QUESTIONS, PLANS, AND THE UTILITY OF ANSWERS

Anton Benz*
Syddansk Universitet, Kolding

Abstract

The goal of this paper is to derive a measure of utility for questions and answers from a game theoretic model of communication. We apply this measure to account for a number of judgements about the appropriateness of partial and mention–some answers, e.g. that a partial answers to a question can be as appropriate as a strongly exhaustive answer. Under the assumption that interlocutors are Bayesian utility optimisers we see questioning and answering as a two-person sequential decision problem with complete coordination of preferences. Our approach builds up on the work of A. Merin and R. v. Rooy on measures of *relevance*. We will compare it in detail with their ideas.

1 Introduction

Given a question $?x.\phi(x)$, what is the most useful answer? This question becomes especially interesting in connection with the problem of *partial* and *mention–some answers*. There are a number of judgments about the appropriateness of partial answers that seem to be due to their utility in a specific pragmatic context. In our examples, we write 'I' for the *inquirer*, and 'E' for the answering *expert*:

(1) I: Where can I buy an Italian newspaper?

In addition to requesting for information, this reveals a future plan of the inquirer, namely to buy an Italian newspaper. Lets assume that it doesn't matter for him where to buy it. The following answers are equally useful with respect to informativity:

(2) E: There are Italian newspapers at the station and at the Palace but nowhere else. (GS)

There are Italian newspapers at the station. (A)

There are Italian newspapers at the Palace. (B)

A and B are example for *mention–some* answers. All three seem to be equally useful. Answer A is not inferior to answer $A \wedge \neg B$:

E: There are Italian newspapers at the station but none at the Palace.

If E knows only that $\neg A$, then $\neg A$ is an optimal partial[1] answer:

E: There are no Italian newspapers at the station.

---

*This paper presents ongoing work that was done in part during my time at the ZAS Berlin; most of it at Syddansk Universitet in Kolding.

[1]Henceforth, I use *partial* answer as a cover term for both, mention-some and partially resolving answers.

In: Cécile Meier, Matthias Weisgerber (ed.) 2004: *Proceedings of the Conference "sub8 – Sinn und Bedeutung". Arbeitspapier Nr. 177, FB Sprachwissenschaft, Universität Konstanz, Germany*

The fact that the answers *GS*, *A* and *B* are equally *useful* for pursuing the inquirer's plan of buying an Italian newspaper seems to account for their being equally *appropriate* as answers. Our problem is to find a game theoretic model for the communication situation that provides for a measure of utility of answers and can account for our intuitive judgments about their quality.

If a question or answer is to be called *useful*, then there must be an end to which it can contribute as a means. Hence, we see the activity of asking and answering as embedded in a pragmatic situation where the inquirer follows a plan that specifies his ends. Under the assumption that interlocutors are Bayesian utility optimisers we see questioning and answering as a two-person sequential decision problem with complete coordination of preferences. The goal of our investigation is to derive an appropriate measure of usefulness from this game theoretic model of communication. There may be additional pragmatic principles that rule out some of the answers; e.g. the Gricean principle of manner would lead to a preference for *A* and *B* over *GS*. We concentrate on the aspect of utility only. This aspect is linguistically important as it captures at least a substantial part of what is called *relevance*. We don't discuss whether the Bayesian principle of optimising expected utility explicates the Gricean principle of relevance or not. Nor do we claim that it provides an explanation for the different uses of the adjective 'relevant'. But, of course, our investigation is closely related to some game and decision theoretic explications of *relevance*, and indeed deeply indebted to them. We build up on the work of A. Merin (1999) and especially R. v. Rooy (2001; 2003; 2003a; 2003b). Merin measures the relevance of a proposition *E* by its ability to make us believe that a certain hypothesis *H* holds. The appropriate pragmatic situation is one where a speaker wants to convince his addressee of the truth of *H* and seeks for evidence that is most effective for this purpose. According to van Rooy, the basic pragmatic situation is one where an inquirer faces a *decision problem* and seeks for information to resolve it. This type of problem has been well studied in decision theory and van Rooy derives a measure for the relevance of *answers* in terms of their ability to influence the inquirer's decision. We build up on van Rooy but in contrast we consider questioning and answering as a sequential *two*–person game and use backward induction for deriving a solution for the expert's and inquirer's decision problems.

2 Partial Answers — Pragmatic Background Assumptions

What counts as an answer to a question? Does the set of answers depend only on the question itself or does it also depend on the inquirer's underlying reasons for asking the question? There has been a controversial debate about these questions and we don't intent to decide them here[2]. This section should make clear our background assumptions and our main motives for adopting them. Little depends on these assumptions in our formal model, and whatever does so can easily be reformulated such that it fits to other assumptions.

Following Gronendijk & Stokhof (1984) we identify the set of answers to a question $?x.\phi(x)$ with the set of all *strongly exhaustive answers*. E.g. if the question is 'Who came to the party yesterday?', and only *John*, *Jane* and *Jeff* came, then the strongly exhaustive answer states that exactly they and nobody else came. If Ω is a set of possible worlds with the same domain D, and $[\![\phi]\!]^v$ denotes the extension of predicate ϕ in v, then a strongly exhaustive answer to question $?x.\phi(x)$ is a proposition of the form $[v]_\phi := \{w \in \Omega \,|\, [\![\phi]\!]^w = [\![\phi]\!]^v\}$; i.e. it collects all worlds where predicate ϕ has the same extension. In our example, it collects all worlds where exactly John, Jane and Jeff came to the party. The set of all possible answers is then given by

[2]See (Groenendijk & Stokhof, 1997)[Sec. 6.2.3] for a short survey of positions regarding *mention–some* interpretations.

$[\![?x.\phi(x)]\!]^{GS} := \{[v]_\phi \mid v \in \Omega\}$. This approach poses a problem for partial answers: They are not elements of $[\![?x.\phi(x)]\!]^{GS}$, hence no answers at all.

We don't try to show that this is the best approach; we just state our main motivations for following Groenengijk & Stokhof. It has been noted that partial answers are possible only if the question is embedded in a situation where they are subordinated to an inquirer's goal. If a question is asked only for gathering information, i.e. in a pragmatically neutral context, then a strongly exhaustive answer is expected:

(3) **a)** Which animals have a good sense of hearing?
b) Where do coral reefs grow?
c) When do bacteria form endospores?

Without an explanation for the possibility of partial answers, this observation alone would not suffice to justify an identification of the set of answers with the set of strongly exhaustive answers. So, we need in addition a pragmatic explanation for partial answers. Let us consider a situation where asking a question is subordinated to further ends:

(4) *Somewhere in the streets of Berlin...*

I: I want to take the next train to Potsdam. Where can I buy a ticket?
a) *E: Lists all places where to buy a ticket*/At the main station/At this shop over there.
b) *E:* Come with me! (*Takes him to the next ticket-shop*)
c) *E:* (*Hands him a ticket*)
d) *E:* There are no controllers on the trains today.

The responses in a) are partial answers. The response in b) contributes to a goal (*Get to a ticket–shop* (G_2)) immediately super–ordinated to the goal of *getting to know a shop that sells tickets* (G_1). The third option in c) contributes to a goal (*Getting a ticket* (G_3)) which is again super–ordinated to the plan of buying a ticket. The response in c) contributes to a project (G_4) that is again super-ordinated to getting a ticket. We wouldn't call the responses in b) and c) *answers*. A more appropriate name is probably *reaction*. In b) it is a mixture of a verbal command and an action, in c) a pure action. But we may replace both by pure assertions, e.g.:

(4) b') *E:* I go to a ticket–shop right now.
c') *E:* I've already bought a ticket for you.

We assume that a question $?x.\phi(x)$ itself introduces the immediate goal of providing the strongly exhaustive answer (G_0). Writing the sub-ordination relation as $<$ we find in Example **(4)** that this immediate goal is embedded in a hierarchy of goals $G_0 < G_1 < G_2 < G_3 < G_4$. We might call such a hierarchy a *plan*. The basic assumption that explains the possibility of responses as in **(4)** is: Super-ordinated goals can override the immediate goal of providing a strongly exhaustive answer. Partial answers differ from other verbal responses by the relative distance of their goal from the basic goal G_0. They contribute to a goal that is directly super–ordinated to it. Hence there is only a gradual difference in how good the responses in **(4)** b'), c') and d) are if we evaluate them as answers. This is our main motivation for identifying the set of proper answers with the set of strongly exhaustive answers.

As we address only the case of partial answers, we concentrate on situations where there is only one goal super–ordinated to the goal of providing the strongly exhaustive answer. The goals are represented by an utility function. A natural way to do this is by setting $U(v,a) := 1$ if we reach the goal in situation v after execution of action a, and $U(v,a) = 0$ if we don't reach it. If in Example **(1)** a is the act of *going to the station* and v a world where there are Italian newspapers at the station, then act a leads to success, and hence $U(v,a) = 1$. Of course, utility measures can represent more fine–grained preferences over the outcomes of actions; e.g. if the inquirer wants to buy an Italian newspaper but prefers to buy it at the Palace because it's closer to his place, then this can be represented by different values for buying Italian newspapers at the station and at the Palace.

3 Merin's Measure of Relevance

Before I go to present my model of questioning and answering that puts the principle of optimising utility in its centre, I first discuss two approaches that introduce game and decision theoretic explications of the Gricean principle of relevance. In addition to utility, relevance measures the (psychological) impact of an assertion on the addressees believes. The Gricean principle of relevance is, of course, a natural candidate for explaining our judgments about the appropriateness of various partial answers. Hence, game and decision theoretic explications of this principle are of immediate relevance to our investigation.

Merin derives his measure of relevance of assertions from measures of the relevance of experimental data in empirical science. The fact that the barometer is rising (E) provides evidence that the weather is becoming sunny. We can see the situation as a competition between two hypotheses: (H) The weather will be sunny, and $(\overline{H})$ The weather will be rainy. For simplicity we may assume that H and $\overline{H}$ are mutually exclusive and cover all possibilities. E, the rising of the barometer, does not necessarily imply that H, but our expectations that the weather will be sunny are much higher after learning E than before. This change of degree of belief can be captured by conditional probabilities. Let P^i represent the given expectations before learning E, i.e. P^i is a probability distribution over possible states of the world (in context i). Let $P^{i'}$ represent the expectations obtained from epistemic context i when E, and nothing else but E, is learned. Modelling learning by conditional probabilities we find that $P^{i'}(H) = P^i(H|E)$, where $P(H|E) := P(H \cap E)/P(E)$ for $P(E) \neq 0$, the probability of H given E. With Bayes' rule we get:

$$P^{i'}(H) = P^i(H|E) = P^i(H) \cdot (P^i(E|H)/P^i(E)). \tag{e.1}$$

$\overline{H}$ denotes the complement of H. Then learning E influences our beliefs about $\overline{H}$ in the same way as it influences our beliefs about H: $P^{i'}(\overline{H}) = P^i(\overline{H}|E)$. We find:

$$P^{i'}(H)/P^{i'}(\overline{H}) = P^i(H|E)/P^i(\overline{H}|E) = (P^i(H)/P^i(\overline{H})) \cdot (P^i(E|H)/P^i(E|\overline{H})). \tag{e.2}$$

Hence, $\log(P^{i'}(H)/P^{i'}(\overline{H})) = \log(P^i(H)/P^i(\overline{H})) + \log(P^i(E|H)/P^i(E|\overline{H}))$. We can see the term $\log(P^i(E|H)/P^i(E|\overline{H}))$ as a measure for the ability of E to make us believe H.

Merin (1999) transfers this measure from empirical sciences to the communication situation. In its new domain we can see $\log(P^i(E|H)/P^i(E|\overline{H}))$ as the (possibly negative) *argumentative force* of E to make the addressee believe that H. Consider a situation where the speaker wants to convince the hearer that H is the case. If P^i represents the epistemic state of the hearer, then an assertion E is the more effective, or relevant, the bigger $\log(P^i(E|H)/P^i(E|\overline{H}))$.

Merin defines *relevance* as a relation between a probability function P representing expectations in some given epistemic context i and two propositions: a proposition H, the *hypothesis*, and a proposition E, the *evidence*. This leads to the following definition[3]:

Definition 1 (Relevance, Merin) *The* relevance $r^i_H(E)$ *of proposition E to proposition H in an epistemic context i represented by a conditional probability function $P^i(.|.)$ is given by* $r^i_H(E) := \log(P^i(E|H)/P^i(E|\overline{H}))$.

Relevance can be positive or negative according to whether E influences the addressee to believe or disbelieve H. In the same way it favours H it disfavours $\overline{H}$, i.e. $r^i_H(E) = -r^i_{\overline{H}}(E)$. Hence, we can model the situation as a zero–sum game between hypotheses H and $\overline{H}$. This fits into Merin's outlook that sees competition and the aim to convince the communication partner of some fact H as the dominant features of conversation.

We don't follow Merin in this respect. For the situations described in Example **(1)** it seems to be more appropriate to model it by a game of complete coordination, i.e. a game where the inquirer's and expert's payoffs completely coincide. But this does not decide about Merin's general attitude, and a discussion of whether we should see conflict or coordination at the bottom of conversation needs more care and space than is available here. What is of more immediate importance is the fact that Merin measures relevance from the perspective of the *receiver* of the information. We will argue below that we need to switch to the *provider's* perspective in order to get the appropriate measure of the utility of answers.

Who's probability is P? It is the purpose of an assertion to influence the expectations of the addressee, hence P must represent the subjective probabilities of the receiver of information, or if the measure is used by the speaker, it must be the subjective probability that the speaker ascribes to the hearer. Approximately, we can identify the addressee's perspective with the *common ground*. In experimentation, the *speaker* is nature and the scientist performing the experiment is the *hearer*. Hence, in both cases, in scientific experimentation and in communication, relevance is defined from the receiver's perspective, i.e. information E is the more relevant the more it influences the receiver's expectations about some hypothesis H. We will argue that in answering a question the dominant goal is not to change the inquirer's, i.e. the receiver's, expectations but to provide information that maximises the expert's, i.e. provider's, expectations about how much E increases the chances of I's success. We will see that van Rooy, who directly addresses questioning and answering using a decision theoretic model, too measures the relevance of an answer from the receiver's perspective. We think that this is the main reason for its inadequateness.

4 Van Rooy's Measure of Relevance

Why do we ask questions? Because we want to resolve a decision problem! That is van Rooy's answer[4]. We will follow him in our analysis of situations like **(1)** which allow for partial answers in Section 6.

Let us first consider whether we can apply Merin's measure for the relevance of assertions to questioning and answering situations. If the inquirer I asks whether ϕ, then we can set $H := \{v \in \Omega \,|\, v \models \phi\}$, and $\overline{H} := \{v \in \Omega \,|\, v \not\models \phi\}$. Assuming that the answering expert applies Grice' principles and selects a proposition with maximal relevance, he has to select a proposition A as answer that maximally affects the inquirer's expectations[5]. Of course, such an answer may

[3](Merin, 1999), Definition 4.
[4](v. Rooy, 2003a, p. 727)
[5]Where we measure an answer's *pure relevance* by the absolute of $r^i_H(A)$; compare Merin (1999), Definition 5.

be highly misleading, even if it has to be truthful. In case of a question like '*Who came to the party yesterday*', we have to consider many competing hypotheses, in fact, all the strongly exhaustive answers in the sense of Groenendijk & Stokhof. It should be possible to generalise Merins approach as long as the set of answers is countable. It becomes more of a problem if the questioning is embedded in a decision problem where the inquirer has to choose between several alternative actions with results that bear different value for him. In this case we can't just measure the amount of information provided by an assertion; we also need to consider the expected gain of profit. Van Rooy's idea was to look at the communicative situation as a problem of decision theory and thereby to derive a criterion for the relevance of questions and answers.

Lets consider an example. An oil company has to decide where to build a new oil production platform. Given the current information it would invest the money and build the platform at a place off the shores of Alaska. An alternative would be to build it off the coast of Brazil. So the ultimate decision problem is to decide whether to take action a and build a platform off the shores of Alaska, or take action b and build it off the shores of Brazil. Now, should the company invest time and explore the off shore fields of Alaska and Brazil more thoroughly before deciding about its actions? If yes, then the company has to find the most efficient way to do it. This type of situation has been thoroughly studied in statistical decision theory[6]. Lets simplify the situation and assume that investigating the oil fields goes without costs. We can capture the essentials of the situation by the following model:

Let Ω be a set of states, $\mathcal{A}$ a set of actions, $U : \Omega \times \mathcal{A} \longrightarrow \mathbf{R}$ an utility measure, and P a probability measure for Ω. Then, the *expected utility* of an action a given P is defined by:

$$EU(a) = \sum_{v \in \Omega} P(v) \times U(v,a). \tag{e.3}$$

The effect of learning a proposition A is again modelled by conditional probabilities. The *expected utility after learning A* is given by:

$$EU(a,A) = \sum_{v \in \Omega} P(v|A) \times U(v,a). \tag{e.4}$$

What a manager wants to have is a criterion that tells him whether or not it is reasonable to investigate the off shore fields before finally deciding the question where to build the platform. As he is a Bayesian utility maximiser, additional explorations are only rational if he can expect that they lead him to choose an action with higher payoff than the action that he would choose now. It can only be higher if newly learned information can induce him to change his decision to build the oil platform off the shores of Alaska, i.e. if it changes his decision for action a. This leads to the following definitions of *relevance*: A proposition A is relevant if learning A induces the inquirer to change his decision about which action a to take. Let a^* denote the action where the expected payoff relative to information represented by P is maximal. Then the *utility value*[7] of proposition A is defined as:

$$UV(A) = \max_{a \in \mathcal{A}} EU(a,A) - EU(a^*,A). \tag{e.5}$$

A is relevant for the decision problem if $UV(A) > 0$. Exactly then the inquirer has a decisive reason to choose another action than a^*. The *expected utility value* of an investigation is then defined by:

$$EUV(Q) = \sum_{q \in Q} P(q) \times UV(q); \tag{e.6}$$

[6] See e.g. (Raiffa & Schlaifer, 1961, Ch. 4), (Pratt et. al., 1995, Ch. 14).

[7] Compare e.g. (Raiffa & Schlaifer, 1961, Sec. 4.5) and (Pratt et. al., 1995).

where Q is the set of all possible results of the investigation. It is reasonable to do additional investigations before finally deciding if $EUV(Q)$ is positive. So we can say that investigating is *relevant* if $EUV(Q) > 0$. Utility value UV and expected utility value EUV are defined from the investigator's perspective. Metaphorically speaking, we can call an experiment a *question*, and a result an *answer* to it. The answering *person*, nature, is not providing information with respect to the investigator's decision problem. There is only one real person involved in this decision model, namely the inquirer. Nature shows oil, or doesn't show oil, according to whether there is oil where the exploration drilling takes place. It does not show it *in order to contribute* to a decision problem, or because it thinks that this is *relevant*. The model does not predict that nature will only give relevant answers, and it does not even say that this were desirable. E.g. assume that there is indeed a very large oil field in the area near Alaska where the company wanted to build the platform given its old information, and a very small oil field in the Brazilian area. If the exploration drilling confirms that the original decision was right, then this is, according to our criterion, irrelevant. Only if by some bad luck the drilling in the Brazilian area gives rise to the hope that there is more oil than in Alaska, we got relevant information.

In (2001; 2001a) van Rooy introduced (e.5) as a measure for the relevance of answers[8]. I hope it became clear, that I don't follow him here. The whole model is a model for a one–person decision problem. The relevance of information is evaluated only from the inquirer's perspective. Hence it is not a trivial claim that this approach provides a measure that can be used by the answering expert E to select the best answer. But, of course, we may try to turn the model into a model for a two-person game. This makes it necessary to reinterpret the formulas above. So we ask: *Who's* probability could P be? There are three possibilities:

1. It is the inquirer's subjective probability.
2. It is the expert's subjective probability.
3. It is the subjective probability that E assigns to I.

Alternatives 1. and 2. are unsatisfactory. If 1., then (e.7) cannot be applied by E. If we assume that (a) the expert can only give answers that he believes to be true, and if we define (b) '*Expert believes q*' by $P_E(q) = 1$, then 2. implies that *any* answer will do as long as E believes it to be true. In order to turn the model into a model for a two-person game we have to choose interpretation 3.[9]. In this case (e.5) advises the answering expert only to choose answers that can make I change his decision.

(5) Assume that it is common knowledge between I and E that there are Italian newspapers at the station with probability $2/3$, and at the Palace with probability $1/3$. What should E answer if he is asked **(1)**: *Where can I buy an Italian newspaper*? According to the initial epistemic state, I decided to go to the station. Lets consider three possible answers: (A) There are Italian newspapers at the station; (B) There are Italian newspapers at the Palace, and $(A \wedge B)$. All three should turn out to be equally relevant but some calculation shows that B is the *only* relevant answer according to (e.5).

[8]"And indeed, it seems natural to say that a cooperative participant of the dialogue only makes a *relevant* assertion in case it makes John *change* his mind with respect to which action he should take. It also seems not unreasonable to claim that in a cooperative dialogue 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, $UV(A) > UV(B)$." (v. Rooy, 2001a, p. 78), emphasises are van Rooy's. 'John' refers to an inquirer in a previous example.

[9]Of course, that's van Rooy's intended interpretation.

This is clearly not intuitive. The point here is the same as in the oil-drilling example. Van Rooy therefore[10] replaces (e.5) by (e.7) in his later papers:

$$UV(A) = \max_{a \in \mathcal{A}} EU(a,A) - \max_{a \in \mathcal{A}} EU(a). \quad \text{(e.7)}$$

(e.7) gives the advice: '*Increase the hopes of the inquirer as much as you can*!' This fixes the problem with Example **(5)** but it's easy to see that we run into a similar problem with *negative* information: Assume that in the scenario of Example **(5)** E knows that there are no Italian newspapers at the station $(\neg A)$; in this case (e.7) implies that $\neg A$ is not relevant because it does not increase the inquirer's expectations. This seems to be quite unintuitive. But the problem can be easily fixed again by taking the absolute $|\,|$ of the right side of (e.7). And again, this runs into problems. An answer that increases, or changes, the hopes of the inquirer as much as possible is not necessarily a good answer:

(6) Assume a scenario like that in **(5)**. There is a strike in Amsterdam and therefore the supply with foreign newspapers is a problem. It is common knowledge between I and E that the probability that there are Italian newspapers at the station is slightly higher than the probability that there are Italian newspapers at the Palace. Now E learns that the Palace has been supplied with British newspapers but not with Italian ones. In general, it is known that the probability that Italian newspapers are available at a shop increases significantly if the shop has been supplied with foreign newspapers. What should E answer when asked: *Where can I buy an Italian newspaper*?

Some calculation shows that, according to (e.7), E should answer that the Palace has been supplied with foreign newspapers. The same holds for the improved version of (e.7) with the absolute difference. As the probability that there are Italian newspaper at the Palace, *given* that the Palace has been supplied with foreign newspaper, is much higher than the assumed probability for there being Italian newspapers at the station, this answer should lead the inquirer to go to the Palace. But this is the wrong choice as there are no Italian newspapers at the Palace. A good answer should maximise the inquirer's chances for real success, and not maximally increase or change his expectations about success.

In its form of (e.7) the criterion is very close to Merin's criterion of relevance. It may be the proper generalisation of Merin's approach for cases where the speaker does not only want to influence the hearer subjective probabilities but also his expectations about payoffs. So it should be better understood as a measure of the *argumentative force* of an assertion.

As indicated above, I see the main methodological flaw of this approach in its attempt to apply a model for a one-person decision problem to a real communication situation involving two persons. The expectations of the answering expert about the real state of affairs are treated as irrelevant. Of course, this is understandable as the model is derived from a theory that accounts for the value of experimental data where *questions* are directed to nature and where *answers* can only be evaluated from the experimentator's perspective.

[10] See (v. Rooy, 2001a, Sec. 4.3) for a comparable example; but this phenomenon has nothing to do with whether or not we see the situation as a game of complete coordination or of conflict. In (2001) van Rooy considers only scenarios with trivial probability distribution, i.e. where all possible states of affairs have equal probability. See also (v. Rooy, 2003, Sec. 3.1) and (v. Rooy, 2003a, Sec. 3.3).

5 Van Rooy's Order of Relevance

In (v. Rooy, 2003a, Sec. 3.1) van Rooy introduces an *order of relevance* as a simplified version of the measure of relevance introduced before. It is used to define the *set of answers* to a question. This is necessary in order to provide the semantics for embedded interrogatives as in *John knew who came to the party*, or *John knew where to buy an Italian newspaper*. In (v. Rooy, 2003b) this approach is generalised such that arbitrary orders of relevance are covered; i.e. it provides for the definitions of *optimal propositions* and *set of answers* for arbitrary orders of relevance. We restrict our discussion here to (v. Rooy, 2003a, Sec. 5.2)[11], and we discuss it only in as far as to whether it provides a justification of our judgments on appropriateness of partial and strongly exhaustive answers. Although the new order of relevance is introduced as a special case of the order of relevance defined by (e.7), I think that it is interesting on its own. It has an independent intuitive basis and prefers answers that eliminate more possible choices of actions in the inquirer's decision problem.

As mentioned before, Groenendijk & Stokhof define the set of answers to be the set of strongly exhaustive answers:

$$[\![?x.\phi(x)]\!]^{GS} = \{\{v \in \Omega \,|\, [\![\phi]\!]^v = [\![\phi]\!]^w\} \,|\, w \in \Omega\}.$$

Van Rooy identifies the set of answers with the set of *relevant* answers. One answer is more relevant than another if it helps more to resolve the inquirer's decision problem (v. Rooy, 2003a, p. 753). The decision problem consists in choosing an action from a set of actions $\mathcal{A}$. For each $a \in \mathcal{A}$ we can define the set of worlds where a is optimal:

- $a^* = \{v \in \Omega \,|\, \neg\exists b \in \mathcal{A}\, U(v,a) < U(v,b)\}$;
- $\mathcal{A}^* = \{a^* \,|\, a \in \mathcal{A}\}$.

As v. Rooy writes, the ordering relation on propositions induced by their utility value under certain conditions comes down to the claim that a proposition A is better to learn than a proposition B if A eliminates more cells from $\mathcal{A}^*$ than B. For some special cases this leads to the following order of relevance:

$$\phi(g') >_{\mathcal{A}^*} \phi(g) \text{ iff } \{a^* \in \mathcal{A}^* \,|\, a^* \cap [\![\phi(g')]\!] \neq \emptyset\} \subset \{a^* \in \mathcal{A}^* \,|\, a^* \cap [\![\phi(g)]\!] \neq \emptyset\}. \qquad \text{(e.8)}$$

Here, ϕ has two arguments: $\phi(v)(g)$ means that group g is such that '$\phi(g)$' holds in v. Lets consider our question: *Where can I buy an Italian newspaper*? We can consider an answer as specifying a *group* of places where it is possible to buy them; e.g. the answer '*At the palace and at the station*' says that the group $\{Palace, Station\}$ is in the actual world in the extension of the predicate *Place-where-Inquirer-can-buy-Italian-newspapers* (ϕ). With help of the order of relevance defined in (e.8) the most relevant group in each world is determined by:

$$[\![\mathrm{Op}(\phi)]\!] = \{\langle v,g\rangle \,|\, \phi(v)(g)\,\&\,\neg\exists g'(\phi(v)(g')\,\&\,\phi(g') >_{\mathcal{A}^*} \phi(g))\}. \qquad \text{(e.9)}$$

This leads van Rooy to define the set of answers as:

$$[\![?x.\phi(x)]\!]^R = \{\{v \in \Omega \,|\, g \in \mathrm{Op}(\phi)[v]\} \,|\, \exists w \in \Omega\, g \in \mathrm{Op}(\phi)[w]\}. \qquad \text{(e.10)}$$

I.e. a proposition is a possible answer if it is a set of worlds v such that there exists a world w and an optimal group g for w such that g is also optimal in v.

[11] This is mainly because (v. Rooy, 2003b) is only available as preprint. The reader is encouraged to check Section 3 there.

(7) My decision problem might be, for example, to find out which way is best for me to go to get an Italian newspaper. It could be, for instance, that the *best* way to buy an Italian newspaper is at the station in u, at the palace in v, and that buying one at the station and at the palace is equally good in w. (v. Rooy, 2003a, p.753)

In this case we get $Op(\phi)(w) = \{Palace, Station\}$[12]. The answer set for '*Where can I buy an Italian newspaper?*' is then: $[\![?x.\phi(x)]\!]^R = \{\{u,w\},\{v,w\}\}$[13]. We are here only concerned with whether or not this provides a justification for the judgment that the partial answers '*At the station*' (A) and '*At the Palace*' (B) are equally relevant with respect to information as the exhaustive answer '*At the station and at the Palace*' $(A \wedge B)$. Intuitively, all three answers should be equally good. $A \wedge B$ is more complex, so A or B should be preferred, but this needs an additional pragmatic principle, Grice' Principle of Manner. If $[\![?x.\phi(x)]\!]^R$ is the set of answers the addressee of the question can choose from, then $A \wedge B$ is not even an answer. But this point may turn out to be not important. What is more relevant for our concerns is whether or not there is a way to extend the order of relevance in order to cover answers that include negations as in '*At the station but not at the Palace*'. As soon as an answer entails the negation of A or B it eliminates more possibilities than either A or B. Hence, it should turn out that $A \wedge \neg B$ is more relevant than A, B or $A \wedge B$. But this contradicts our initial judgment that all these answers are equally good with respect to utility. Given the intuitions that underlie (e.8) this generalisation would be natural. This shows that it is not so much important to eliminate unsuccessful options in order to solve a decision problem but to show actions that are successful.

6 Measuring the Utility of Answers and Questions

As the discussion of Merin's and van Rooy's approaches did show, it is essential to take into account the perspectives of the interlocutors. As there are two interlocutors involved in questioning and answering, the inquirer and the addressee of the question, we conclude that we need to model it as a two person game. Under the assumption that interlocutors are Bayesian utility optimisers we see questioning and answering as a two-person sequential decision problem with complete coordination of preferences. This fits well within a dialogue theory that sees collaboration towards joint goals at the heart of communication. The best known such theory is that of Herbert H. Clark (1996). He predominantly analyses dialogue in terms of joint projects. This implies that knowledge about each other, and especially the notion of *common ground*, gets some prominence. Fortunately, we don't need mutual knowledge in our communication model, but we represent some knowledge about the knowledge of others. We follow van Rooy and see questioning and answering situations that allow for partial answers as subordinated to a final decision problem of the inquirer. Hence, we find three successive decision problems:

Inquirer I Asks		Expert E answers		I decides for action		Evaluation
↓		↓		↓		↓
•	$\xrightarrow{Q?}$	•	$\xrightarrow{A}$	•	$\xrightarrow{a}$	•
↑		↑		↑		↑
subject. expectations of I		subject. expectations of E		subject. expectations of I		utility measure
$\langle \Omega_I, P_I \rangle$		$\langle \Omega_E, P_E \rangle$		$\langle \Omega, P \rangle$		$U(v,a)$

[12] This is van Rooy (2003a, p. 753).
[13] This is again van Rooy's calculation.

We denote the inquirer again by I, and the answering person by E. As before, let Ω be a fixed set of states of the environment, $\mathcal{A}$ a set of actions, $U : \Omega \times \mathcal{A} \longrightarrow \mathbf{R}$ an utility measure. The model will represent (a) I's final beliefs about the world, (b) expectations of E about I's beliefs about the state of the environment, and (c) expectations of I about E's answering situation.

Why do we want to represent E's expectations about I's beliefs? Consider again a scenario like that in **(6)**.

(8) There is a strike in Amsterdam and therefore the supply with foreign newspapers is a problem. It is common knowledge between I and E that I has a clear preference to buy his newspapers at the station. Now E learns that the Palace has been supplied with Italian newspapers and he knows that both, the station and the Palace, get their newspapers from the same supplier. As the supplier favours none of his customers, the probability that the station too got Italian newspapers is quite high. Now, E knows that I thinks that the Palace and the station probably get their Italian newspapers from different sources, hence the fact that the Palace got them does not indicate to I that the station got them too. We assume that E should only say what he believes to be true, i.e. we assume that only propositions A with $P_E(A) = 1$ are admissible as answers. If he does not take into account I's beliefs, then it seems the best answer he can give to the question '*Where can I buy an Italian newspaper*?' is '*The Palace got Italian newspapers today*'. But this would induce I to go to the Palace, although E knows that I's expected utility is higher if he goes to the station.

(a) The inquirer I has to decide in the final situation about which action a to take according to his expectations about the actual states of affairs. We assume that his decision does not depend on what he believes that the expert E believes. Hence we can represent his epistemic state by a pair $\langle \Omega, P \rangle$, where P represents his knowledge about the actual world.

I's Decision Situation: It is given by a probability space $\langle \Omega, P \rangle$ and an utility measure U.

(b) We assume that the answering expert E wants to maximise I's final success. Hence, E's payoff is identical with I's. E has to choose his answer in such a way that it optimally contributes towards I's decision. In general, he has to calculate how I will decide if he provides him with some information A. Therefore, our model must take into account what E expects about what I knows. Hence, we represent the possible states of affairs in E's decision situation by pairs $\langle v, \langle \Omega, P \rangle \rangle$, where $v \in \Omega$ is a possible state of the environment and $\langle \Omega, P \rangle$ a possible information state of I. The states $\langle \Omega, P \rangle$ are intended to represent I's knowledge *before* he learned E's answer. I assume that all probability spaces that I will introduce are finite.

E's Answering Situation:

- $\Omega_E := \{\langle v, \langle \Omega, P \rangle \rangle \mid v \in \Omega,\, P \text{ a probability measure on } \Omega\}$.
- P_E is a probability measure on Ω_E.

Note that any probability measure P_E on Ω_E induces a probability measure on Ω by $P_E(v) := P_E(\{\langle u, \langle \Omega, P \rangle \rangle \in \Omega_E \mid v = u\})$.

(c) The inquirer I in his initial situation may take into account what he believes that E believes. Hence we represent the possible states in I's initial decision situation by pairs $\langle v, \langle \Omega_E, P_E \rangle \rangle$, where $\langle \Omega_E, P_E \rangle$ is a possible information state of E.

I's Questioning Situation:

- $\Omega_I := \{\langle v, \langle \Omega_E, P_E \rangle \rangle \mid v \in \Omega,\, P_E \text{ a probability measure on } \Omega_E\}$.

- P_I is a probability measure on Ω_I.

Again, we get a probability measure on Ω by $P_I(v) := P_I(\{\langle u, \langle \Omega_E, P_E \rangle\rangle \in \Omega_I \mid v = u\})$.

How to determine the utility of answers? The idea is to calculate backward from the final evaluation situation, i.e. by backward induction. This is indeed the most natural and straightforward solution to our decision problems.

Calculating Backward Expected Utilities

The final Decision Situation: The expected utility of an action $a \in \mathcal{A}$ is calculated according to (e.3):

$$EU_{\langle \Omega, P \rangle}(a) := \sum_{v \in \Omega} P(v) \times U(v, a). \tag{e.11}$$

Hence E calculates I's expected utilities in situation $v = \langle w, \langle \Omega, P_v \rangle\rangle$ after learning A by:

$$EU_{\langle \Omega, P_v \rangle}(a, A) = \sum_{u \in \Omega} P_v(u|A) \times U(u, a). \tag{e.12}$$

Let $a_A^v := \text{indmax}_{a \in \mathcal{A}} EU_{\langle \Omega, P_v \rangle}(a, A)$. If it is not unique, we assume some mutually known tie breaking rule. According to our assumption, E's payoff function is identical with I's payoff function U, i.e. questioning and answering is a game with complete coordination. In order to maximise his own payoff, E has to choose an answer such that it induces I to take an action that maximises their common payoff.

The Answering situation: We use again (e.3) for calculating the expected utility of an answer $A \subseteq \Omega$:

$$EU_{\langle \Omega_E, P_E \rangle}(A) := \sum_{v \in \Omega_E} P_E(v) \times U(v, a_A^v). \tag{e.13}$$

We add here a pragmatic constraint: An answer is admissible only if $P_E(A) = 1$. This means that we only allow for answers that the expert E believes to be *true*. For $v = \langle w, \langle \Omega_E, P_E^v \rangle\rangle \in \Omega_I$ let $Adm(v) := \{A \subseteq \Omega \mid P_E^v(A) = 1\}$, the set of *admissible*, i.e. true, answers. This leads to the following definition of the set of optimal answers in situation v:

$$\text{Op}(v) = \{A \in Adm(v) \mid \forall B \in Adm(v)\, EU_{\langle \Omega_E, P_E^v \rangle}(B) \leq EU_{\langle \Omega_E, P_E^v \rangle}(A)\}. \tag{e.14}$$

If there are several optimal answers, then we assume again that E's choice from $\text{Op}(v)$ is determined (in a mutually known way). We denote this unique answer by A^v.

Going back to I's initial querying situation, we have to switch perspectives again. In order to calculate his expected utilities for questions, I has to take into account which action he would choose if he gets some information A. We denote this action by a_A. I can calculate for all possibilities $v = \langle u, \langle \Omega_E, P_E \rangle\rangle \in \Omega_I$ the answer A^v that is optimal from E's perspective given by $\langle \Omega_E, P_E \rangle$. Hence, he can conclude that in situation v he will be led to take action $a_{A^v} = \text{indmax}_{a \in \mathcal{A}} EU_{\langle \Omega, P_I \rangle}(a, A^v)$, where P_I is the probability measure induced on Ω and where we have to use the tie braking rule if a_{A^v} is not unique.

The Querying situation: The expected utility of a question Q can then be calculated by:

$$EU_{\langle \Omega_I, P_I \rangle}(Q) := \sum_{v \in \Omega_I} P_I(v) \times U(v, a_{A^v}). \tag{e.15}$$

$EU_{\langle \Omega_I, P_I \rangle}(Q)$ does not depend on Q. This is a consequence of calculating utilities with respect to a fixed decision problem. This is the point where our pragmatic backgrounds assumptions of Section 2 enter. There we claimed that a question introduces the immediate goal of providing a strongly exhaustive answer; this goal may be subordinated to further ends that provide for additional targets that can override the immediate goal of exhaustively resolving the question. These further ends may be given in the background, or they may be inferred from the question by help of some plan recognition mechanism. All these goals and further ends have to be represented in our model by the utility function U. Hence, in (e.15), U should have a subscript Q. Lets consider e.g.:

(9) **a)** Where can I buy an Italian newspaper? (Q)
b) Are there Italian newspapers at the station? (Q')

Let $\mathcal{A} = \{a, b\}$, the actions of going to the station and going to the Palace. We can represent the difference between Q and Q' by the assumption that $U_Q(w,a) = U_{Q'}(w,a)$ and $U_Q(w,a) = 1$ iff in w there are Italian newspapers at the station; $U_Q(w,b) = 1$ iff in w there are Italian newspapers at the Palace and $U_{Q'}(w,b) = 0$ for all w. How U_Q has to be defined in general given the common background and the inquirer's question lies outside of our investigation.

The Examples Reconsidered

Let us first consider the answers of Example **(2)**. Let $\mathcal{A} = \{a, b\}$, the actions of going to the station and going to the Palace. Let $A \subseteq \Omega$ be the set of worlds where there are Italian newspapers at the station, and $B \subseteq \Omega$ where they are at the Palace. Let $\overline{A}$ and $\overline{B}$ denote the respective complements. We represent the payoffs as follows: $U(w,a) = 1$ iff $w \in A$, $U(w,b) = 1$ iff $w \in B$.

If E knows that A, then A is an optimal answer:

$$EU_{\langle \Omega_E, P_E \rangle}(A) = \sum_{v \in \Omega_E} P_E(v) \times U(v, a_A^v) = \sum_{v \in A} P_E(v) \times U(v,a) = 1$$

No other answer can yield a higher payoff[14]. In the same way it follows that B is optimal if E knows that B. The same holds for $A \wedge B$ and the strongly exhaustive answer.

If E knows only $\neg A$, hence $P_E(\overline{A}) = 1$, then $\neg A$ is an optimal answer:

$$EU_{\langle \Omega_E, P_E \rangle}(\overline{A}) = \sum_{v \in \Omega_E} P_E(v) \times U(v, a_{\overline{A}}^v) = \sum_{v \in \overline{A}} P_E(v) \times U(v,b) = P_E(\overline{A} \cap B) = P_E(B).$$

If $P_E(C) = 1$, then for $v \in \Omega_E$ either $a_C^v = a$ or $a_C^v = b$. Let $B_C := \{v \in \Omega_E \mid a_C^v = b\}$. Then

$$EU_{\langle \Omega_E, P_E \rangle}(C) = \sum_{v \in B_C} P_E(v) \times U(v,b) = P_E(B_C \cap B) \leq P_E(B).$$

Here enters: $P_E(C) = 1 \Rightarrow P_E(C \cap A) = 0$. Hence, no other answer than $\neg A$ can be better[15].

Lets consider Example **(6)**. We use the same utility function as before. Let N denote the set of all $u \in \Omega$ where the Palace has been supplied with British newspapers. We model the epistemic states described in **(6)** by the following condition: For all $v = \langle u, \langle \Omega, P \rangle \rangle \in \Omega_E$ it holds that:

[14] There could be a problem if I believes that B and E has evidence that B is unlikely. Hence, the result stated above holds in full generality only if we assume in addition e.g. that E believes that I can't be convinced of B. In all calculations there are additional pragmatic assumptions that should be made explicit in a more rigorous presentation. E.g. we repeatedly have to assume that for conditional probabilities $P(u|A)$ it holds that $P(A) > 0$.

[15] It is important that I can only choose between actions a and b. The result holds even if $C = B$.

1. $P(A) > P(B)$ and $P(B \cap N) > P(A \cap N)$;
2. $P_E(A) > P_E(B)$, $P_E(A \cap N) = P_E(A)$ and $P_E(B \cap N) = P_E(B)$.

Is N a good answer? Let $v \in \Omega_E$:

$$EU_{\langle \Omega, P_v \rangle}(a, N) = \sum_{u \in N} P_v(u|N) \times U(u, a) = P(A \cap N).$$

and

$$EU_{\langle \Omega, P_v \rangle}(b, N) = \sum_{u \in N} P_v(u|N) \times U(u, b) = P(B \cap N).$$

Hence $a_N^v = b$. We get

$$EU_{\langle \Omega_E, P_E \rangle}(N) = \sum_{v \in N} P_E(v) \times U(v, b) = P_E(N \cap B) = P_E(B) < P_E(A)$$

It is easy to see that $EU_{\langle \Omega_E, P_E \rangle}(\Omega) = P_E(A)$. Hence, N cannot be the best answer. Let C be such that $P_E(C) = 1$. Let $A_C := \{v \in \Omega_E \mid a_C^v = a\}$ and $B_C := \{v \in \Omega_E \mid a_C^v = b\}$. Then:

$$EU_{\langle \Omega_E, P_E \rangle}(C) = \sum_{v \in A_C} P_E(v) \times U(v, a) + \sum_{v \in B_C} P_E(v) \times U(v, b) = P_E(A_C \cap A) + P_E(B_C \cap B) \quad \text{(e.16)}$$

Is it possible that E has a better answer than saying nothing although he does not know more about A and B than the inquirer? Well, it is.

(10) Imagine that there is a causal relation between who is delivering Italian newspapers this morning and whether or not there are newspapers at the station and Palace. E knows that a man named *van den Berg*[16] delivered them this morning. Assume that E knows that I can infer from who delivered the newspapers whether there are Italian newspapers at the station or at the Palace. In this case E has a better answer than saying nothing:

I: Where can I buy an Italian newspaper?

E: A man named 'van den Berg' delivered the newspapers this morning. (C)

In addition, this example shows that our theory does not say that the relevance measured from the receiver's perspective, i.e. the inquirer's, is irrelevant for determining the best answer. C carries no information for E, only for I. It only says that, in general, the inquirer's perspective is not enough. It takes two for communicating.

References

H.H. Clark (1996): *Using Language;* Cambridge.

J. Groenendijk, M. Stokhof (1984): *Studies in the Semantics of Questions and the Pragmatics of Answers;* Ph.D. thesis, University of Amsterdam.

J. Groenendijk, M. Stokhof (1997): *Questions;* in: J. v. Benthem, A. ter Meulen (eds.): *Handbook of Logic and Language*, Amsterdam, the Netherlands; pp. 1055-1124.

[16] I was told that — in spite of its complete flatness — this is the most common name in the Netherlands.

A. Merin (1999): *Information, Relevance, and Social Decisionmaking: Some Principles and Results of Decision-Theoretic Semantics;* In: L.S. Moss, J. Ginzburg, and M. de Rijke (eds.): *Logic, Language, and Information* Vol. 2; Stanford.

J.W. Pratt, H. Raiffa, R. Schlaifer (1995): *Introduction to Statistical Decision Theory;* The MIT Press, Cambridge, Massachusetts.

H. Raiffa, R. Schlaifer (1961): *Applied Statistical Decision Theory;* Harvard.

R. v. Rooy (2001): *Utility of Mention-Some Questions;* To appear in *Journal of Language and Computation*.

R. v. Rooy (2001a): *Questions and Answers in cooperative and non-cooperative settings;* in: P. Kühnlein, A. Newlands, H. Rieser (eds.): *Proceedings of the Workshop on Coordination & Action*, pp. 71–90.

R. v. Rooy (2003): *Quality and Quantity of Information Exchange; Journal of Logic, Language, and Computation* 12, pp. 523–451.

R. v. Rooy (2003a): *Questioning to Resolve Decision Problems; Linguistics and Philosophy* 26, pp. 727–763.

R. v. Rooy (2003b): *Questions and Relevance;* In: Questions and Answers: Theoretical and Applied Perspectives (Proceedings of 2nd CoLogNET-ElsNET Symposium, pp. 96–107.