Values of questions, and Q-A games

Robert van Rooy
ILLC/University of Amsterdam

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. In this paper I argue that questions are asked because theirs answers help to resolve the questioners decision problem. By relating questions to decision problems I show (i) how we can measure the values of questions, and (ii) how with the help of these values and certain other assumptions we might model the strategic deliberations of participants in a question-answer dialogue game.

Values of completely answered questions

In Savage's (1954) statistical decision theory a distinction is made between states of the world, acts, and consequences; states of the world together with acts determine the consequences, and the relevant expectation is unconditional expected utility. If we assume that the utility of doing action a in world w is $U(w, a)$, we can say that the *expected utility* of action a, $EU(a)$, with respect to probability function P is $EU(a) \quad = \quad \sum_w P(w) \times U(w, a)$.

Let us now assume that the agent faces a *decisison problem*, i.e. he wonders which of the alternative actions in A he should choose. A decision problem of an agent can be modeled as a triple, $\langle P, U, A \rangle$, containing (i) the agents probability function, P, (ii) his utility function, U, and (iii) the alternative actions he considers, A. You might wonder why we call this a decision *problem*; should the agent not simply choose the action with the greatest expected utitility? Yes, he should, if he *chooses now*. We might say that the *value* of choosing now is the expected value of the act with the greatest expected utility (where i varies over the actions in A): $U(\text{Choose now}) \quad = \quad max_i \sum_w P(w) \times U(w, a_i)$.

But now suppose that the agent doesn't have to choose now, but can try to receive some useful information by first *asking a question*. Suppose that the other participant of the dialogue answers this question by giving answer C, and that, as a good Bayesian, the agent himself updates his probability state by conditionalizing on C. Then we can say that the value of making an informed decision conditional on learning C is the expected utility conditional on C of the act that has highest expected utility: $U(\text{Learn C, choose later}) \quad = \quad max_i \sum_w P(w/C) \times U(w, a_i)$.

It should be obvious that when C is non-trivial with respect to P, the utility of choosing after learning the answer is always greater than the utility of choosing before learning. Let us call this difference the *utility of the answer*, a value that is always dependent on a decision problem. What we are after, however, is something different; we want to know what the *utility of a question* is. But here we face a problem, because even if we assume that the questioner will accept the given answer, the questioner doesn't know which answer this other participant will give. It turns out, fortunately, that we can determine the utility of a question easily when we assume that a question *partitions* the state space. First, we can determine the value of an informed decision; this value is the present expectation of expected value of the choice made after learning the answer (where k varies over the alternative answers):

$$U(\text{Learn answer, choose later}) \quad = \quad \sum_k P(C_k) \times max_i \sum_w P(w/C_k) \times U(w, a_i)$$

In the decision theory of Savage this value is equal to $\sum_k max_i \sum_{w \in C_k} P(w) \times U(w, a_i)$, which, as shown by Skyrms (1990),[1] is on certain natural assumptions always at least as great as our earlier $U(\text{Choose now})$. But now we can also determine the *value of the question*. This value, usefulness, or relevance, of a question, given a decision problem, can be defined as the difference between $U(\text{Learn answer, choose later})$ and $U(\text{Choose now})$, which is a natural measure, if we assume that learning via questioning is cost-free.

[1]Skyrms, B. (1990), *The Dynamics of Rational Deliberation*, Cambridge Mass.: Harvard University Press.

The question-answer game

Suppose now that an agent faces a decision problem, which question should he then ask? The answer seems obvious; the question that gives rise to the highest value with respect to that decision problem. Notice that if two questions both give rise to partitions, we can say that the first partition, P, is a *refinement* of the other partition, Q, iff $\forall X \in P : \exists Y \in Q : X \subseteq Y$. It should be obvious that the question that gives rise to a more refined partition has always a value that is at least as high as the other one.

But notice that this fact is based on two non-trivial assumptions; (i) asking a question is always cost-free, in particular, the questioner doesn't mind to make his intentions fully explicit and never tries to hide them, and (ii) the respondent will always answer the question by giving the *complete* answer; it is assumed that the respondent is fully informed and fully cooperative. Let us now give up these assumptions, i.e. let us assume that the value of a question might decrease if it makes explicit your intentions, and that the questioner has to take into account the fact that, for whatever reasons, the respondent might give a partial answer to a question even though he knows the complete answer. If we now can determine the values of the strategies of giving full or partial answers by the respondent, and if we can determine the value of a question for the questioner in case the respondent gives only a partial answer, we end up with a game between questioner and respondent.

Suppose, for instance, that John, the questioner, is in love with Mary, but, everything else being equal, prefers not to make this clear to Bill, the respondent. That is, if Bill would give a complete true answer, he would prefer to ask the question *Who will come to the party?* to asking explicitly *Will Mary come?* Let us say that being explicit has a cost of ϵ. Bill, on the other hand, doesn't mind to give a full answer to the question *Will Mary come?*, but prefers to give a partial answer to the question *Who will come?* Thus, both players have two strategies, and let us assume that the game is one with perfect information.

What then are the payoffs? Assume that there only 2 relevant individuals over which the *wh*-word varies, Mary and Sue, and that it is assumed that at least one of the two will come to the party. In that case, the question *Who will come?* has three complete answers: *Only Mary*, *Only Sue*, and *Mary and Sue*. Assuming that all worlds have equal probability, that the value of an uninformed decision is 3, and that the utility of the best action in each world where the questioner knows whether Mary will come is 6, it's easy to see that the value of *Who will come?* is $(1/3 \times 6) + (1/3 \times 6) + (1/3 \times 6) - 3 = 6 - 3$, if Bill gives a complete answer, and that the value of asking the explicit question *Will Mary come?* is $(1/2 \times 6) + (1/2 \times 6) - \epsilon - 3 = 6 - \epsilon - 3$. To complete the payoff function for John, we have to determine the value of *Who will come?* in case Bill gives a partial answer. Let us also limit ourselves to the following two partial answers *Mary will come*, M, and *Sue will come*, S. Note that the latter (mention-some) answers, unlike the earlier complete answers, will not partition the state space. It seems reasonable then to determine the value of the first question in these circumstances as the difference between the following utility and 3:

$$
\begin{aligned}
U(Q_p) \quad = \quad & P(M) \times max_i \textstyle\sum_w P(w/M) \times U(w, a_i) && (= U(\text{Learn M, choose later})) \\
& + P(S) \times max_i \textstyle\sum_w P(w/S) \times U(w, a_i) && (= U(\text{Learn S, choose later})) \\
& \text{-}\ \ P(M \wedge S) \times max_i \textstyle\sum_w P(w/(M \wedge S)) \times U(w, a_i) && (= U(\text{Learn M} \wedge \text{S, choose later}))
\end{aligned}
$$

On the assumption that the act with the highest value in case John doesn't know whether Mary will come is 3 (in case Bill answers *Sue*), the value of $Q_p - 3$ will be $(2/3 \times 6) + (2/3 \times 3) - (1/3 \times 6) - 3 = 4 - 3$, which is lower than the value of the same question in case Bill gives a complete answer, and higher or equal to the explicit question *Will Mary come?* just in case $\epsilon \geq 2$. If we call the difference between the value of *Who will come?* in case Bill gives a complete or partial answer the *risk* of asking an implicit answer, we might say that John should be explicit in his questioning just in case the *cost* of asking the explicit question, ϵ, is *smaller* than the *risk* of asking the implicit question, 2, which confirms our intuition about the situation.

To make the description of the game complete, we have to know whether Bill prefers the strategy of giving a complete answer to *Who will come?* in case John actually asks *Will Mary come?* or not. If not, then depending on whether the cost of being explicit for John is greater than the risk of being implicit, the game will have either one or two Nash-equilibria.