Pc]$

So c must be A if i is exhaustive.

The exhaustivity operator is defined on the basis of the above. The operator will take the discourse referents of a formula and deliver an exhaustive reading for all of them if such a reading exists.

(7) $\sigma[q(\varphi)] = \{i \in \sigma[\varphi] : \forall j, k \in \sigma[\varphi]\ (j =_{dm(\sigma,\sigma[\varphi])} i \wedge k$ is object-identical to $i \exists l \in \sigma[\varphi]\ (l =_{dm(\sigma,\sigma[\varphi])} k \wedge l$ is object-identical to $j))\}$

We are now in a position to state the updates corresponding for some short dialogues. Space prevents me from a treatment of the relation with surface forms.

(8) a. Who sleeps? John and Bill sleep.
$\sigma[q(c \wedge sleep(c))][union(b, j, c) \wedge sleep(c)]$
b. Who sleeps? Nobody.
$\neg(\sigma, \sigma[q(c \wedge sleep(c))])$
c. Which boy danced with which girl? John with Suzy.
$\sigma[q(c \wedge boy(c) \wedge d \wedge girl(d) \wedge dance\_with(c, d))][j = c \wedge s = d]$

2 Questions

Standard answers to questions are in the theory of Groenendijk and Stokhof characterised by three properties:

(9) They are true
They are rigid
They are exhaustive

This leads to two semantic characterisations of questions: as concepts of their standard answers (i.e. functions that assign to a possible world the standard answer in that world) and -equivalently- as partitions over possible worlds, induced by the relation $\{< i, j >$:the answer denoted by the question in i contains $j\}$).

In principle, the theory of answering that is natural for this theory is the statement *the extension of the question is the answer.* This will only work if indeed we receive rigid and full answers. In real life, this is not always the case: non-rigid answers occur and many answers tell us less than what is required. The theory needs therefore to be enriched by a pragmatics of answers. In Groenendijk & Stokhof this involves the restriction to an information state on the one hand and the development of the notion of a partial answer.

Within update semantics, there is no need for a restriction to an information state: this is how we start anyway. But do we need partial answers? I think that there is a case for assuming that they may be dispensed with completely, if we drop rigidity. A proper partial answer is then one where the answerer indicates that she is not giving a full answer to the question that was asked, but a standard answer to a weaker question. It is the task of the person interpreting the answer to work out the weaker question on the basis of the formal properties of the answer and the original question. A mechanism for finding hidden questions has now become a fairly standard assumption in the theory of topic and focus and there is no good reason why we should not appeal to it here. In particular, so-called non-exhaustive answers would be a good candidate for this treatment.

The reason why rigidity can be dropped in our current framework is that we have (by using discourse markers for Wh-elements) made sure that questions "denote" rigid answers. Identifying these with the markers of an answer makes it superfluous to have an additional rigidity demand in the answer, beyond such rigidity as can be derived from Gricean informativeness.

Much the same holds for the need to make the answer exhaustive by means of an extra operator. The mechanism that identifies Wh-element markers with the corresponding markers in the answer guarantees that the relevant discourse markers all have an exhaustive interpretation with respect to the question and the information state. Further exhaustivity is not needed.

So in the pragmatics of answers, we depart in three ways from Stokhof & Groenendijk: all answers are standard, but possibly to a derived question and all exhaustivity and rigidity derives from the question itself. Truth also disappears, but this is purely a question of the framework: within the update semantics there is no need for a particular information carrier which represents the real world[2].

We have so far not considered what it means for an information state to answer a question. (In terms of this notion, complete answers can be defined: does $\sigma[Q][A]$ answer Q?) This is easy in Groenendijk & Stokhof (1984): an information state σ answers a question if the question partitions it into the partition $\{\sigma\}$. Here we use a similar idea, but things are one degree more complex, due to the fact that we require the Wh-elements of the question to be new and that we do not have a unique structured domain.

The question update eliminates the carriers that lack an appropriate exhaustive value for the Wh-elements. What they then do for the other worlds is select the appropriate *dm*-variant: this is the answer according to a carrier i for which an answer exists. So the question update divides the carriers into two classes: one where there is no answer, another where the carrier gets mapped to its exhaustive *dm*-variant. That an information state answers a question means first of all that the first class is empty: every carrier should determine an answer. Second, they should all determine the same answer.

This can be expressed as conditions on $\sigma[q(\varphi)]$. If σ answers $q(\varphi)$, the only effect of the update is that the allowed values for the discourse markers of $q(\varphi)$ are restricted. This implies that every $i \in \sigma$ should have a *dm*-variant in $\sigma[q(\varphi)]$. It follows from the definition of the question-update that only one *dm*-variant survives the update. So, sets of carriers that are *dm*-variants of each other collapse to a single element (this follows from the nature of the question update). In this case, we have a function f_q mapping σ to $\sigma[q(\varphi)]$. This does not yet guarantee that the same answer is given. But we can express this adequately by the demand that f_q preserves object-identity: f_q maps object-identical carriers to object-identical carriers.

In (10), a definition is given that expresses these two conditions directly in terms of object-identity and *dm*-variance. If σ does not answer the question, there are carriers without an answer or there are two carriers that have the same ideas about the structure of the world, but different ideas about the facts which determine different answers.

(10) σ answers $q(\varphi)$ iff
1. $\forall i \in \sigma\ \exists ! j \in \sigma[q(\varphi)]\ i =_{dm(\sigma,\sigma[q(\varphi)])} j$
2. $\forall i, j \in \sigma\ \forall k, l \in \sigma[q(\varphi)]\ (oid(i,j) \wedge i =_{dm(\sigma,\sigma[q(\varphi)])} k \wedge j =_{dm(\sigma,\sigma(q(\varphi)])} l \rightarrow oid(k,l)))$

In essence, our definition is the same as that of Groenendijk & Stokhof. Object-identity composed with *dm*-variance for a fixed set of discourse markers is an equivalence relation.

[2] Following Heim (1983), we could add a privileged information carrier in standard states. This would not help us with the problem at hand however as we have no explicit notion of the denotation of a question (its answer) in a carrier. The function f_q can be considered as an approximation of this notion, but it gives carriers, not propositions as its result. There is no guarantee that any carrier will be mapped to itself by this function.

(11) $\quad iRj \Leftrightarrow_{def} \exists k\ (oid(i,k) \wedge k =_X j)$

Within an equivalence class of that relation we can form a partition by looking at the actual values for the discourse markers.

(12) $\quad iQj \Leftrightarrow_{def} \forall x \in X\ ix = jx$

Consider $\tau = \sigma[q(\varphi)]$ and let $X = dm(\sigma, \tau)$. Consider $i/R/Q$ with respect to τ. Now, $i/R = i/R/Q$, i.e. the partition determined by Q over i/R is $\{i/R\}$, if σ answers the question.

If we demand that all carriers are related by our relation R, answers are rigid: if σ answers $q(\varphi)$, the question update guarantees that the discourse markers of φ are defined and rigid over σ. What we have in fact for answering a question is rigidity modulo R.

The main difference with the theory of Groenendijk and Stokhof is that the denotation of a question in an information state is implicit rather than explicit. Explicit denotation plays an important role in the analysis of embedded questions. The problem with embedded questions is primarily that no fully satisfactory semantics for the attitudes has so far been given in update semantics.

Let me sketch a possible approach to knowledge. We assume that for $i \in \sigma$ we have a set of carriers K_i representing John's knowledge. We can then define knowledge updates in the following way[3]:

(13) $\quad \sigma[K\varphi] = \{i \in \sigma[\varphi] : K_i[\varphi] = K_i\}$

This definition guarantees by means of an extra update that φ also holds in σ after the update. It would be more proper to arrange this by presupposition, as the projection behaviour is not correctly characterised by putting an extra update inside the larger update, but we will not bother about this here. We demand that $i \in K_i$. We must also guarantee that K_i and i have the same ideas about discourse markers. A tentative demand lets i and $j \in K_i$ be related by the composition of object-identity and variation for markers outside the discourse markers of σ that are also markers of K_i.

How do things work out with an example like (14)?

(14) $\quad$ John knows who sleeps.

We must form $\sigma[q(c \wedge sleep(c))]$ and reduce that to those i where K_i is a fixpoint for the update with $q(c \wedge sleep(c))$. Since i is in $\sigma[q(c \wedge sleep(c))]$, i agrees with K_i on the value of c.

[3]This is the same treatment as the one suggested by Kamp for belief according to Heim (1992). There are alternatives for this definition.

In σ, however, it need not be clear who sleeps, so there may be object-identicals j of i that have an c-variant in $\sigma[q(c \wedge sleep(c))]$ which is also a fix-point for the update with $Kq(c \wedge sleep(c))$.

If the answer is known this cannot happen. So like in GS, we have (15):

(15) Bill sleeps.
John knows who sleeps.
ergo
John knows that Bill sleeps.

and its negative variant.

How about other environments of indirect questions? I only have a version of *wonder* here, that is definable as the negation of our knowledge operator (in fact, this works only because we did not treat knowledge as presuppositional).

(16) John wonders who sleeps

(16) comes out as an update that eliminates the carriers where John knows who sleeps and their *dm*-variants. It is clear that this time the corresponding inferences are blocked.

So, I conclude that nothing indicates that it is impossible to develop indirect questions on the lines we have indicated so far for direct questions. The difference, so elegantly captured, by making knowledge an extensional predicate of indirect questions and wondering an intensional predicate, returns here as a difference between concepts that presuppose the existence (truth) of their arguments. It may well be that other uses of a denotation mechanism can be reduced to a distinction between presupposing or not.

References

Groenendijk, J. & Stokhof, M. *Studies on the Semantics of Questions and the Pragmatics of Answers.* Academisch Proefschrift, Universiteit van Amsterdam 1984.

Heim, Irene. On the Projection Problem for Presuppositions. *WCCFL*2 (1983), p. 114-126.

Heim, Irene. Presupposition Projection and the the Semantics of Attitude Verbs. In: *Journal of Semantics* **9**, 3, (1992), p. 183-222.

Kamp, H. A Theory of Truth and Semantic Representation. In: J.A.G. Groenendijk, T.M.V. Janssen & M.B.J. Stokhof (eds.). *Formal Methods in the Study of Language.* Mathematical Centre, Amsterdam 1981. p. 277-322.

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

Veltman, F. Defaults in Update Semantics. Accepted by *Journal of Philosophical Logic.* (ms. 1994, University of Amsterdam)

Zeevat, H. & Scha, R. Integrating Pragmatics in Update Semantics. In: A. Ortony, J. Slack & O. Stock (eds.) *Communication from an Artificial Intelligence Perspective: Theoretical and Applied Issues.* Springer Verlag Heidelberg, 1992.