Utility and Relevance of Answers

Anton Benz

February 16, 2005

1 Introduction

How to answer a question? If the inquirer asks it in order to make a decision about something, then a wide range of reactions can be appropriate. If asked '*Who of the applicants is qualified for the job?*', reactions may range from '*Only Müller and Schmidt*', '*At least Müller*', over '*Müller has working experience in this field*', '*Schmidt needs extra training*', to '*The younger ones show more enthusiasm*', or even '*The job needs an expert in PCF Theory*'. This paper divides into two parts. The goal of the first part is to derive a measure of utility for 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. Under the assumption that interlocutors are Bayesian utility optimisers we see questioning and answering as a two-person game with complete coordination of preferences. Our approach builds up on work by A. Merin and R. v. Rooij on decision theoretically formulated measures of *relevance*.[1] In the second part we study the relation between their approaches and our game theoretic model of answering. We are aiming for principled characterisations, and are especially interested in clarifying when and why we have to model this type of communication as a *two*–person game.

There are a number of judgements about the appropriateness of 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*. We use the following question **(1a)** as our main example:

(1) Somewhere in the streets of Amsterdam...

- **a)** I: Where can I buy an Italian newspaper?
- **b)** E: At the station and at the Palace but nowhere else. (SE)
- **c)** E: At the station. (A) / At the Palace. (B)

The answers b) and c) are equally useful with respect to conveyed information and the inquirer's goals. The answer in b) is called *strongly exhaustive*; it tells us for every location whether we can buy there an Italian newspaper or not. The answers in c) are called *mention–some* answers. In general, a mention–some answer like A is not inferior to an answer like $A \wedge \neg B$:

- **d)** E: There are Italian newspapers at the station but none at the Palace.

[1] I.e. (Merin, 1999) and quite a number of papers by R. v. Rooij listed in the bibliography.

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

e) E: There are no Italian newspapers at the station.

We call this type of answers *partial* answers.

In Section 2 we work out our model. We provide explicit definitions of the underlying structures because we need them for later comparison with work by Merin and v. Rooij. We will pay attention to how it incorporates Gricean Maxims. Specifically: (1) The model incorporates the *Cooperation Principle* as our games are games of perfect coordination, i.e. the answering expert E is always cooperative. (2) Throughout we will assume that the maxim of *Quality* cannot be violated, i.e. E can only answer with a proposition that she thinks to be true. (3) We ignore the *Maxim of Manner*, i.e. if two propositions turn out to be equally useful, then we treat them as equally good answers even if one of them needs a much more complex sentence to be expressed. The main difference shows up with respect to the *Maxims of Quantity* and *Relevance*. We replace them by the assumption that interlocutors are Bayesian utility maximisers.

The answers in **(1b)**–**(1e)** form only the core of phenomena that have to be explained. We study a number of examples in Section 3. An especially interesting group results from situations where the answering expert has to take into account the possibility of misleading expectations. An answer like '*Müller worked as a student in a software company*' may produce incorrect expectations if he did only subordinate jobs in the reception. In general, this type of examples is roughly characterised by (1) the existence of an answer C that favours but doesn't decide a certain hypothesis, (2) another answer C' which disfavours the same hypothesis, and the fact that E knows C and C', or, at least, believes C' to be very probable. We call answers like C *non–trivial* partial answers. We will show that they provide a principal problem for approaches that determine optimal answers according to a decision theoretically formulated measure of relevance. We prove that no such theory can be empirically adequate.

We present decision theoretic explications of the maxim of relevance in Section 4. The results about their relation to our game theoretic definition of best answers are presented in Section 5. From Theorem 5.3 it will follow that no decision theoretically formulated criterion for selecting maximally relevant answers can avoid selecting misleading answers.

2 Optimal Answers

Background

There has been a controversial debate about whether or not strongly exhaustive answers have a prominent status among the set of all possible answers. Groenendijk & Stokhof (1984) are most prominent defenders of the view that they constitute the basic answer, whereas other types of answers have to be accounted for pragmatically. For a constituent question like '*Who came to the party?*' it has to tell us for each person whether he or she came to the party or not. This is of some importance for the interpretation of embedded interrogatives: If Peter knows who came to the party, then Peter knows whether John came to the party, or whether Jeff came to the party, or whether Jane came to the party, etc. The set of all possible answers is then the set of all strongly exhaustive answers.[2] On the other side there are examples like 'Peter knows where to buy an Italian newspaper' which does not seem to imply that Peter knows whether he can buy an Italian newspaper at X, where X ranges over all kiosks in Amsterdam. The same difference shows up for the respective unembedded, or direct, questions. This leads to a position that sees questions as ambiguous or underspecified.[3]

It is not our aim to solve this controversy here. We just indicate how we like to position our approach in its context. Hence we state our background assumptions and our main motives for adopting them. But, in a technical sense, our game theoretic analysis of questioning and answering does not depend on these assumptions.[4]

Following Groenendijk & Stokhof (1984) we identify the set of answers to a question $?x.\phi(x)$ with the set of all *strongly exhaustive answers*. If we take this approach, then we have to find a pragmatic explanation for the possibility of mention–some answers. Our main motivation for adopting their view comes from the observation that only questions that are subordinated to special goals of the inquirer allow for mention–some answers. If a question is asked only for gathering information, i.e. in a pragmatically neutral context, then a strongly exhaustive answer is expected:

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

In situations where asking a question is subordinated to further ends we find a wide range of other reactions:

(3) *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 and where not to buy a ticket.*
b) E: At the main station / At this shop over there.
c) E: Come with me! (*Takes him to the next ticket-shop*)
d) E: (*Hands him a ticket*)
e) E: There are no controllers on the trains today.

The response in **(3a)** is the strongly exhaustive answer, those in **(3b)** are mention–some answers. The response in **(3c)** 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 response in **(3d)** contributes to a goal (*Getting a ticket* (G_3)) which is again super–ordinated to the plan of buying a ticket.

[2] 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 $[v]_\phi := \{w \in \Omega \mid [\![\phi]\!]^w = [\![\phi]\!]^v\}$; i.e. it collects all worlds where predicate ϕ has the same extension. The set of all possible answers is then given by $[\![?x.\phi(x)]\!]^{GS} := \{[v]_\phi \mid v \in \Omega\}$. This poses a problem for mention–some answers as they are not elements of $[\![?x.\phi(x)]\!]^{GS}$, hence not answers at all.

[3] For a short survey of positions regarding *mention–some* interpretations see (Groenendijk & Stokhof, 1997)[Sec. 6.2.3].

[4] Hence, the reader who disagrees with me on the relation between mention–some and mention–all answers will still get the full value out of the game theoretic model.

The response in **(3e)** contributes to a project (G_4) that is again super-ordinated to getting a ticket. We wouldn't call the responses in **(3c)** to **(3e)** *answers*. A more appropriate name is probably *reaction*.

Due to our assumption that strongly exhaustive answers are basic, 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 **(3)** that this immediate goal is embedded in a hierarchy of goals $G_0 < G_1 < G_2 < G_3 < G_4$. The following mechanism explains the possibility of responses as in **(3)**: Super-ordinated goals can override the immediate goal of providing a strongly exhaustive answer. Mention–some answers contribute to a goal that is super–ordinated to the basic goal G_0. The information conveyed by them is *optimal* with respect to this super–ordinated goal. It is the aim of this paper to precisely characterise this optimality in game theoretic terms. This provides the general idea how to derive the possibility of mention–some answers.

In a game theoretic model we can represent a goal by a 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 $u(v, a) = 1$. 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 is closer to his place, then this can be represented by assigning higher values to buying Italian newspapers at the Palace than to buying them at the station.

Finally, we should emphasise that we consider only direct questions, i.e. no syntactically embedded questions. If we know what is the optimal answer to a question Q?, then we do not necessarily know how to interpret it if it occurs as a syntactically embedded question.

(4) **a)** Peter knows where to buy an Italian newspaper.

b) Peter knows where to buy best an Italian newspaper.

These two sentences are not equivalent. If the set of optimal answers were identical with the meaning of the embedded sentence, then they should be equivalent. Hence, by determining optimal answers, we make no claim about embedded questions.

The Utility of Answers

As mentioned before, our background assumptions do not immediately affect the following analysis of partial and mention–some answers. The key–idea is to see questioning and answering as embedded in a decision problem. It is due to Robert van Rooij who explored it in quite a number of papers.[5] I see it as one of the most interesting contributions of game and decision theory to pragmatics until now. To motivate this move, we look at some examples. In **(1)**, the

[5] E.g. (v. Rooij, 2001, 2003,a,b). Why do we ask questions? Because we want to have some information. But why this particular kind of information? Because only information of this particular kind is helpful to resolve the decision problem that the agent faces. (v. Rooij, 2003a, p. 727).

inquirer has to decide where to go to in order to buy an Italian newspaper. Other examples that allow a mention–some answer are:

(5) **a)** (*In a job centre*) I am a computer expert. Where can I apply for a job?

b) I like to go skiing in the Alps. What places can you recommend?

c) (*In a job interview*) What are your qualifications?

In **(5a)**, the inquirer has to decide where to apply for a job. In **(5b)**, he has to decide where to go to for skiing; in **(5c)**, whether or not to employ a candidate. In each case there is a finite set of actions $\{a_1, \ldots, a_n\}$ and the inquirer asks for information that helps him to make an optimal choice among them. In **(1)** this set may be represented as {go-to(x) | x a newspaper kiosk in Amsterdam}; in **(5a)** as {send-application-to(x) | x a group of regional software companies}; in **(5b)** as {travel-to(x) | x a valley in the Alps}; in **(5c)** as {employ, not-employ}. The decision depends on the preferences of the decision maker over the outcomes of these actions, and on his/her information about the state of the world. We assume that there is a fixed set Ω that collects all possible states. If the decision maker does not have complete information, then he has to rely on his expectations about the world. We can represent them by the probabilities he/she assigns to the different possible states. In order to keep things simple, we assume that there are only countably many states of the world, i.e. that Ω is countable. In this case, a probability distribution is just a real valued function $P : \Omega \longrightarrow \mathbf{R}$ such that (1) $P(v) \geq 0$ for all $v \in \Omega$ and (2) the sum of all $P(v)$ equals 1. For sets $A \subseteq \Omega$ it is usual to set $P(A) = \sum_{v \in A} P(v)$. Hence $P(\Omega) = 1$. We collect these elements in a structure:

Definition 2.1 *A* decision problem *is a triple* $\langle(\Omega, P), \mathcal{A}, u\rangle$ *such that* (Ω, P) *is a countable probability space,* $\mathcal{A}$ *a finite, non–empty set and* $u : \Omega \times \mathcal{A} \longrightarrow \mathbf{R}$ *a function.* $\mathcal{A}$ *is called the* action *set, and its elements* actions*.* u *is called a* payoff *or* utility *function.*

A decision problem represents the inquirer's problem. By asking his question he makes it common knowledge. Again for simplicity, we assume here that his beliefs represented by the probability space (Ω, P) are mutually known, too. In order to indicate that a probability distribution represents the *inquirer*'s beliefs we write P_I.

How does the situation change if we include the answering expert in our model? The only parameter that we add to our formal structure is a probability distribution P_E that represents her expectations about the world:

Definition 2.2 *A* support problem *is a five–tuple* $\langle\Omega, P_E, P_I, \mathcal{A}, u\rangle$ *such that* (Ω, P_E) *and* (Ω, P_I) *are countable probability spaces, and* $\langle(\Omega, P_I), \mathcal{A}, u\rangle$ *is a decision problem. We call a support problem* well–behaved *if (1) for all* $A \subseteq \Omega$: $P_I(A) = 1 \Rightarrow P_E(A) = 1$ *and (2) for* $x = I, E$ *and all* $a \in \mathcal{A}$: $\sum_{v \in \Omega} P_x(v) \times u(v, a) < \infty$*.*

The first condition for well–behavedness is included in order to make sure that E's answers cannot contradict I's beliefs,[6] the second is there in order to

[6] See Fact 3.1 below on p. 8.

keep the mathematics simple. A support problem represents just the fixed static parameters of the answering situation. We assume that I's decision does not depend on what he believes that E believes. Hence his epistemic state (Ω, P) represents just his expectations about the actual world. E's task is to provide information that is optimally suited to support I in his decision problem. Hence, E faces herself a decision problem, where her actions are the possible answers. The utilities of the answers depend on the way how they influence I's final choice. We look at the dependencies in more detail. We find two successive decision problems:

Expert E answers		I decides for action		Evaluation
$\downarrow$		$\downarrow$		$\downarrow$
$\bullet$	$\xrightarrow{A}$	$\bullet$	$\xrightarrow{a}$	$\bullet$
$\uparrow$		$\uparrow$		$\uparrow$
expectations of E (Ω, P_E)		expectations of I (Ω, P_I)		utility function $u(v, a)$

We assume that the answering expert E is fully cooperative and 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. Due to our assumption that I's information is mutually known, E is able to calculate how I will decide. Hence, we represent the decision process as a sequential two-person game with complete coordination of preferences. We find a solution, i.e. optimal answers and choices of actions by calculating backward from the final outcomes.

The following model will be worked out using standard techniques of game and decision theory. We concentrate on *ideal* dialogue. By this we mean that all participants have only true beliefs and adhere to the Gricean maxims — as far as they are necessary. The *Cooperation Principle* e.g. is represented by the fact that we consider only games of pure coordination. We will introduce other maxims together with our analysis.

Calculating Backward Expected Utilities

First we have to consider the final decision problem of I. The probability P_I in our support situation is intended to represent his beliefs before E has given her answer. Hence, we have to say how an answer will change these beliefs. In probability theory the effect of learning a proposition A is modelled by *conditional probabilities*. The related learning model is known as *Bayesian learning*. Let H be any proposition, e.g. the proposition that there are Italian newspapers at the station; or that software companies x, y, z offer jobs for computer experts. H collects all possible worlds in Ω where these sentences are true. Let C be some other proposition, e.g. the answer given by our expert. Then, the probability of H *given* C, written $P(H|C)$, is defined by:

$$P(H|C) := P(H \cap C)/P(C). \tag{2.1}$$

This is only well–defined if $P(C) \neq 0$. Lets consider an example. I thinks that the probability that the station has Italian newspapers *and* that the owner of the newspaper store there is Italian is 1/50, i.e. $P_I(H \cap C) = \frac{1}{50}$. (2) He thinks that the probability that in general the owner is Italian is $P_I(C) = \frac{2}{50}$. It follows that in every second world where the owner is Italian, he must also sell Italian newspapers. If now I learns from E's answer that the owner is indeed Italian then he should belief that there are Italian newspapers at the station with probability 1/2: With (2.1) we calculate $P(H|C) = \frac{1}{50} : \frac{2}{50} = \frac{1}{2}$.

I's Decision Situation

How does I choose his action? It is standard to assume that rational agents try to maximise their expected utilities: Let $\langle(\Omega, P), \mathcal{A}, u\rangle$ be any decision problem. Then, the *expected utility* of an action a is defined by:

$$EU(a) = \sum_{v\in\Omega} P(v) \times u(v, a). \tag{2.2}$$

As the effect of learning a proposition A with $P(A) > 0$ is modelled by conditional probabilities, we get the *expected utility after learning* A by:

$$EU(a, A) = \sum_{v\in\Omega} P(v|A) \times u(v, a). \tag{2.3}$$

As we assumed that the decision maker I tries to maximise expected utilities by his choice, it follows that he will only choose actions that belong to $\{a \in \mathcal{A} \mid EU_{\langle\Omega,P_I\rangle}(a, A)\, \textit{is maximal}\}$. In addition we assume that I has always a preference for one action over the other, or that there is a mutually known rule that tells I which action to choose if this set has more than one element. In this case we can write a_A for this unique element. In short, we assume that the function $A \mapsto a_A$, for $P_I(A) > 0$, is known to E.

E's Decision Situation

According to our assumption, E's payoff function is identical with I's payoff function u, i.e. questioning and answering is a game of complete coordination (Principle of Cooperation). 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. We use again (2.2) for calculating the expected utility of an answer $A \subseteq \Omega$. With a_A defined as in the previous paragraph we get:

$$EU_E(A) := \sum_{v\in\Omega} P_E(v) \times u(v, a_A). \tag{2.4}$$

We add here a further Gricean maxim, the *Maxim of Quality*. We call an answer *admissible* if $P_E(A) = 1$. The Maxim of Quality is represented by the assumption that the expert E does only give admissible answers. This means that she believes them to be *true*. For a support problem $S = \langle\Omega, P_E, P_I, \mathcal{A}, u\rangle$ we set:

$$Adm_S := \{A \subseteq \Omega \mid P_E(A) = 1\} \tag{2.5}$$

Hence, the set of optimal answers for S is given by:

$$\mathrm{Op}_S = \{A \in Adm_S \mid \forall B \in Adm_S\, EU_E(B) \leq EU_E(A)\}. \tag{2.6}$$

3 Examples

We consider only well–behaved support problems $\langle \Omega, P_E, P_I, \mathcal{A}, u\rangle$, i.e. for all $A \subseteq \Omega : P_I(A) = 1 \Rightarrow P_E(A) = 1$. As mentioned before, the condition $P_I(A) = 1 \Rightarrow P_E(A) = 1$ entails that E's answers cannot contradict I's beliefs. More precisely, we find:

Fact 3.1 *Let* $(\Omega, P_E, P_I, \mathcal{A}, u)$ *be a given support problem, then the condition* $\forall A \subseteq \Omega\, (P_I(A) = 1 \Rightarrow P_E(A) = 1)$ *entails for all* $A, B \subseteq \Omega$*:*

1. $P_E(A) = 1 \Rightarrow P_I(A) > 0$,
2. $P_I(A|B) = 1 \,\&\, P_E(B) = 1 \Rightarrow P_E(A) = 1$.

We start with our main example **(1)**, i.e. with the question: *Where can I buy an Italian newspaper*? We first describe the general scenario before justifying the different types of answers.

We denote by a, b the actions of going to the station and going to the Palace. There may be other actions too. 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. We represent the payoffs as follows: For every possible action $c \in \mathcal{A}$ the utility value is either 1 (success) or 0 (failure); especially we assume that $u(v, a) = 1$ iff $v \in A$, else $u(v, a) = 0$; $u(v, b) = 1$ iff $v \in B$, else $u(v, b) = 0$.

Mention–Some Answers

We have to show that the mention–some answers are equally good as the strongly exhaustive answer (SE). Technically this means that we have to show that: if A, B and SE are admissible answers, then $EU_E(A) = EU_E(B) = EU_E(SE)$.

We start with answer A: If E knows that A, then A is an optimal answer. If learning A induces I to choose action a, i.e. if $a_A = a$, then the proof is very simple:

$$EU_E(A) = \sum_{v\in\Omega} P_E(v) \times u(v, a_A) = \sum_{v\in A} P_E(v) \times u(v, a) = 1.$$

Clearly, no other answer could yield a higher payoff. If we want to prove the claim in full generality, i.e. for all cases, may they be as complicated as they can be as long as our previously formulated restrictions hold, then we need some more calculation. We first note the following fact: Lets assume that I chooses after learning A an act c different from a, i.e. $a_A = c \neq a$. Then let C denote the set where action c is successful, i.e. $C = \{v \in \Omega \mid u(v, c) = 1\}$. Then either (i) $P_E(C) = 1$ or (ii) $P_E(C) < 1$. In the first case (i) it follows again that $EU_E(A) = 1$, and our claim is proven. Case (ii) leads to a contradiction by Fact 3.1: If I chooses c, then $EU_I(c|A) = \max_{c'\in\mathcal{A}} EU_I(c'|A) = EU_I(a|A) = 1$; hence $P_I(C|A) = 1$, and therefore $P_E(C) = 1$ in contradiction to (ii). It follows that only (i) is possible.

In the same way it follows that B is optimal if E knows that B. The same result follows for any stronger answer, including the strongly exhaustive answer SE, $A \wedge B$ or $A \wedge \neg B$. This shows that their expected utilities are all equal as long as they are admissible answers. We have no condition that represents the *Maxim of Manner*, hence all these answers are equally good and E can freely choose between them.

Partial Answers

We now turn to Example **(1e)**. Let $\overline{A}$ and $\overline{B}$ denote the complements of A and B. We assume here in addition that I can only choose between a and b, i.e. between going to the station and going to the Palace. We show: If E knows only $\neg A$, hence $P_E(\overline{A}) = 1$, then $\neg A$ is an optimal answer. We first assume that learning $\neg A$ leads I to choose action b, i.e. if I learns that there are no Italian newspapers at the station, then he will go to the Palace:

$$\begin{aligned} EU_E(\overline{A}) &= \sum_{v\in\Omega} P_E(v) \times u(v, a_{\overline{A}}) = \sum_{v\in\overline{A}} P_E(v) \times u(v, b) \\ &= P_E(\overline{A} \cap B) = P_E(B). \end{aligned}$$

Let C be any proposition. If $P_E(C) = 1$, then either $a_C^v = a$ or $a_C^v = b$; hence either $EU_E(C) = 0$ or $EU_E(C) = P_E(B)$. Here enters: $P_E(C) = 1 \Rightarrow P_E(C \cap A) = 0$. Hence, no other answer than $\neg A$ can be better. It is important that I can only choose between actions a and b. The result holds even if $C = B$.

What happens, if I doesn't choose b but a? This means that $0 = EU_I(a|\overline{A}) = \max\{EU_I(a|\overline{A}), EU_I(b|\overline{A})\}$, hence $EU_I(b|\overline{A}) = 0$. This entails that $P_I(B\cap\overline{A}) = 0 = P_E(B \cap \overline{A})$. Hence, E believes that there are neither Italian newspapers at the station nor at the Palace, hence no answer can rise expected utilities above 0. This concludes discussion of Example **(1e)**.

Non–trivial Partial Answers

Non–trivial partial answers will play a significant role when we discuss relevance based approaches. We consider the same general setting as before; especially the utility functions are defined as before, and I can only choose between a and b. Lets consider the following example:

(6) There is a strike in Amsterdam and therefore the supply with foreign newspapers is a problem. The probability that there are Italian newspapers at the station is slightly higher than the probability that there are Italian newspapers at the Palace, and it might be that there are no Italian newspapers at all. All this is common knowledge between I and E. Now E learns that (N) the Palace has been supplied with foreign newspapers. 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.

We model the epistemic states described in **(6)** by the following condition:

$$P_I(A) > P_I(B) \text{ and } P_x(B \cap N) > P_x(A \cap N) \text{ for } x = I, E. \tag{3.7}$$

As before, P_E describes E's beliefs when choosing her answer, i.e. *after* learning N, and P_I describes I's beliefs *before* learning E's answer. Is N an optimal answer? Lets first calculate I's reaction:

$$EU_I(a, N) = \sum_{v\in N} P_I(v|N) \times u(v, a) = P_I(A \cap N);$$

and

$$EU_I(b, N) = \sum_{v\in N} P_I(v|N) \times u(v, b) = P_I(B \cap N).$$

Hence, he will choose b, i.e. $a_N = b$. With (2.4) we get for E:

$$EU_E(N) = \sum_{v \in N} P_E(v) \times u(v, b) = P_E(N \cap B) = P_E(B) > P_E(A).$$

It is easy to see that for any answer C either $EU_E(C) = P(A)$ or $EU_E(C) = P(B)$. Hence, N is an optimal answer.

Now we change the scenario slightly:

(7) We assume the same scenario as in **(6)** but E learns this time that (M) the Palace has been supplied with *British* newspapers. Due to the fact that the British delivery service is rarely affected by strikes and not related to newspaper delivery services of other countries, this provides no evidence whether or not the Palace has been supplied with Italian newspapers.

The fact that M provides no evidence whether or not there are Italian newspaper at the station (A) or the Palace (B) means that $P_E(A) = P_E(M \cap A) > P_E(M \cap B) = P_E(B)$. I's epistemic state has not changed from **(6)**, hence we assume again that $P_I(A) > P_I(B)$ and $P_I(A \cap N) < P_I(B \cap N)$. What are the optimal answers? If E says nothing, i.e. if she answers Ω, then the expected payoff is $EU_E(\Omega) = P_E(A)$. Is there a better answer than saying nothing? Let C be such that $P_E(C) = 1$. Then either I will go to the station, i.e. choose a, or go to the Palace, i.e. choose b. Hence, $EU_E(C) = P_E(A)$ or $EU_E(C) = P_E(B) < P_E(A)$. This shows that I cannot provide any information that does better than Ω. Is N still an optimal answer? We find that answering with N leads I to go to the Palace, and therefore $EU_E(N) = P_E(B) < EU_E(\Omega)$. Hence, N cannot be a best answer. It would be misleading if E replied that *the Palace has been supplied with foreign newspapers.*

Now, lets finally consider:

(8) We assume the same scenario as in **(6)** where E learns that (N) the Palace got supplied with foreign newspapers but her intuition tells her that, if there have been Italian newspapers among them, then they are sold out before I can get there. Of course, this is only a conjecture of hers.

We assume here that her expectations are so strong that still $P_E(A) > P_E(B)$. As in **(6)**, we assume that $P_I(A) > P_I(B)$ and $P_I(A \cap N) < P_I(B \cap N)$. We find again that answering N, i.e. that the Palace got foreign newspapers, will induce I to do b. Again it follows that E thinks that this is misleading. Hence, again N is not an optimal answer.

In our model, the place of Grice' Maxim of Relevance was taken over by the assumption that interlocutors are Bayesian utility maximisers, i.e. by the assumption that rational agents choose actions that maximise their expected payoffs. This principle is somewhat alien to the linguistic pragmatic tradition, hence, we may ask: Isn't it possible to replace it again by the more familiar Gricean maxim? In order to answer this question we need a halfway precise formulation of this principle. We discuss in the next section some explications in decision theoretic terms. We prove in Section 5 that no possible decision theoretic explication of the Maxim of Relevance can explain examples **(6)**–**(8)**.

4 Relevance

The Gricean Maxim of Relevance is, of course, a natural candidate for explaining our judgements about the appropriateness of various partial and mention–some answers. Hence, game and decision theoretic explications of this maxim are of immediate interest for our investigation. Roughly, relevance measures the (psychological) impact of an assertion on the addressee's beliefs. In decision theory there is only one decision maker. He may be uncertain about the state of the world but there are no other players who's moves or beliefs he has to take into account.[7] In a proper game theoretic problem, the payoffs of one player's moves depend on the moves of the other players, and vice versa. Hence, the more general question behind the discussion of the next two sections is whether or not it is essential that we model questioning and answering as a *two*–person game.

We divide our discussion of explications of relevance into two sections. The first one addresses proposals that measure relevance in terms of the amount of information carried by an utterance, the second one on proposals that, in addition, take into account the expected utilities. The first one concentrates mainly on the approach by Arthur Merin (Merin, 1999), the second one on work by Robert van Rooij.[8]

Information Based Measures of Relevance

If the stock market is rising (D), it indicates that the economy prospers, and therefore probably the unemployment rate will sink. We can see the situation as a competition between two hypotheses: (H) The unemployment rate sinks, and ($\overline{H}$) the unemployment rate doesn't sink. H and $\overline{H}$ are mutually exclusive and cover all possibilities. D, the rising of the stock market, does not necessarily imply that H, but our expectations that the unemployment rate will sink are somewhat higher after learning D than before. We can say, that fact D is *positively relevant* for our belief that H. We can model the change of degree of belief by conditional probabilities as indicated previously, and based on them, it is possible to derive measures of relevance. We discuss here one proposal by Merin in more detail.

Merin (1999) defines *relevance* as a relation between a probability function P representing expectations in some given epistemic context $i = (\Omega, P)$ and two propositions: a proposition H, the *hypothesis*, and a proposition D, the *evidence*. This leads to the following definition:[9]

Definition 4.1 (Relevance, Merin) *The* relevance $r^i_H(D)$ *of proposition D to proposition H in an epistemic context i represented by a conditional probability function $P^i(.|.)$ is given by* $r^i_H(D) := \log(P^i(D|H)/P^i(D|\overline{H}))$.

Merin applies this measure to communication situations. In its new domain we can see $\log(P^i(D|H)/P^i(D|\overline{H}))$ as the (possibly negative) *argumentative force* of D to make the addressee believe that H. The details of the definition are not of concern here. $r^i_H(D)$ can be positive or negative according to whether D influences the addressee to believe or disbelieve H. In the same way it favours

[7] Sometimes *nature* is considered to be a second player.

[8] We concentrate on his earlier work (2001; 2003) and (2003a; 2003b).

[9] (Merin, 1999), Definition 4.

H it disfavours $\overline{H}$, i.e. $r^i_H(D) = -r^i_{\overline{H}}(D)$. In a situation where the speaker wants to convince the hearer that H an assertion D is the more effective, or relevant, the bigger $r^i_H(D)$. If $r^i_H(D) = 0$ then it neither favours nor disfavours any of the two hypotheses, and it is reasonable to call D irrelevant.

What is of concern for us now is only the fact that r^i_H takes as parameters the elements of a tuple $\langle \Omega, P, H, D\rangle$. Merin takes an argumentative attitude towards communication, i.e. he sees the aim of convincing the conversational partner of some hypothesis H as the basic goal of communication. For this purpose it is reasonable to choose the proposition that has the greatest impact on the addressee's beliefs, i.e. the most relevant proposition. This means that we can define by means of r^-_- a decision function R that selects for each context $\langle \Omega, P, H\rangle$ a proposition D with maximal argumentative force.

Let us consider how to 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 $H := \{v \in \Omega \mid v \models \phi\}$, and $\overline{H} := \{v \in \Omega \mid v \not\models \phi\}$. Assume we are in a job interview, I wants to know whether E is qualified for the job (H) or not ($\overline{H}$). Hence, he asks her about her qualifications. E has to find the strongest argument that is not already known to I for convincing him of H. Maybe, she didn't mention in her resume that (D) she worked as a student regularly in a similar company, which indicates that she knows the business. If this is indeed her strongest argument, then she should use it. But the selected answer may be highly misleading, even if it is truthful. E.g. it may be that E didn't mention this information in her resume because she worked only at the telephone switchboard.

If we decide to use a decision theoretic model for the conversational phenomenon we investigate, then this implies that one person plays the role that *nature* plays in an experiment, i.e. the decision function R does not depend on his preferences and expectations. In both cases, in scientific experimentation and in communication, relevance is then only defined from the receiver's perspective. If the measure of relevance is based only on information, it does not even take his utilities into account. Hence it shows its limits in situations like the following: Lets call our applicant Eve. We assume that Ω consists of four worlds $\{v_1, \ldots, v_4\}$ of equal probability. In world v_1 Eve is a highly qualified and experienced applicant who can start right at the job. In v_2 she is qualified but needs some more training. In v_3 and v_4 she is not qualified with only minor differences between the two states. I, who read her resume, asks a colleague, E, whether she knows more about this applicants abilities. Now assume that E knows that $D = \{v_2, v_3\}$. Is D relevant? If the decision maker learns D, then, using Merin's measure, it turns out that $r_H = \log(P(E|H)/P(E|\overline{H})) = 0$. Hence, D is irrelevant. But, intuitively, it is relevant for the decision maker to learn that the most favoured situation v_1 cannot be the case.

Utility Based Measures of Relevance

The last example shows that, in general, we have to consider the interlocutors preferences. Van Rooij'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.

Getting the results of the exploration drilling, the company has to decide whether to go ahead and follow their old plans and built the production platform in the north, or to redesign them and build it off Brazil. One heuristic says that information can only be relevant if it induces the company to choose an action that promises higher payoff than the action it would have chosen before getting this information. This heuristic 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, and is the *more* relevant the more it increases the inquirer's expectations. Let a^* denote the action where the expected payoff is maximal relative to the information available before drilling, represented by P. Then the *utility value*[10] of proposition A is defined as:

$$UV(A) = \max_{a \in \mathcal{A}} EU(a, A) - EU(a^*, A). \tag{4.8}$$

A is relevant for the decision problem if $UV(A) > 0$.

In our example, $UV(A)$ can only be higher if newly learned information can induce the company to build the oil platform off the shores of Brazil (action b), and not off the shores of Alaska ($a = a^*$).

The utility value UV is defined from the investigator's perspective. Metaphorically speaking, we can call an experiment a *question* to nature, and a result an *answer* from 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 himself. Nature shows oil, or doesn't show oil, according to whether there is oil where the exploration drilling takes place or not. It does not deliberate and show it in order to help the investigator, 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. So even from the companies, i.e. the receiver's perspective, relevant information is not the same as desirable information.

In (2001) van Rooij used (4.8) as a measure for the relevance of answers.[11]

---

[10] This type of situation has been thoroughly studied in statistical decision theory. Compare e.g. (Raiffa & Schlaifer, 1961, Sec. 4.5) and (Pratt et al., 1995).

[11] In later work v. Rooij tested other measures, e.g. (4.9) in (2003), (2003a), and compares quite a number of possible definitions in (2004). Prashant Parikh (1992) seems to be the first one who introduced (4.8) as a measure of linguistic *relevance*. Rohit Parikh (1994) used it for measuring the usefulness of communicated (vague) information.

It should be noted here that, in order to derive the semantics for *embedded* interrogatives, v. Rooij introduces an operator that combines the effects of the principles of relevance and quantity, see e.g. (2003a, Sec. 5.2). This was worked out further under the name of *exhaustification*, e.g. (2004). The operator is based on an order of relevance that is introduced as a special case of the order based on $UV(A)$ in the sense of (4.9). A due discussion of the relation between v. Rooij's exhaustification operator and the results of this paper must wait

We concentrate on this early proposal because, as we think, it shows quite clearly the principled limitations of a relevance based approach. Information is evaluated only from the inquirer's perspective, and as our example shows, this value is not identical with its desirability.

Although, a measure like (4.8) is defined only from one person's perspective we can apply it to the communication situation. We have to ask: *Who's* probability is P? 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 measures like (4.8) cannot be applied by E, except 1. and 3. coincide. If we assume that (a) the expert can only give answers that he believes to be true, then 2. implies that *any* answer A will do because then $EU(a, A) = EU(a)$ for all a. In order to turn the model into a model for a two-person game we have to choose interpretation 3.[12] In this case (4.8) advises the answering expert only to choose answers that can make I change his decision. The same problem that we found in the example with the oil company and its exploration drilling, we find with respect to questioning and answering:

(9) 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. Now, E learned privately that they are in stock at both places. What should E answer if she 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 ($A \wedge B$). Intuitively, all three are equally good. Some calculation shows that B is the *only* relevant answer according to (4.8). What (4.8) shows us is that B has the largest practical impact, but this is not the same as maximising joint payoff. The generalisation we are after is to show that the same problem shows up with *any* measure of relevance. We provide a formal proof in the next section.

As a further example we look at the following definition of utility value, also proposed by van Rooij[13] as an explication for relevance:

$$UV(A) = \max_{a \in \mathcal{A}} EU(a, A) - \max_{a \in \mathcal{A}} EU(a). \tag{4.9}$$

(4.9) gives the advice: '*Increase the hopes of the inquirer as much as you can*!' This fixes the problem with Example **(9)** but it's easy to see that we run into a similar problem with *negative* information: Assume that in the scenario of Example **(9)** E learns that there are no Italian newspapers at the station ($\neg A$); in this case (4.9) 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

for another occasion.

[12] Of course, that is the intended interpretation.

[13] See also (v. Rooij, 2003, Sec. 3.1) and (v. Rooij, 2003a, Sec. 3.3).

can be easily fixed again by taking the absolute | | of the right side of (4.9). But even here we find an example that shows a difference between the so–defined relevance and desirability. An answer that increases, or changes, the hopes of the inquirer as much as possible is not necessarily a good answer. We consider again Example **(7)**. Some calculation shows that, according to (4.9), E should answer that the Palace has been supplied with foreign newspapers. The same holds for the improved version of (4.9) 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 probabilities are higher that there are Italian newspapers at the station. A good answer should maximise the inquirer's chances for real success, and not maximally increase or change his expectations about success.

Van Rooij was aware that relevant answers might be misleading. In Sec. 4.2 of (v. Rooij, 2003) he discusses two reasons: The answering person (1) lacks important information or (2) has a reason to withhold information, e.g. due to opposing interests. The situations described in our examples differ in both respects: In **(7)**, **(8)** and **(9)** the interests completely coincide, and the expert has all the information necessary to decide that the answers picked out by criteria (4.8) and (4.9) are not optimal.

Relevance based Decision Functions

There are a number of reasonable decision theoretic explications for the notion of *relevance*. If we call information *relevant*, then the meaning of this depends on the special circumstances of the situation including the purposes of the interlocutors. In his recent work,[14] van Rooij discusses not one specific measure of relevance but compares whole groups of interesting explications and their merits and demerits in special applications. Hence, the measures that we discussed so far are only examples for types of measures.

Let $\mathcal{S}$ be the set of all support problems over a given set Ω, and for $S \in \mathcal{S}$ let D_S denote its associated decision problem. Adm_S denotes, as defined in (2.5), the set of admissible answers of S. Let $\mathcal{D} := \{\langle D_S, Adm_S\rangle \mid S \in \mathcal{S}\}$.

For the purpose of our paper, we can divide relevance measures into two groups: (i) in measures that depend only on a decision problem $\langle(\Omega, P), \mathcal{A}, u\rangle$ and pick out a relevant answer that depends on the admissible sets, i.e. measures that define a decision function $R : \mathcal{D} \longrightarrow \mathcal{P}(\Omega)$ such that every $R(D_S, Adm_S)$ is optimally relevant; (ii) in measures that depend in addition on a given hypothesis H such that they define a decision function $R : \mathcal{D} \times \mathcal{P}(\Omega) \longrightarrow \mathcal{P}(\Omega)$ such that $R(D_S, Adm_S, H)$ is of optimal argumentative force with respect to H. The second group corresponds roughly to the argumentative view of communication defended by Merin. Hence, we call the first group of decision functions *non–argumentative*, and the second group *argumentative* decision functions. The distinction cross classifies with our distinction between information and utility based measures, but, as a contingent matter of fact, the information based measure discussed previously belongs to the second class, and the utility based measures to the first class.

[14]See (v. Rooij, 2004).

Our aim in the next section is (1) to show that no non–argumentative measure of relevance can always pick out an optimal answer as defined by (2.6). Following our previous analysis, this implies that no non–argumentative measure can be empirically adequate. For every measure there will be a well–behaved support problem where the most relevant answer is not optimal. (2) We will show how our construction in Section 2 can be used to define an adequate argumentative decision functions in the sense of (*ii*).

5 Relevance and Best Answers

Given a support problem $S = \langle \Omega, P_E, P_I, \mathcal{A}, u \rangle$ we call any answer in Op_S as defined by (2.6) a *best answer*. The construction in Section 2 defines functions $BA : \mathcal{S} \longrightarrow \mathcal{P}(\Omega)$ that pick out optimal answers for each support problem S, i.e. with $BA(S) \in \mathrm{Op}_S$ for all $S \in \mathcal{S}$. It is the purpose of this section to study the relation between such best answer decision functions and functions that choose optimally *relevant* answers. This section does not need more mathematical skills than the previous ones, but its presentation is necessarily more formal and compact. We made a number of assumptions when we constructed optimal answers in Section 2. We start with a summary. We considered support problems $S = \langle \Omega, P_E, P_I, \mathcal{A}, u \rangle$ such that

1. S is well–behaved;
2. The answering person can only choose answers from $Adm_S = \{A \subseteq \Omega \mid P_E(A) = 1\}$ (Maxim of Quality);
3. There is a commonly known function $a : Adm_S \longrightarrow \mathcal{A}$, $A \mapsto a_A$, that chooses for each admissible answer A an action a_A that is optimal from I's perspective. a_A is the action that I will perform after learning A.[15]

The last assumption was necessary in order to guarantee that E can calculate the effects of an answer in cases where there are several optimal choices for I.

For our purposes we need a precise definition of *misleading answer*. An answer is misleading, if it induced the inquirer to perform an action of which E believes that it is not optimal:

Definition 5.1 (Misleading Answer) *Let $\langle \Omega, P_E, P_I, \mathcal{A}, u \rangle$ be a given support problem and $a : Adm_S \longrightarrow \mathcal{A}$, $A \mapsto a_A$, as above, then an answer $A \subseteq \Omega$ is* misleading, *iff $EU_E(a_A) \neq \max_{a \in \mathcal{A}} EU_E(a)$.*

Partial answers are answers like '*There are no Italian newspapers at the station*' in **(1e)**. This answer still rules out one of the actions, namely going to the station. Roughly, we call a partial answer *non–trivial* if there are at least two actions a, b such that it doesn't rule out both of them. There are still possible worlds where a is best, and others where b is best. This holds for both interlocutors. The answers in examples **(6)**–**(8)** are of this type.

[15] The last assumption was introduced on p. 7. For the definitions of well–behavedness and admissibility of answers see Definition 2.2 and Equation (2.5).

Definition 5.2 *Let $S = \langle \Omega, P_E, P_I, \mathcal{A}, u \rangle$ be a given support problem, then S has a* non–trivial partial answer *C, iff there exist actions $a, b \in \mathcal{A}$ and $a : Adm_S \longrightarrow \mathcal{A}$, $A \mapsto a_A$, as above, such that for*

$$\begin{aligned} A &:= \{v \in \Omega \mid u(v,a) > u(v,b)\} \\ B &:= \{v \in \Omega \mid u(v,b) > u(v,a)\} \end{aligned}$$

it holds that (1) $P_x(C), P_x(A|C), P_x(B|C) > 0$, for $x = I, E$, and (2) $a = a_A$.

Think we have a support problem and we have successfully defined a measure of relevance that picks out an answer C that happens to be an optimal answer too. As a decision theoretic model does only take account of the preferences and beliefs of one player, we can redefine the beliefs of the other player without changing the value of the measure of relevance. If C is partial and nontrivial, we can do it in such a way that C becomes thereby a misleading answer.

Theorem 5.3 *For every well–behaved support problem $S = \langle \Omega, P_E, P_I, \mathcal{A}, u \rangle$ with a non–trivial partial admissible answer C exists a probability distribution P'_E on Ω such that for the support problem $S' = \langle \Omega, P'_E, P_I, \mathcal{A}, u \rangle$ (1) the admissible answers are the same as for S and (2) C is a misleading admissible answer.*

Proof: Let $\langle \Omega, P_E, P_I, \mathcal{A}, u \rangle$ and C be given as in the theorem. Then, by Definition 5.2 there exist $a, b \in \mathcal{A}$ such that for

$$A := \{v \in \Omega \mid u(v,a) > u(v,b)\} \text{ and } B := \{v \in \Omega \mid u(v,b) > u(v,a)\}$$

it is $P_x(A|C), P_x(B|C) > 0$, $x = I, E$, and $EU_I(a|C) = \max_{c \in \mathcal{A}} EU_I(c|C)$. As abbreviations we use $\overline{B} := \{v \in \Omega \mid u(v,a) \geq u(v,b)\} = \Omega \setminus B$. We write:

$$P''_E(v) := \begin{cases} P_E(v|\overline{B}) \text{ for } v \in \overline{B} \\ P_E(v|B) \text{ for } v \in B \end{cases}$$

Then we set:

$$\begin{aligned} N_a &:= \sum_{v \in A} P''_E(v) \cdot (u(v,a) - u(v,b)) > 0 \\ N_b &:= \sum_{v \in B} P''_E(v) \cdot (u(v,b) - u(v,a)) > 0. \end{aligned}$$

Clearly, $N_a, N_b > 0$. Now we can define the probability distribution P'_E:

$$\begin{aligned} P'_E(v) &:= P''_E(v) \cdot \frac{N_b}{2(N_a + N_b)} \text{ for } v \in \overline{B}; \\ P'_E(v) &:= P''_E(v) \cdot \left(1 - \frac{N_b}{2(N_a + N_b)}\right) \text{ for } v \in B. \end{aligned}$$

P'_E is obviously a probability distribution over Ω. Let $N := N_b/(2(N_a + N_b))$. We first show (1), i.e. that the admissible answers are the same. This follows from elementary calculations; we show only $P'_E(D) = 1 \Rightarrow P_E(D) = 1$. Let $\alpha := P_E(D|\overline{B})$ and $\beta := P_E(D|B)$. If $\alpha < 1$ or $\beta < 1$, then $P'_E(D) = \alpha N + \beta(1 - N) < N + (1 - N) = 1$. Hence $\alpha = \beta = 1$, and it follows that $P_E(D) = 1$. Next we show (2). By assumption we know that $EU_I(a|C) = \max_{c \in \mathcal{A}} EU_I(c|C)$. Hence, in order to show that C is misleading we have to prove that $EU_{P'_E}(a) < EU_{P'_E}(b)$. It is

$$EU_{P'_E}(c) = \sum_{v\in\overline{B}} P'_E(v)\cdot u(v,c) + \sum_{v\in B} P'_E(v)\cdot u(v,c) \text{ for all } c\in\mathcal{A}.$$

We find that $EU_{P'_E}(b) - EU_{P'_E}(a) =$

$$\sum_{v\in A} P''_E(v)N(u(v,b)-u(v,a)) + \sum_{v\in B} P''_E(v)\,(1-N)\,(u(v,b)-u(v,a)) \;=$$

$$= \; N\cdot(-N_a) + (1-N)\cdot(N_b) \;=\; \frac{-N_aN_b}{2(N_a+N_b)} + N_b - \frac{N_b^2}{2(N_a+N_b)}$$

$$= \; N_b - N_b\cdot\frac{1}{2}\cdot\frac{N_a+N_b}{N_a+N_b} \;=\; \frac{1}{2}N_b \;>\; 0.$$

This shows that from E's perspective b would be the preferable action. Hence, C is misleading and (2) is proven. ■
From this theorem it follows immediately:

Corollary 5.4 *Let $\mathcal{S}$ be the set of all support problems over Ω. For $S\in\mathcal{S}$ let D_S denote its associated decision problem. Let $\mathcal{D} := \{\langle D_S, Adm_S\rangle \mid S\in\mathcal{S}\}$ where Adm_S is the set of admissible answers of S. Then there exists no function $R:\, \mathcal{D}\longrightarrow\mathcal{P}(\Omega)$ such that for all $S\in\mathcal{S}:\; R(D_S, Adm_S)\in \mathrm{Op}_S$.*

In **(6)**–**(8)** we saw examples where an empirically adequate criterion for optimal answers must pick out a non–trivial partial answer. Hence, if a measure of relevance is empirically adequate, then there are examples where it must choose non–trivial partial answers. But then there must also be an example where it chooses a misleading answer.

This shows that no decision theoretically defined non–argumentative measure of relevance can be adequate for all support problems. What about argumentative measures as that proposed by Merin? The following proposition shows that there are *argumentative* decision functions that always select optimal answers we can presuppose a function that provides for each support problem S a suitable hypothesis H_S for which E has to argue.

Proposition 5.5 *Let $\mathcal{S}$ and $\mathcal{D}$ be defined as in Cor. 5.4, and assume that the previous conditions for construction best answer functions are fulfilled. Then there exists a function $H:\, \mathcal{S}\longrightarrow\mathcal{P}(\Omega)$, $S\mapsto H_S$ and a function R such that for all $S\in\mathcal{S}$ $R(D_S, Adm_S, H_S)\in\mathrm{Op}_S$.*

Proof: Let $S = \langle\Omega, P_E, P_I, \mathcal{A}, u\rangle$ be a given support problem. Let $BA:\, \mathcal{S}\longrightarrow\mathcal{P}(\Omega)$ be such that for all $S\in\mathcal{S}$ $BA(S)\in\mathrm{Op}_S$, i.e. BA is a best answer decision function. Then we simply set $H_S := BA(S)$ and $R(D_S, Adm_S, H_S) = H_S$. Clearly, this decision function has the desired properties. ■

For us the main importance of this proposition lies in the fact that it makes absolutely clear the relation between our game theoretic model of questioning and answering and explanations based on decision theory. The latter need an externally given hypothesis as a goal for which an interlocutor could argue. In our model, this hypothesis is provided theory internally. But this remains the 'only' difference between the two approaches. Hence, the proposition provides for us a bridge to applications of pure decision theory. In the previous sections we have emphasised the differences and, of course, the weaknesses and shortcomings. This obscures somewhat the sheer usefulness of this apparatus. In all cases where we don't need to bother about the argumentative goal, a decision theoretic criterion of relevance may be completely adequate.

From Corollary 5.4 we know that a non–argumentative decision function cannot guarantee that we always select optimal answers. Proposition 5.5 seems to show that the argumentative conception of communication builds the proper basis for an explication of relevance. Whether it really warrants this conclusion, we have to investigate another time. We leave it at this point.

The same holds for the question concerning the wider significance of our results for Relevance Theory. Sperber & Wilson (1986) are well known for their claim that the Gricean maxims can be reduced to the Maxim of Relevance. If our arguments that applied to decision theoretic explications carry over to Relevance theoretic explications, then the consequences are indeed far reaching. They would amount to a proof that no such reductionist approach can be empirically adequate. But whether or not this conclusion is warranted or not must again be left for future research.

6 Conclusion

We set out with the aim to derive a measure of utility of answers from a game theoretic model of communication that accounts for a number of judgements about the appropriateness of partial and mention–some answers. In general, we looked at communication as a sequential two–person game of complete coordination. We presented a sketch how to explain the existence of mention–some answers even if one assumes that the basic answer to a question is the strongly exhaustive answer. We argued that mention–some answers contribute to goals of the inquirer that are super–ordinated to the immediate goal of getting an answer to their question. Relative to these super–ordinated goals they provide *optimal* information. The set of best answers can be calculated by backward induction. We applied our model to a number of examples. A sub–group of partial answers, which we called *non–trivial* partial answers, turned out to be especially interesting; they make it necessary for the answering person to take into account the possibility of misleading information. We showed in the second part of the paper that our model improved here over previous explanations based on decision theoretically formulated relevance measures. The main goal of the second part was to provide a principled characterisation of the relation between our game theoretic model and approaches that use a decision theoretically defined measure of relevance for finding optimal answers. They define relevance from the perspective of the receiver of information. Choosing maximally relevant answers then means trying to maximise his expectations about responses. As is to be expected, this runs the risk of providing misleading information. We found that no decision function based on maximal relevance can be successful in avoiding this risk.

This brings us to the final question: What conclusions does our analysis allow about the status of Gricean maxims? Our model incorporated the Cooperation Principle and the Maxim of Quality, which principles were supplemented by the assumption (Utility) that interlocutors are Bayesian utility maximisers, i.e.

choose actions that maximise expected utilities. It is clear, that we needed in addition the Maxim of Manner in order to rule out overly complex answers. As we provided here only a case study, we can formulate only tentative conjectures about the other maxims:

Conjecture 1: The Maxim of Relevance is not among the basic axioms of pragmatics.

Conjecture 2: The first sub–maxim of Quantity (*Say as much as you can*) is superfluous — as a consequence of (Utility).

Previously, we wrote that the more general question behind our discussion of explications of relevance is the question whether or not it is essential to model communication as a *two*–person game. I hope, it could show that it is.

References

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.

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.

P. Parikh (1992): *A Game–Theoretic Account of Implicature;* Proceedings of the 4th Conference on Theoretical Aspects of Reasoning about Knowledge, Morgan Kaufmann, Monterey, CA.

R. Parikh (1994): *Vagueness and Utility: The Semantics of Common Nouns; Linguistics and Philosophy* 17.

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

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

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

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

R. v. Rooij (2004): *Relevance of complex sentences;* To appear in Proceedings of LOFT 04.

D. Sperber, D. Wilson (1986): *Relevance : Communication and Cognition;* Blackwell, Oxford.