Research on Language and Computation **2:** 401–416, 2004.

Utility of Mention-Some Questions

ROBERT VAN ROOY
ILLC/University of Amsterdam, The Netherlands (E-mail: vanrooy@hum.uva.nl)

Abstract. In this paper, I argue that the 'ambiguity' between mention-all and mention-some readings of questions can be resolved when we relate it to the *decision problem* of the questioner. By relating questions to decision problems, I (i) show how we can measure the utilities of both mention-all and mention-some readings of questions, and (ii) give a natural explanation under which circumstances the mention-some reading is preferred.

Key words: decision theory, resolving underspecified questions

1. Introduction

According to most approaches towards questions, the meaning of a question is its set of possible (complete) answers. Groenendijk and Stokhof (1982) argue that an answer to a *wh*-question should be *exhaustive*, and should *mention all* the relevant individuals, while Hamblin (1973) assumes that in answering a *wh*-question one only needs to *mention some* positive instance, which seems particularly convincing for a question like *Where can I buy an Italian newspaper?* More recently it has been argued that *wh*-questions are in general *ambiguous* between a mention-some and a mention-all interpretation. But what, then, is the contextual parameter that resolves the ambiguity?

In this paper, I will propose that the *decision problem* of the questioner is crucial here. The intuition behind this proposal is the natural assumption that we only ask questions to receive some particular kind of information; the kind of information that would help to resolve the *decision problem* that the questioner faces. By relating questions to decision problems, I (i) show how we can measure the utilities of both mention-all and mention-some readings of questions, and (ii) give a natural explanation under which circumstances the mention-some reading is preferred.

2. Questions as Sets of Answers

According to Hamblin (1958, 1973), we answer a question by making a statement that expresses a proposition. Just as it is normally assumed that you know the meaning of a declarative sentence when you know under which circumstances this sentence is true, Hamblin argues that you know the meaning of a question when you know what counts as an appropriate answer to the question. Taking both assumptions together, this means that the meaning of a question as linguistic object (interrogative sentence) can be equated with the set of propositions that would be expressed by the appropriate linguistic answers. This gives rise to the problem what an appropriate linguistic answer to a question is.

According to almost all formal analyses of questions it is assumed that a *yes/no*-question like *Does somebody walk?* has only two appropriate answers: *Yes*, i.e., *Somebody walks*; and *No*, *Nobody walks*. Although polar questions have two appropriate answers, it is clear that only one of these two answers can be *true*. This means that with respect to each world a *yes/no*-question simply expresses a proposition: the proposition expressed by the true appropriate answer in that world. If we represent a *yes/no*-question simply by a formula like $?A$, where A is a first-order formula, and assume that $[[A]]^w_g$ denotes the truth value of A in w with respect to assignment function g, the proposition expressed by question $?A$ in world w is:[1]

$$[[?A]]^E_{w,g} = \{v \in C : [[A]]^v_g = [[A]]^w_g\}$$

We might call the above proposition the *extension* of question $?A$ in world w. To determine the *intension* of the *yes/no*-question, we simply abstract away from the real world:

$$[[?A]]^I_g = \lambda w.\{v \in C : [[A]]^v_g = [[A]]^w_g\}$$

Notice that this function from worlds to propositions is simply equivalent to the following set of propositions:

$$[[?A]]^I_g = \{\{v \in C : [[A]]^v_g = [[A]]^w_g\} \mid w \in C\}$$

and that this set of propositions *partitions* the worlds in C.

Given this analysis of polar interrogative sentences, the question arises what the meaning of a *wh*-question is; i.e., what counts in a world as an appropriate true answer to a question like *Who walks?* Hamblin (1973) makes the following proposal:

> [...] a question sets up a choice-situation between a set of propositions, namely, those propositions that count as answers to it. [...] we shall regard 'who walks' as denoting a set, namely, the set whose members are the propositions denoted by 'Mary walks', 'John walks', ... and so on for all individuals (p. 48).

Notice that in distinction with the above approach towards *yes/no*-questions, a *wh*-question might have more than one true appropriate answer according to Hamblin's analysis. In the above quote Hamblin talks only about single *wh*-questions, but we obviously don't need to restrict ourselves to them, and can analyze multiple *wh*-questions in a similar way. Let us assume that if A is an (open) first order formula and $\vec{x}$ the sequence of variables $x_1, \ldots, x_n$, we will represent (multiple) *wh*-questions by formulae like $?\vec{x}A$. Its Hamblin-intension (HI) is then given by the following *function* from worlds to the set of propositions that correspond to the set of true answers of the question in that world, or the equivalent *set* below that, where $\vec{d}$ denotes an n-ary sequence of objects.

$$\begin{aligned}[[?\vec{x}A]]_g^{HI} &= \lambda w. \left\{\left\{v \in C : [[A]]_g^v[^{\vec{x}}/_{\vec{d}}] = 1 \ \&[[A]]_g^w\left[^{\vec{x}}/_{\vec{d}}\right] = 1\right\} \mid \vec{d} \in D^n\right\} \\ &= \left\{\left\{v \in C : [[A]]_g^v[^{\vec{x}}/_{\vec{d}}] = 1 \ \&[[A]]_g^w[^{\vec{x}}/_{\vec{d}}] = 1\right\} \mid \vec{d} \in D^n \ \& \ w \in C\right\}\end{aligned}$$

Notice that the intention of a question according to Hamblin's analysis does *not* form a partition, because several elements of the set might overlap each other. This, of course, is due to the fact that according to Hamblin a *wh*-question might have more than one true appropriate answer in a world. But this means that a *wh*-question leaves to the answerer in several worlds a non-trivial *choice* how to answer the question. This choice will turn out to be important later.

We might call the above function the *intension* of a *wh*-question. To determine the *extension* of a question, we simply apply the function to the actual world. The extension of the *wh*-question in world w will then be $[[?\vec{x}A]]_g^{HI}(w)$, which is equal to

$$[[?\vec{x}A]]_{w,g}^{HE} = \left\{\left\{v \in C : [[A]]_g^v[^{\vec{x}}/_{\vec{d}}] = 1 \ [[A]]_g^w[^{\vec{x}}/_{\vec{d}}] = 1\right\} \mid \vec{d} \in D^n\right\}$$

Notice that this set is the set of *true* answers to a *wh*-question, and for John to know who walks it seems reasonable to demand that the set of worlds that represents his knowledge state in w has to be a subset of an element of the extension corresponding to the question.

Groenendijk and Stokhof (1982) have argued, however, that knowing for one individual who walks *that* he walks is not enough for John to know who walks. They claim that to know the answer to the question *Who walks?* John needs to know of *each* single individual *whether* he or she walks. In general, Groenendijk & Stokhof argue that John knows in world w the answer to the question that is represented by $?\vec{x}A$ if and only if the set of worlds that represents his knowledge-state is a subset of the denotation of $[[?\vec{x}A]]_{w,g}^E$:

$$[[?\vec{x}A]]_{w,g}^E = \{v \in C \mid \ [[\lambda\vec{x}A]]_g^v = [[\lambda\vec{x}A]]_g^w\}$$

where the lambda term $\lambda\vec{x}A$ denotes the following set of n-ary sequences with respect to world w and assignment function g:

$$[\lambda\vec{x}A]]_g^w = \left\{\vec{d} \in D^n \mid [[A]]_g^w[^{\vec{x}}/_{\vec{d}}] = 1\right\}$$

This above denotation might be called the *extension* of a question. To determine the corresponding *intension* we can, as always, simply abstract from the world. What results is the following *function* from worlds to propositions, or, equivalently, the *set* of propositions below:

$$\begin{aligned}[[?\vec{x}A]]_g^I &= \lambda w.\{v \in C \mid [[\lambda\vec{x}A]]_g^v = [[\lambda\vec{x}A]]_g^w\} \\ &= \{\{v \in C \mid [[\lambda\vec{x}A]]_g^v = [[\lambda\vec{x}A]]_g^w\} \mid w \in C\}\end{aligned}$$

Notice that this set of propositions gives rise to a *partition* of the state space C. The intension of a question is a set of mutually exclusive propositions thought of as the set of all alternative *exhaustive* answers to the question.

Groenendijk and Stokhof's analysis of questions has a number of nice properties not shared by the analysis of Hamblin (1973), nor by Karttunen's (1977) that is built on it. First of all, on their assumption that the extension of a question is a proposition, they can straightforwardly explain why questions can freely be conjoined with declaratives when embedded under verbs like *know*. In particular, to account for *wh*-complements like *John knows who came to the party*, they don't need to postulate two separate verbs of knowledge, as Karttunen (1977) had to. Second, their analysis has the consequence that not only single and multiple *wh*-questions have denotations of the same category, but that also *yes/no*-questions are analyzed in the same way as *wh*-questions. This has the important consequence, third, that they can give a general definition of *entailment* between all kinds of interrogatives simply by inclusion of intension. Thus, if Q and Q' are the intensions denoted by two questions, question Q is said to entail question Q' iff $\forall q \in Q : \exists q' \in Q' : q \subseteq q'$.[2]

3. Mention-some Questions and Human Concerns

Although the partition analysis of questions has a number of satisfying features, there are also some worries with the approach. The main worry, perhaps, is that according to Groenendijk and Stokhof's (1982) *mention-all* analysis of questions, it is predicted that each question has at most one true and appropriate answer in a world. Groenendijk and Stokhof argue that this is not so problematic for *choice readings* of questions, or for questions *coordinated* by a *disjunction*, for they just express more than one question.

Unfortunately, however, this way out doesn't seem to work satisfactorily for other uses of interrogative sentences that can be truly and appropriately answered in more than one way. These uses of interrogative sentences are the sentences that get interpreted as *mention-some* questions.[3]

First, there are the examples discussed recently by Beck and Rullmann (1999) which contain expressions that explicitly mark non-exhaustivity:

(1) Who, for example, came to the party?

It is clear that you can completely answer this question without giving the exhaustive list of people who came to the party.

Second, there are questions like (2)–(4) that typically get a mention-some reading, although they are not explicitly marked as such:

(2) Who has got a light?

(3) Where can I buy an Italian newspaper?

(4) How can I get to the station?

Just like (1), also these questions can intuitively be answered appropriately by mentioning just one individual, place, or manner, i.e., you don't have to give an exhaustive list of persons that have got a light, place where you can buy an Italian newspaper, or way to go to the station, respectively.

On first thought it might seem that also if interrogative sentences on their mention-some reading just express one question, they are not really problematic for a mention-all analysis of interrogative sentences. The reason is that one can claim that although mentioning all relevant individuals would completely answer the *wh*-question, in practice it normally suffices to give only a *partial* answer.

However, this proposal is unsatisfactory, because even if I just mention one individual, place, or way, I have intuitively resolved the question, i.e. satisfactorily answered question (1)–(3) or (4). Moreover, as argued for by Groenendijk and Stokhof (1984, p. 532), it seems that not all partial answers to a question like (2) intuitively count as satisfactory answers. Although

(5) John hasn't got a light

would be a partial answer to (2) when it has a mention-all interpretation, the answer does intuitively not resolve the question, in case the *wh*-phrase ranges over more than two individuals.

A further fact that suggests that the appropriateness to answer questions like (2)–(4) by mentioning just one individual/place/manner should not be explained by suggesting that giving a *partial* answer normally suffices is the fact that also when (2)–(4) are embedded, like in

(6) John knows where he can find an Italian newspaper

the *wh*-phrase has still typically the mention-some interpretation. That is, John needs to know only one (relevant) place where he can find an Italian newspaper in order for the sentence to be true.

We can conclude that it doesn't seem to be a good strategy to explain mention-some answers to *wh*-questions by saying that in linguistic practice partial answers normally suffice. The natural question that arises is whether we can say something more about the kind of circumstances under which a mention-some interpretation of a *wh*-question arises.

It has also been noted by several authors that mention-some interpretations of a question like *Where is a P?* typically arise only in peculiar situations. Situations where the questioner has a problem, or goal, and learning one (relevant) place where a *P* is would already suffice to solve the problem how to reach the goal. Question (3), for instance, is typically asked by an Italian tourist in Amsterdam with the goal of getting an Italian newspaper in mind. The tourist doesn't really mind where he can buy one, all he is interested in is where he should go to buy one. Mentioning just one element of the set of alternative 'equally best' places will perfectly resolve the question.

Whether a *wh*-question has a mention-some or a mention-all reading thus seems to depend on whether a, and what kind of, *human concern* lies behind the fact that the question was asked. Groenendijk and Stokhof (1984) even argue that *wh*-questions typically have a mention-all reading, and that they *can* only get a mention-some reading when some particular human concerns are at stake. They notice that when we embed (3) under verbs like *wonder*, *ask* or *know*:

(7) John *wonders/asks/knows* where he can buy an Italian newspaper

the embedded question can, and typically will, have a mention-some reading. However, *wh*-complements embedded under verbs which are not related to human concerns only seem to allow for a mention-all interpretation:

(8) Where you can get gas *depends* on what day it is.

(9) Who will come is partly *determined* by who is invited

In this section, we have seen that *wh*-questions can have a mention-some reading, and in particular when human concerns, or goals, are at stake. Whether a mention-some answer suffices to resolve the question or not depends on how *useful* the answer is. The usefulness of the answer, in turn, should be related to the goals of the questioner (cf. Ginzburg, 1995). In Section 4, I propose to make this precise by using tools of a well developed theory of rational behavior: Bayesian decision theory.

4. Utilities of Questions

4.1. Utilities of Questions Represented by partitions

In Savage's (1954) classical formulation of Bayesian decision theory, a distinction is made between states of the world, acts, and consequences; states of the world together with acts determine the consequences, each act-world pair has exactly one consequence, and the consequence of an act includes all features that are relevant to the decision maker's values. If we assume that the utility of doing action a in world w is $U(w, a)$, we can say that the *expected utility* of action a, $EU(a)$, with respect to probability function P is

$$EU(a) \quad = \quad \sum_w P(w) \times U(w, a)$$

Let us now assume that the agent faces a *decision problem*, i.e., he wonders which of the alternative actions in A he should choose. A decision problem of an agent can be modeled as a triple, $\langle P, U, A \rangle$, containing (i) the agents probability function, P, (ii) his utility function, U, and (iii) the alternative actions he considers, A. In case the set of worlds and the set of actions are finite, we might represent such a decision problem as a decision table like the one below:

World	Prob.	Actions		
		a_1	a_2	a_3
u	1/3	4	−2	0
v	1/3	1	7	1
w	1/3	1	4	4

In this decision problem there are three relevant worlds, u, v, and w, and three relevant actions, a_1, a_2, and a_3. For each of these actions we can now determine its expected utility. The expected utility of action a_1, for instance, is $(P(u) \times U(u, a_1)) + (P(v) \times U(v, a_1)) \times (P(w) \times U(w, a_1)) = (1/3 \times 4) + (1/3 \times 1) + (1/3 \times 1) = 4/3 + 1/3 + 1/3 = 6/3 = 2$. In a similar way we can see that the expected utility of action a_2 is 3, while action a_3 has a utility of 5/3.

The problem that the agent faces is which action he should perform. You might wonder why we call this a decision *problem*; should the agent not simply choose the action with the greatest expected utility, i.e., action a_2? Yes, he should, if he *chooses now*. But now suppose that our agent doesn't have to choose now, but has the opportunity to first receive some useful information by *asking question Q*.

Before we can determine the utility of Q, we first have to say how to determine the expected utility of an action conditional on learning some new information. For each action a_i, its conditional expected utility with respect to new proposition C, $EU(a_i, C)$ is

$$EU(a_i, C) = \sum_w P(w/C) \times U(a_i, w)$$

When John learns proposition C, he will of course choose that action in A which maximizes the above value. Then we can say that the utility value of making an informed decision conditional on learning C, UV(Learn C, choose later), is the expected utility conditional on C of the action that has highest expected utility:

$$UV(\text{Learn } C, \text{ choose later}) = \max_i EU(a_i, C)$$

In terms of this notion we can determine the value, or *relevance*, of the assertion C.[4] Referring to a^* as the action that has the highest expected utility according to the original decision problem, $\langle P, U, A\rangle$, i.e., $\max_i EU(a_i) = EU(a^*)$, we can determine the *utility value* of the *assertion* C, $UV(C)$, as follows:

$$\begin{aligned} UV(C) &= UV(\text{Learn } C, \text{ choose later}) - UV(\text{Learn } C, \text{ still do } a^*) \\ &= \max_i EU(a_i, C) - EU(a^*, C) \end{aligned}$$

This value, which in statistical decision theory (cf. Raiffa and Schlaifer, 1961) is known as the *value of sample information* C, $VSI(C)$, can obviously never be negative. In fact, it predicts that an assertion only has a positive utility value in case it influences the action that the agent will perform. And indeed, it doesn't seem unnatural to say that a cooperative participant of the dialogue makes a *relevant* assertion just in case it makes our agent *change* his mind with respect to which action he should take. In our above example, for instance, we can see that proposition $\{v\}$ has a utility value of 0, because the best action to perform after learning that v is the case is the same action as the one that would have been performed with respect to the original decision problem. Proposition $\{u, w\}$, on the other hand, has a positive utility, because if the agent would learn this proposition, he would change his mind and would perform action a_1 instead of action a_2. The utility value of $\{u, w\}$ would be $(1/2 \times (4 - (-2))) + (1/2 \times (1 - 4)) = 6/2 + (-3/2) = 3/2$. It seems not unreasonable to claim that in a cooperative dialogue the assertion that expresses $\{u, w\}$ is 'better' than the assertion that expresses $\{v\}$, because the former has a higher utility value.[5] In general, we can say that one assertion, A, is 'better' than another, B, just in case the utility value of the former is higher than the utility value of the latter, $UV(A) > UV(B)$.

In terms of the utility value of assertions/answers, we can now determine the utility values of *questions*. Suppose that question Q is represented by the partition $\{q_1, \ldots, q_n\}$. Then we can determine the *expected* utility value of a question, $EUV(Q)$ as the *average* utility value of the possible answers:

$$EUV(Q) = \sum_{q \in Q} P(q) \times UV(q)$$

Suppose that for our above example a question was asked that could be represented by the partition $\{\{u, w\}, \{v\}\}$. The expected utility value of the question would then be $(P(\{u, w\}) \times UV(\{u, w\})) + (P(\{v\}) \times UV(\{v\})) = (2/3 \times 3/2) + (1/3 \times 0) = 1$. Notice that because the utility values of assertions can never be negative, the above determined expected value of a question, which in statistical decision theory is known as the *expected value of sample information*, $EVSI$, can also never be negative. In fact, the value will only be 0 in case a^* *dominates* all other actions in A with respect to the question. An action dominates the other actions in A with respect to the question in case no answer to the question would have the result that the agent will change his mind about which action to perform, i.e., for each answer q it will be the case that $\max_i EU(a_i, q) = EU(a^*, q)$. In these circumstances the question really seems irrelevant and, assuming that asking questions is *cost free*, it seems natural to say that question Q is *relevant* just in case $EUV(Q) > 0$. It should be obvious that this measure function totally orders all questions with respect to their expected utility values.[6]

4.2. Utility of mention-some questions

We have seen above that mention-some interpretations leave a choice to the answerer in several worlds how to answer the question, because several answers in the intension of a question might overlap each other. Let us make this a bit more concrete by defining for a particular question whose intension can be represented by $\{\{u, w\}, \{v, w\}\}$ the different *answer rules* that represent the different ways the answerer could answer this question. Notice that in this simple example the answerer has a non-trivial choice only in world w, and, thus, there are only two answer rules relevant. According to the first answer rule, f, the answerer answers in both u and w by a sentence that expresses $\{u, w\}$, and in v he answers by a sentence that would express $\{v, w\}$. According to the second answer rule, g, on the other hand, the answerer answers only in u by a sentence that expresses $\{u, w\}$, but answers in both v and w by a sentence that expresses $\{v, w\}$. Notice that although the question represented by $\{\{u, w\}, \{v, w\}\} = Q$ is not a partition, if we look for each answer rule at the set of worlds in which a particular answer is given, this latter set will form a *partition*. For answer rule f, for instance, this latter set will be $\{f^{-1}(q) | \ q \in Q\} = \{\{u, w\}, \{v\}\}$.

For our question above only two answer rules were possible, but depending on how much overlap there exists between the possible answers a mention-some reading of a question can get, many more answer rules can be relevant. For a particular question Q', let us denote this set by F. Because the answerer might use any element in F, the questioner doesn't know which answer rule the answerer will actually use. Let us temporarily assume, however, that he *does* know which f will be used. In that case, the utility of choosing after he learned the answer q, UV_f(Learn q, choose later), should be determined as follows:

$$UV_f(\text{Learn } q, \text{ choose later}) = \max_i EU(a_i, f^{-1}(q))$$

In terms of this notion, we can now also define the utility of answer q, $UV_f(q)$:

$$\begin{aligned} UV_f(q) &= UV_f(\text{Learn } q, \text{ choose later}) - EU(a^*, f^{-1}(q)) \\ &= \max_i EU(a_i, f^{-1}(q)) - EU(a^*, f^{-1}(q)) \end{aligned}$$

This value will never be negative.

The problem that we want to solve is how to determine the utility of question Q that is not represented by a partition. I will do this in terms of the notion of $UV_f(q)$, and thus indirectly in terms of answer rules. Intuitively, to determine the utility of question Q we want to find out for each answer q in Q the probability that it will be given, i.e., $P(\text{get } q)$. The utility of the question Q is then equal to

$$EUV(Q) = \sum_{q \in Q} P(\text{get } q) \times UV(\text{`get } q\text{'})$$

where UV('get q') is the utility value of the proposition corresponding to the worlds in which answer q is given. If it is clear what the relevant answer rule is, f for instance, it is clear how to determine this utility: $UV(\text{`get } q\text{'}) = UV_f(q)$, and probability: $P(\text{get } q) = P(f^{-1}(q))$, i.e., the utility and probability of the set of worlds in which answer q will be given according to answer rule f. Because it is unclear, however, which answer rule is used, the probability that answer q will be given, $P(\text{get} q)$, cannot be set equal to $P(f^{-1}(q))$, but must rather be equated with $\sum_{f \in F} P(f) \times P(f^{-1}(q))$, assuming that the questioner's uncertainty about the answer rule that will be used can be quantified by probability function P.

If we agree on the proposal that the probability that answer q will be given, $P(\text{get } q)$, should be equated with $\sum_{f \in F} P(f) \times P(f^{-1}(q))$, the utility of our question Q with respect to the answer rules in F can be determined as follows:

$$EUV_F(Q) = \sum_{q \in Q} \sum_{f \in F} P(f) \times P(f^{-1}(q)) \times UV_f(q)$$

This formula looks rather complicated, but can, fortunately, be simplified considerably. First, notice that the probability that answer rule f will be chosen, $P(f)$, does not depend on any particular element of Q. This means that the above formula is equal to

$$EUV_F(Q) = \sum_{f \in F} P(f) \times \sum_{q \in Q} P(f^{-1}(q)) \times UV_f(q)$$

Remember now that $UV_f(q)$ is the same as $UV(f^{-1}(q))$, and that for each answer rule f, the set $\{f^{-1}(q)|\ q \in Q\}$ is a partition, even if Q itself is not. Let us call this partition Q^f. This partition can be thought of as the denotation of a mention-all question and has an expected utility value: $EUV(Q^f)$. Because this value $EUV(Q^f)$ is the same as $\sum_{q \in Q} P(f^{-1}(q)) \times UV_f(q)$, we can now redefine the value of $EUV_F(Q)$ also as

$$EUV_F(Q) = \sum_{f \in F} P(f) \times EUV(Q^f)$$

This redefinition is not only simpler to write down than the one we started out with, it also makes clear that we can easily compare the utilities of the mention-all and mention-some readings of *wh*-questions. This comparison is based on the easy to prove fact that if Q and Q' are the partitions denoted by two questions such that $Q \sqsubseteq Q'$, i.e., $\forall q \in Q: \exists q' \in Q': q \subseteq q'$, the expected utility of Q will be at least as high as the expected utility of Q', $EUV(Q) \geq EUV(Q')$. Notice that this means that the question that is represented by partition Q has a utility at least as high as the perhaps non-partitional question Q', when answer rule f is used, if the following condition is fulfilled:

$$\forall q \in Q: \ \exists q' \in Q': \ q \subseteq f^{-1}(q')$$

It is not difficult to see, fortunately, that for any answer rule this relation is guaranteed to exist between the partition induced by a mention-all reading of question $?\vec{x}A$ and the intension of the question on its mention-some reading, when for each sequence of individuals $\vec{d}$ in the relevant domain of 'quantification' of the sentence represented by $?\vec{x}A$, there exists a cell in $[[?\vec{x}A]]^I_g$ that denotes the set of worlds where $\vec{d}$ is the only element of $[[\lambda\vec{x}A]]_g$. Notice that when questions are interpreted with respect to an 'empty' context, this will always be the case.

Because the above fact holds for any arbitrary answer rule $f \in F$, also the *average* utility of the mention-some reading of the question, $EUV_F([[?\vec{x}A]]^{HI})$, can never be higher than the utility on the corresponding mention-all reading. From this we can conclude that under the above mentioned condition *the utility of a mention-some reading of a question can never be higher than the expected utility of the corresponding mention-all reading*.

This result is obviously relevant to understanding in which situations a *wh*-question has a mention-all or a mention-some reading. On the assumption that the questioner is rational and the answerer cooperative and knows the decision problem of the questioner, this suggests that a *wh*-question will usually get a mention-all interpretation, because usually the question has a utility that is strictly higher on this interpretation.

To see how things work, consider our earlier discussed example, again, where the alternative actions are a_1–a_3, and where the probabilities of the worlds u, v, and w, and the utilities of the actions in these worlds are given in the table below:

$Sick(x)$	World	Prob.	Actions		
			a_1	a_2	a_3
Only C	u	1/3	4	−2	0
Only D	v	1/3	1	7	1
C & D	w	1/3	1	4	4

Notice that this time I have assumed that in the three different worlds the property *being sick* has a different extension: in world u only individual C is sick, in world v only D, while in w both are sick. This means, obviously, that the *wh*-question *who is sick?*, represented by the formula $?xSick(x)$, should on its mention-all interpretation be represented as $\{\{u\}, \{v\}, \{w\}\}$. The expected utility of the question on this interpretation can then be calculated as $\sum_q P(q) \times UV(q) = (1/3 \times (4 - (-2)) + (1/3 \times 0) + (1/3 \times 0) = 6/3 = 2$. In section 4.1 we have determined the utility of question $\{\{u, w\}, \{v\}\}$ with respect to the same decision table, and we found that this question has a utility of 1. Notice that the question represented by $\{\{u\}, \{v\}, \{w\}\}$ is *finer-grained* than the question represented by $\{\{u, w\}, \{v\}\}$, and that—in accordance with what we have said above—the former question has indeed a *higher utility*, i.e., 2 versus 1.

Whereas the question represented by $?xSick(x)$ should be represented by the partition $\{\{u\}, \{v\}, \{w\}\}$ on its mention-all reading, on a mention-some interpretation the question can be represented by the following set of propositions: $\{\{u, w\}, \{v, w\}\}$. This set of propositions does not form a partition, because the answers overlap each other. Because the answers overlap each other, we should analyze the utility of this question in terms of *answer rules*.

Notice that just as in the example discussed in Section 4.2, the answerer has a non-trivial choice only in world w, and, thus, there are only two answer rules relevant. According to the first answer rule, f, the answerer answers in both u and w by a sentence that expresses $\{u, w\}$, and in v he answers by a sentence that would express $\{v, w\}$. According to the second answer rule, g, on the other hand, the answerer answers only in u by a sentence that expresses $\{u, w\}$, but answers in both v and w by an sentence that expresses $\{v, w\}$. We have seen above that although the question represented by $\{\{u, w\}, \{v, w\}\} = [[?xSick(x)]]^{HI}$ is not a partition, if we look for each answer rule at the set of worlds in which a particular answer is given, these latter sets will form partitions: for answer rule f this will be $\{\{u, w\}, \{v\}\}$, and for answer rule g it is $\{\{u\}, \{v, w\}\}$. We have determined already that the former partition has a utility of 1, and the utility of the latter partition is $(P(\{u\} \times UV(\{u\})) + (P(\{v, w\}) \times UV(\{v, w\})) = (1/3 \times (4 - (-2))) + (2/3 \times ((1/2 \times 0) + (1/2 \times 0))) = 1/3 \times 6 = 2$. Because each answer rule is equally likely, the *average* expected value of the question is $(1/2 \times EUV(\{\{u, w\}, \{v\}\})) + (1/2 \times EUV(\{\{u\}, \{v, w\}\})) = (1/2 \times 1) + (1/2 \times 2) = 3/2$. This value of question *Who is sick?* on its mention-some interpretation is *lower* than the corresponding utility value of the question on its mention-all interpretation. This is in accordance with our earlier findings where we saw that a *wh*-question on its mention-some reading will never have a higher utility than the question on its mention-all interpretation.

5. Decision Problem as Contextual Parameter

Notice, however, that our above result does not rule out the possibility that in some particular situations the utility of the mention-some interpretation of a question will be equal to the utility of the corresponding mention-all reading. I claim that due to pragmatic reasoning, exactly in these circumstances the interrogative sentence will get a mention-some reading. The reason is that providing a mention-some answer causes less *effort* than providing a mention-all answer.[7]

Pragmatics can be seen as the study of the interaction between context and utterance. A context should represent enough information to be able to determine both *what* is said (or meant) by an utterance, and whether it was used *appropriately*. We have seen above that the decision problem of the agent who asks a question is the crucial contextual parameter that helps to determine whether the interrogative sentence was used appropriately, i.e., whether the question was *relevant* in its context of interpretation. In van Rooy (1999) I argued that the decision problem of the questioner is also the crucial contextual parameter to determine what it takes for an assertion to resolve the question. Just like for other contextual parameters, also the interaction between decision problem (i.e., the relevant contextual parameter) and interrogative used might go in two directions. If you don't know the decision problem, i.e., the *intentions* of the speaker, you might learn something (by *accommodation*) about it from the interrogative sentence used. For linguistic applications of our framework, however, we will concentrate ourselves in this paper on the other side of the interaction. If you *do* know the relevant decision problem of the questioner, you typically will be able to find out what it takes to resolve a question.

Suppose now that a question is used that allows for several interpretations. Which of those interpretations was actually intended by the questioner? The answer is simple: the interpretation with the highest utility. On the assumption that we (more or less) know the questioner's decision problem, we can calculate for each interpretation of the question its expected utility. On the assumption that it is common ground that the speaker is rational, and thus a utility maximizer, the hearer can infer that the question—interrogative sentence—should have the most relevant/useful interpretation with respect to the questioner's decision problem.

I have claimed above that in some particular situations the utility of the mention-some and mention-all readings of *wh*-questions coincide, and that in these situations it suffices for the answerer to give only a mention-some answer. In these situations the question receives a mention-some interpretation in order to minimize effort. As a typical example where this is the case, consider the table below.

$?xBIN(x)$	world	P	n	s
Only N	u	1/3	6	0
Only S	v	1/3	0	6
N & S	w	1/3	4	4

In this example, we consider three relevant worlds where the extension of the predicate *places where you can buy an Italian newspaper* differ: world u, where you can buy an Italian newspaper only in the North, world v, where you can buy one only in the South, and world w, where you can buy one at both places. The decision problem also contains two relevant actions: action n which denotes the action of walking north; and action s which denotes the action of walking south. The decision table represents a situation where (i) the agent has no preference for learning that he can only buy an Italian newspaper in the N(orth) or only in the S(outh), because walking n(orth) and s(outh), respectively, would in those cases have an equal utility, and (ii) it's indifferent to him to walk either n(orth) or s(outh) when he learns that he can buy a paper in both parts of the city. I claim that this is the typical kind of situation in which the relevant *wh*-question can receive a mention-some reading: in those situations the question can intuitively be resolved equally well by a mention-some answer as by a mention-all answer.

It turns out that in these situations also the utilities of the mention-some and mention-all readings of the question coincide. In our example above, for instance, the utilities of the two readings of the question *Where can I buy an Italian newspaper?*, modeled in the formal language by $?xBIN(x)$, turn out to be 2 for both. This is easy to see for the utility of the question on its mention-all reading, while it can also be easily checked that the mention-some reading has an expected utility of 2 with respect to each of its two answer rules, and thus also has an *average* utility of 2. Because the mention-some answer is (known to be) equally useful as the mention-all one, and shorter, the interrogative will get the mention-some interpretation.

I have discussed above when a mention-some reading arises of a non-embedded question in terms of the decision problem that the questioner faces. But it should be clear that the same reasoning can be used to determine when an embedded *wh*-question receives a mention-some reading. The only difference is that this time it need not be the decision problem of the questioner, or speaker, that is relevant, but it can also, and typically will, be the decision problem that the agent denoted by the subject of the embedding clause faces.

6. Conclusion

Following an idea of van Rooy (1999), I have shown in this paper that by relating questions to decision problems we can determine the utility of unambiguous questions and use it to resolve the underspecification of interrogative sentences. In this earlier paper only questions were considered that give rise to partitions, i.e., *wh*-questions that give rise to *mention-all* readings. In this paper, however, I have shown that we can also determine the utility of questions on their *mention-some* readings. To determine these latter utilities, I have made crucial use of *answer rules*, rules that determine which answer will be given in which worlds. Making use of these rules, I have shown that the utility of a mention-some reading of a *wh*-question will never be higher than the utility of the mention-all reading of the same question, but that their utilities sometimes coincide. I have argued that these facts are relevant for linguistic applications, because these expected utilities of the different readings of the same interrogative sentence might help to determine the actual question asked by an interrogative sentence, or better perhaps, might determine under which circumstances a mention-some answer suffices to answer a *wh*-question satisfactorily.

Acknowledgements

I would like to thank Balder ten Cate, Paul Dekker, Jeroen Groenendijk and Frank Veltman for discussion, and the anonymous reviewers for JLC for valuable comments.

Notes

[1] Here, and elsewhere in this paper, I will assume that we analyze sentences with respect to a fixed model.

[2] Until now I have used the term 'question' only to denote interrogative sentences. From now on, however, I allow myself to be more liberal and will use it also for the intensions

interrogative sentences denote. I hope this double use of the notion will not lead to confusion.

[3] For a convincing argumentation that mention-some questions differ crucially from choice readings of questions (see Groenendijk and Stokhof, 1984).

[4] Notice that 'relevance' does not denote stochastic dependence here, as it standardly does in probability theory. The standard notion says that C is relevant for B iff $P(B/C) \neq P(B)$.

[5] It is important, however, not to think of v as a single world, but rather as a representative of lots of worlds that are similar enough to treat them as an equivalence class.

[6] In van Rooy (1999), I determined the utility of questions in a somewhat different way. It turns out, however, that the two ways of calculating the utilities of questions are equivalent (cf. van Rooy, to appear).

[7] In Van Rooy (ms.) I argue that interpretation results from balancing *relevance* with *effort*.

References

Beck S., Rullmann H. (1999) A flexible approach to exhaustivity in questions. *Natural Language Semantics*, 7, pp. 249–298.

Ginzburg J. (1995) Resolving Questions, I. *Linguistics and Philosophy*, 18, pp. 459–527.

Groenendijk J., Stokhof M. (1982) Semantic Analysis of Wh-Complements. *Linguistics and Philosophy*, 5, pp. 175–233.

Groenendijk J., Stokhof M. (1984) Studies in the Semantics of Questions and the Pragmatics of Answers. PhD Thesis, University of Amsterdam.

Hamblin C. L. (1973) Questions in Montague English. *Foundations of Language*, 10, pp. 41–53.

Karttunen L. (1977) Syntax and semantics of questions. *Linguistics and Philosophy*, 1, pp. 3–44.

Raiffa H., Schlaifer R. (1961) *Applied Statistical Decision Theory*. MIT Press, Cambridge, MA.

van Rooy R. (1999) Questioning to resolve decision problems. In: Dekker, P. (ed.), *Proceedings of the Twelfth Amsterdam Colloquium*, ILLC, Amsterdam.

van Rooy R. Comparing questions and answers: a bit of logic, a bit of language, and some bits of information. In: Groenendijk et al. (eds.), *Sources and Streams of Information*. CSLI, Stanford, to appear.

van Rooy R. Implicatures and Relevance in Bidirectional OT. University of Amsterdam, The Netherland. (MS.)

Savage L. J. (1954) *The Foundations of Statistics*. Wiley, New York.