The role of preferred outcomes in determining implicit questioning strategies

Jon Scott Stevens*

Center for General Linguistics, Berlin, Germany

Abstract

We look at a particular class of indirect answers to yes/no questions, where the speaker supposes that the hearer's question was motivated by some underlying goal and aims to provide the hearer with an alternative way of accomplishing it. We posit that the questions in these cases are sub-questions in an implicit questioning strategy which is anchored to the hearer's goal, and which must be inferred probabilistically by the speaker. The use of these implicit questioning strategies is licensed when the hearer has a preferred outcome, i.e., better and worse ways to accomplish her goal. We formally model the generation and interpretation of questions and answers in this context as a Bayesian game.

1 Indirectness & implicit questioning strategies

We consider a particular class of phenomenon whereby a yes/no question ('YN-question') is answered indirectly with a particular intonation pattern, lacking the falling pitch contour that typically characterizes the end of a declarative answer.

(1) Q: Does Bob's Market sell turnips?
A: Freshtown sells turnips...
H* L-H%

The placement of the H* pitch accent in the answer is determined by focus. The salience of $\lambda x.sell(x, turnips)$ introduced by the question requires focus on the subject in (1) (Rooth, 1992; Schwarzschild, 1999; Wagner, 2012b), and in a different context, e.g. in (2), we find the pitch accent occurring elsewhere to reflect the different focus structure.

(2) Q: Does Bob's Market sell sour cream and onion chips?
A: Bob's Market sells cheese and onion chips...
H* L-H%

But independently of focus placement, we find a rising boundary tone which signals discourse non-finality, typically suggesting that the question under discussion (Roberts, 1996) is still open (see Lai, 2012). The answer does have a flavor of non-finality in that it is interpreted by default as a mention-some answer, indicating in (1) that perhaps other markets sell turnips as well (see also Wagner, 2012a). But rather than signaling that the question is unresolved, the answers in (1) and (2) both clearly implicate a 'no' answer to the YN-question (assuming a context in which the answerer is assumed to know who sells what), allowing the discourse to potentially conclude without any follow-ups.

The license for these answers is not so straightforward under previous approaches to indirect answerhood which hold that indirect answers must allow for a clear inference of a 'yes' or 'no' answer based on the denotation of the response given (Asher and Lascarides, 2003; de Marneffe et al., 2009), as in (3).

*Work on this paper was supported by the Deutsche Forschungsgemeinschaft (DFG), Grant Nr. BE 4348/3-1. Special thanks to Anton Benz for his feedback.

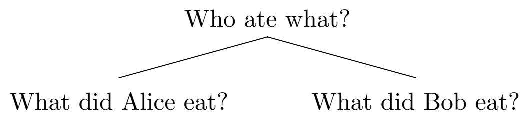

Figure 1: A questioning strategy.

(3) Q: Does Bob's Market sell turnips?
A: Bob's Market has a terrible selection of produce.

Here, there is a clear probabilistic inference from the answer given to the proposition 'Bob's Market does not sell turnips', a valid literal answer to the YN-question. But in (1) and (2), this does not hold. There is no clear inferential link, for example, between Bob's Market selling turnips and Freshtown selling turnips.

Such indirect answers are called 'alternative answers' by Stevens et al. (2015), who model the generation of responses to YN-questions by representing and reasoning about the questioner's *goals*. Intuitively, questions are a means to some end, and often a question serves to gather the information required to accomplish some goal in the real world (van Rooij, 2003; Benz and van Rooij, 2007). Indirect answers can be licensed by anticipating these goals and offering alternative solutions. In (1), for example, one assumes that the questioner likely has the goal to go shopping for turnips, and needs to know which store(s) she can go to in order to do this. It is only in virtue of this assumption that the answer in (1) is licensed. Consider some other logically possible alternative answers.

(4) Q: Does Bob's Market sell turnips?
A: ?Bob's Market sells carrots...

(5) Q: Does Bob's Market sell turnips?
A: #Bob's Market sells hand soap...

Whereas the exchange in (1) is acceptable more or less out of the blue, (4) would require some explicit evidence that carrots were an acceptable alternative to turnips (perhaps the questioner simply wants some root vegetables, and is not picky), and (5) would require a much stranger context, and it is infelicitous out of the blue. This difference can be seen as a difference in probability of goals.

Formally, the difference between (1), (4) and (5) can be analyzed as a constraint on *implicit questioning strategies*, a type of questioning strategy in the sense of Roberts (1996) and Büring (2003). Originally devised to account for contrastive topics in answers to multiple wh-questions, a questioning strategy is a way of breaking questions down into sub-questions. Sub-questions are defined, after Groenendijk and Stokhof (1984), such that Q1 is a sub-question of Q0 iff exhaustively answering Q0 logically provides an answer to Q1. Strategies may be structured as trees as in Fig.1, which represents a possible strategy for the multiple wh-question 'Who ate what?'. Imagine a discourse in which the only two relevant individuals are Alice and Bob. The question 'Who ate what?' can be exhaustively answered by first answering 'What did Alice eat?' and then answering 'What did Bob eat?', resulting in an answer like, "Alice ate turnips, and Bob ate carrots."

We can construct a similar tree for (1), as in Fig.2. But crucially, although the root node question in structures like Fig.1 which motivate the placement of contrastive topics, e.g. a multiple wh-question like 'Who ate what?', is typically taken to be a proper question under discussion—that is, it is known to both interlocutors—the questioning strategy in Fig.2 in the context of (1) is one where the root node question is private to the questioner, and thus must be inferred by the answerer. In other words, only one of the subordinate

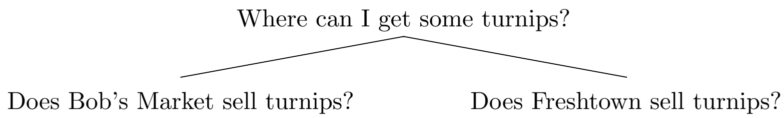

Figure 2: An implicit questioning strategy.

YN-questions is explicitly put on the table in the dialogue. The overarching strategy is implicit.[1]

Intuitively, the answerer in (1) is making two inferences. First, the observed YN-question is likely to have been motivated by an underlying super-question, 'Where can I get some turnips?'. Second, the hearer's goal in asking the question is to find a single place to buy turnips, and therefore a mention-some answer is appropriate. Taken together, these two inferences provide a clear path to an indirect answer: If Bob's Market sells turnips, then say so, and if not, then anticipate any sister questions by offering a single alternative turnip seller. Because the indirect answer is only ever licensed in the event that Bob's Market does not sell turnips, assuming that the speaker knows who sells what, an implicit 'no' can always be inferred.

The infelicity of (4) and (5), then, is reduced to the unavailability in a neutral context of the inference that the observed YN-question is part of an implicit questioning strategy headed by 'Does Bob's Market sell turnips and/or carrots?' or 'Does Bob's Market sell turnips and/or hand soap?', respectively. The availability of these inferences is highly context-dependent—with enough context, almost any implicit questioning strategy can be inferred.

(6) Context: The questioner has a rather strange Secret Santa assignment this year, who wants either turnips or hand soap for Christmas. The questioner is going to Bob's Market later and wants to make a shopping list.

Q: Does Bob's Market sell turnips?
A: Bob's Market sells hand soap...

It is possible to account for these inferences within a game-theoretic framework (Benz et al., 2006; Parikh, 2010; Clark, 2011; Franke, 2011; Benz, 2012) by modeling such QA-exchanges as Bayesian games, where the behavior of both the questioner and answerer is determined by probabilistic reasoning about unknown variables (for the answerer, the questioner's underlying goal, and for the questioner, the true state of the world). Under this approach, it is the common-sense knowledge that finding a store at which to buy a particular item is a likely enough goal to have, combined with a drive toward efficiency in dialogue (anticipating follow-up questions, etc.), which licenses the indirect answer.

Stevens et al. (2015) offer a computational model along these lines, where the problem of producing alternative answers is reduced to the problem of learning a prior probability distribution over possible hearer goals. But that model is oriented toward generating answers to questions, and doesn't offer a full picture of how the questions themselves are selected. Therefore, one important puzzle remains: Why would the questioner ever employ an implicit questioning strategy in the first place? Instead of Fig.2, why not simply ask, 'Where can I get some turnips?' And why ask about Bob's Market first and not Freshtown? In this paper we argue that implicit questioning strategies of this type are to be expected whenever the questioner has a *preferred outcome*. For example, it might be the case that Bob's Market is cheaper or closer to the questioner's house than Freshtown, and thus the

[1] That the question is not shared, but rather merely inferrable, explains why we do not find narrow focus on *Bob's Market* in the question.

best outcome is one where the questioner buys turnips from Bob's. To ask the more general question would risk a mention-some answer that guides the questioner to a sub-optimal market when she could have gone to the preferred market instead. Acknowledging that there are better and worse ways to accomplish a goal gives us a fuller picture of how questions and answers are generated and interpreted in these contexts.

We begin by outlining our game-theoretic model, which combines insights from Franke's (2011) 'iterated best response' approach and Benz's (2012) 'error models'. We then work through a derivation of examples (1), (5) and (6).

2 Model

We model QA-exchanges as signaling games, a class of Bayesian games (Fudenberg and Tirole, 1991). The players in the game are the questioner and the answerer, who both want the questioner's goal to be accomplished. Based on the question that was asked, the answerer strategically selects a *message*, m, which carries some conventional propositional content. The questioner must then reason, based on m, about the current state of the world, ω (e.g., which stores sell which items). Using these inferences about the world state, the questioner can choose an *action*, a (e.g., going to Bob's or Freshtown), that will lead to an optimal outcome. The optimality of an outcome—its *utility*—is determined by whether and how well it accomplishes the questioner's *goal*, γ. The questioner can anticipate what messages might be generated by the answerer in order to strategically choose questions that most efficiently optimize her outcome. The *selection procedures* whereby the players select questions, answers and actions can be represented as functions which yield sets of strategically optimal questions, answers and actions from which the player randomly selects (allowing for a tie between multiple equally beneficial moves). In the end, any set of procedures for selecting questions, answers and actions for which neither player can improve her expected outcome given the behavior of the other player is an *equilibrium*. Our goal is to account for implicit questioning strategies and their associated indirect answers by showing that these are felicitous only when they are part of an equilibrium in a signaling game.

2.1 Goals & expected utility of actions

We begin by representing a space of possible questioner goals, where a goal is a set of instructions for accomplishing some concrete task. More precisely, for our purposes, a goal is an ordered tuple of conditional imperatives, where executing any one of the imperatives under the right condition will lead to the goal being accomplished. For example, instructions for buying turnips could be represented as an ordered tuple ⟨'Go to Bob's if Bob's sells turnips', 'Go to Freshtown if Freshtown sells turnips'⟩, where the second instruction is only executed if the first fails. If both fail, then the questioner should simply do nothing. Consider three possible goals, representations of which are given in Table 1, where each goal involves shopping and is accomplished by going to one of two area stores and/or adding relevant items to a shopping list. The ordering of the instructions encodes the questioner's preferences. γ_t instructs the questioner to go to Bob's if Bob's sells turnips, regardless of what Freshtown sells. Only if that condition fails to hold should the questioner consider going to Freshtown. This encodes a preference to go to Bob's. Similarly, γ_B, which corresponds to a context like in (6), encodes a preference to buy turnips over hand soap (perhaps because turnips are cheaper). Finally, the most specific goal to be inferred from the question 'Does Bob's Market sell turnips?' is the singleton goal γ_0: The questioner wants to buy turnips from Bob's, and no other store-item pair will do.

The first step in building up our model will be to use these goal representations to create a generalized schema for generating *utility functions* which we will call *action utility*

Label	Representation	Description
γ_t	$\langle go.to(B)$ if $sell(B,t)$, $go.to(F)$ if $sell(F,t)\rangle$	Questioner wants to buy turnips from either Bob's or Freshtown, and prefers to go to Bob's.
γ_B	$\langle add(t)$ if $sell(B,t)$, $add(h)$ if $sell(B,h)\rangle$	Questioner wants to go shopping at Bob's for either turnips or hand soap, and prefers to buy turnips.
γ_0	$\langle add(t)$ & $go.to(B)$ if $sell(B,t)\rangle$	Questioner wants to buy turnips from Bob's, or else do nothing.

Table 1: Three possible questioner goals. B stands for 'Bob's Market', F for 'Freshtown', t for 'turnips' and h for 'hand soap'. Possible imperatives are *go.to*—go to a given store—and *add*—add a given item to your shopping list.

functions. An action utility function returns a numerical value for each combination of real-world action a (going to Bob's, buying turnips, etc.), goal γ, and current *world state* ω, where the current world state encodes which relevant information is true in the world, i.e. which stores sell which products. Action utility is higher when the questioner's action leads to a preferred outcome. Therefore, every questioner's ultimate goal, in the most general terms, is to maximize action utility. One simple schema for one- and two-member goal tuples, which suffices for current purposes, is given in (7) below, where action $\emptyset$ corresponds to doing nothing, and where the questioner has the option of asking a follow-up question (Benz, 2012).

(7) For any one- or two-member goal γ of the form $\langle P$ if Φ, $(R$ if $\Psi)\rangle$, action utility $U(\gamma, \omega, a)$ is equal to:
   a. 2 iff $a = P$ and Φ is true in ω
   b. 1 iff $a = R$ and Ψ is true in ω
   c. 1 iff $a = \emptyset$ and $\neg(\Phi \vee \Psi)$ is true in ω
   d. $1 - c$ iff the questioner asks a follow-up question
   e. 0 elsewhere

The highest action utility value, 2, is awarded when the action leads to the questioner's preferred outcome. "Runner-up" outcomes are awarded utility 1. Moreover, doing nothing yields utility 1 in worlds where the questioner's goal is impossible to accomplish; this encodes an avoidance of fool's errands. Finally, the questioner always has the option to ask a follow-up question. To unpack the implications of follow-ups, consider a questioner with goal γ_t. She could ask, "Does Bob's Market sell turnips?", as in (1), and the answerer could simply reply, "no". If the questioner remains uncertain as to whether Freshtown sells turnips, then she can follow up with, "Does Freshtown sell them?" (the sister question in Fig.2). If the answer is 'yes', the questioner will know to go to Freshtown, and if the answer is 'no', the questioner will know to do nothing rather than looking for turnips at either store to no avail. Under the assumption that the answerer will supply a 'yes' or 'no' answer to a follow-up, the outcome is always one that will lead to action utility 1. However, the questioner incurs a small *effort cost* along the way: An additional QA-exchange was required to accomplish the task at hand. We represent this as a small constant c. Action utilities generated for γ_t, γ_B and γ_0 from Table 1, broken down by possible world state, are given in Table 2.

The questioner wants to maximize action utility, but it is impossible to do this directly, because the questioner does not know the current world state ω. Information about ω must be drawn from the message m which was selected by the answerer. Indeed this is

ω	γ_t			γ_B			γ_0	
$\langle B, F \rangle$	$go.to(B)$	$go.to(F)$	$\emptyset$	$add(t)$	$add(h)$	$\emptyset$	$add(t)$ & $go.to(B)$	$\emptyset$
$\langle t+h, t+h \rangle$	2	1	0	2	1	0	2	0
$\langle t+h, t \rangle$	2	1	0	2	1	0	2	0
$\langle t+h, h \rangle$	2	0	0	2	1	0	2	0
$\langle t+h, \emptyset \rangle$	2	0	0	2	1	0	2	0
$\langle t, t+h \rangle$	2	1	0	2	0	0	2	0
$\langle t, t \rangle$	2	1	0	2	0	0	2	0
$\langle t, h \rangle$	2	0	0	2	0	0	2	0
$\langle t, \emptyset \rangle$	2	0	0	2	0	0	2	0
$\langle h, t+h \rangle$	0	1	0	0	1	0	0	1
$\langle h, t \rangle$	0	1	0	0	1	0	0	1
$\langle h, h \rangle$	0	0	1	0	1	0	0	1
$\langle h, \emptyset \rangle$	0	0	1	0	1	0	0	1
$\langle \emptyset, t+h \rangle$	0	1	0	0	0	1	0	1
$\langle \emptyset, t \rangle$	0	1	0	0	0	1	0	1
$\langle \emptyset, h \rangle$	0	0	1	0	0	1	0	1
$\langle \emptyset, \emptyset \rangle$	0	0	1	0	0	1	0	1

Table 2: Values of action utility function $U(\gamma, \omega, a)$, grouped by questioner's goal, where each world state ω is characterized by which relevant items Bob's and Freshtown sells. An additional action "follow-up question" is always available with utility $1 - c$. World states specify which relevant items Bob's sells, then which relevant items Freshtown sells.

a key function of questions (van Rooij, 2003): to obtain information about the world in order to better maximize utility. Because the link between m and ω is conventional, and not absolute, the hearer encodes the effect of m in a conditional probability distribution over world states. This probability, which encodes a *belief* on the part of the questioner, is conditioned by m as well as by the question that was asked, q, and the questioner's assumption about the answerer's procedure, $\mathcal{S}_M$, for generating messages based on q and ω. The best the questioner can do, then, is to choose actions which maximize utility in the aggregate by maximizing the weighted average of action utilities over possible worlds, weighted by $P(\omega|m, q, \mathcal{S}_M)$. This average is the *expected action utility*.

$$EU(a|\gamma, m, q, \mathcal{S}_M) = \sum_{\omega} P(\omega|m, q, \mathcal{S}_M) \cdot U(\gamma, \omega, a) \quad (1)$$

2.2 Expected utility of answers

We now work backwards from actions to answers, to determine how to strategically optimize answers to YN-questions. Assuming a cooperative Gricean scenario, the answerer wishes above all for the questioner to achieve her immediate goal. At the same time, we can reasonably assume a preference for direct answers, all things being equal. To encode this, we posit a *message utility* function U_M which, for any given outcome of the QA-exchange, is the action utility awarded to questioner for that outcome minus the *cost* of the message, $C(m)$, which we take here to be zero for direct answers, and some small constant k for indirect answers.

$$U_M(\gamma, \omega, m, a) = U(\gamma, \omega, a) - C(m) \quad (2)$$

There are two unknown variables: The answerer does not know the identity of γ or the action a which the questioner will take after receiving m. The expected utility of a message must therefore be calculated by first making inferences about a based on values of γ, and then making inferences about γ based on the question that has been asked, q. The latter inference requires a representation of the questioner's procedure $\mathcal{S}$ for using γ, m and q

to select an action, and of the questioner's procedure $\mathcal{S}_Q$ for using γ to select questions. The resulting formulation is more complex than expected action utility, but it can be boiled down as follows: The expected utility of a message is the weighted average of the message utility of all possible outcomes involving that message, where the probability of an outcome depends on assumptions about how the questioner will react to each message—$P(a|\gamma, m, q, \mathcal{S})$—and about how the question was selected to further the questioner's goal—$P(\gamma|q, \mathcal{S}_Q)$.

$$EU_M(m|\omega, q, \mathcal{S}, \mathcal{S}_Q) = \sum_{\gamma,a} P(\gamma|q, \mathcal{S}_Q) \cdot P(a|\gamma, m, q, \mathcal{S}) \cdot U_M(\gamma, \omega, m, a) \quad (3)$$

2.3 Expected utility of questions

Again stepping backward through the QA-exchange, we move from the answerer's expected utility of sending a message given a certain question to the questioner's expected utility of posing a question given a certain goal. This also requires an assumed message selection procedure, $\mathcal{S}_M$.

$$EU_Q(q|\gamma, \mathcal{S}_M) = \sum_{\omega,m} P(\omega) \cdot P(m|\omega, q, \mathcal{S}_M) \cdot U(\gamma, \omega, \arg\max_a\ EU(a|\gamma, m, q, \mathcal{S}_M)) \quad (4)$$

For any world-message pair, we can calculate the utility the questioner will receive in that world if she takes the action which maximizes EU given that message and $\mathcal{S}_M$. The weighted average for all such world-message pairs is the expected utility of the question. More simply, questions are selected by considering what answers they might elicit, and what best-case outcomes will result from receiving those answers.

We now have three quantities, EU, EU_M and EU_Q, corresponding respectively to the optimality of an action given a question/answer pair and goal, the optimality of an answer given a question and world state, and the optimality of a question given a goal. An equilibrium in this sequential game is one where the questioner chooses a question that maximizes EU_Q, the answerer provides a response that maximizes EU_M and the questioner responds to the answer with an action that maximizes EU, where the players' beliefs, $\mathcal{S}_Q$, $\mathcal{S}_M$ and $\mathcal{S}$, on which the expected utility values depend, reflect the assumption that all players maximize their own expected utility.[2]

2.4 Finding equilibrium

Where the model of Stevens et al. (2015) generates answers by optimizing EU and EU_M, taking the question to be background information, the current model must balance these with EU_Q in order to select questions as well as answers. This complicates the procedure somewhat, but it is nonetheless possible to derive an equilibrium via back-and-forth reasoning (see e.g. Franke, 2011). The idea is to find a sensible default set of beliefs to start with, find optimal selection procedures for those beliefs, then incorporate those procedures into a new set of beliefs, iteratively refining procedures until the iteration converges on a fixed point. We propose the following back-and-forth reasoning schema.

1. Set default procedures $\mathcal{S}^0$ and $\mathcal{S}^0_Q$ to stand in for $\mathcal{S}$ and $\mathcal{S}_Q$, respectively
2. Let $\mathcal{S}^1_M(\omega, q)$ be $\arg\max_m\ EU_M(m|\omega, q, \mathcal{S}^0, \mathcal{S}^0_Q)$
3. Let $\mathcal{S}^2_Q(\gamma)$ be $\arg\max_q\ EU_Q(q|\gamma, \mathcal{S}^1_M)$
4. Let $\mathcal{S}^2(\gamma, m, q)$ be $\arg\max_a\ EU(a|\gamma, m, q, \mathcal{S}^1_M)$
5. Let $\mathcal{S}^3_M(\omega, q)$ be $\arg\max_m\ EU_M(m|\omega, q, \mathcal{S}^2, \mathcal{S}^2_Q)$

[2] The solution concept relevant to these games is 'perfect Bayesian equilibrium' (Fudenberg and Tirole, 1991).

6. Iterate until $\langle S_M^n, S_Q^{n+1}, S^{n+1}\rangle = \langle S_M^{n+2}, S_Q^{n+3}, S^{n+3}\rangle$

The tuple $\langle S_M^n, S_Q^{n+1}, S^{n+1}\rangle$ constitutes an equilibrium. The final piece of the puzzle is to specify $\mathcal{S}^0$ and $\mathcal{S}_Q^0$. Following the idea from Franke (2011) and others of a *naïve hearer*, we propose default procedures for determining the questioner's behavior which do not rely on any reasoning about the answerer's utility. For our model, this means that questions and actions are chosen only in accordance with the questioner's goal. The questioner can ask about her goal directly via a wh-question, or, if the goal is a non-singleton, employ a YN-question as part of an implicit questioning strategy; this selection is made at random. Then, after receiving an answer to the question, the naïve questioner simply acts on the first relevant condition that is met, e.g., if her goal is γ_t, γ_B or γ_0 from Table 2, and the answer is "Bob's sells turnips", she goes to Bob's. These naïve strategies can be formulated as follows, assuming a one- or two-member goal of form $\langle P \text{ if } \Phi, (R \text{ if } \Psi)\rangle$.

$$\begin{aligned} \mathcal{S}^0(\gamma, m) &= \{P\} \text{ if } [\![m]\!] \to \Phi,\ \{R\} \text{ if } [\![m]\!] \to \Psi,\ \{\emptyset\} \text{ otherwise} \\ \mathcal{S}_Q^0(\gamma) &= \{q \mid [\![q]\!] \in \{\{\Phi, \neg\Phi\}, (\{\Phi, \Psi, \Phi \wedge \Psi, \neg(\Phi \vee \Psi)\})\}\} \end{aligned} \tag{5}$$

We now apply this model to (1), (5) and (6) in Section 1.

3 Explaining the examples

If we work through the iterated reasoning schema given above to simulate contexts for (1), (5) and (6), we find the following:

- The optimal question for a questioner with goal γ_t is always the YN-question, "Does Bob's Market sell turnips?" rather than the broader wh-question, as long as γ_t has a non-zero probability.
- In a world state where Bob's does not sell turnips, but where Freshtown sells turnips and Bob's sells hand soap, the optimal response to the question "Does Bob's Market sell turnips?" is to give the alternative "Freshtown sells turnips..." when $P(\gamma_t)$ is sufficiently high and higher than $P(\gamma_B)$, to give the alternative "Bob's Market sells hand soap" when $P(\gamma_B)$ is sufficiently high and higher than $P(\gamma_t)$, and to give a direct 'no' answer when the two goals are equiprobable.

The easiest way to show this is via computational simulation. We implemented the model in a simplified context where the questioner can have goals γ_t, γ_t' (same as γ_t but with a preference for Freshtown rather than Bob's), γ_B, γ_B' (where hand soap is preferred over turnips), or one of three singleton goals, γ_0^{Bt}, γ_0^{Ft} or γ_0^{Bh}, depending on what item was sought at which store. The questioner can ask about her goal directly, or, for the non-singletons, ask YN-questions which pertain to her goal. The answerer may respond with direct yes/no answers or else convey any potentially relevant proposition of the form $sell(X, y)$. The questioner may respond to the answer with actions $go.to(B)$, $go.to(F)$, $add(t)$, $add(h)$, $add(t)\&go.to(B)$, $add(t)\&go.to(F)$, $add(h)\&go.to(B)$ or $\emptyset$. In order to calculate expected utility, conditional probabilities $P(\omega|m, q, \mathcal{S}_M)$, $P(a|\gamma, m, q, \mathcal{S})$ and $P(m|\omega, q, \mathcal{S}_M)$ were calculated by considering possible outputs of selection procedures, as follows, where the function BOOL takes a proposition and returns 1 if true and 0 if false.

$$P(\omega|m, q, \mathcal{S}_M) = \frac{\text{BOOL}([m \in \mathcal{S}_M(\omega, q)])}{|\{\omega' \mid m \in \mathcal{S}_M(\omega', q)\}|} \tag{6}$$

$$P(a|\gamma, m, q, \mathcal{S}) = \frac{\text{BOOL}([a \in \mathcal{S}(\gamma, m, q)])}{|\mathcal{S}(\gamma, m, q)|} \tag{7}$$

$$P(m|\omega, q, \mathcal{S}_M) = \frac{\text{BOOL}([m \in \mathcal{S}_M(\omega, q)])}{|\mathcal{S}_M(\omega, q)|} \tag{8}$$

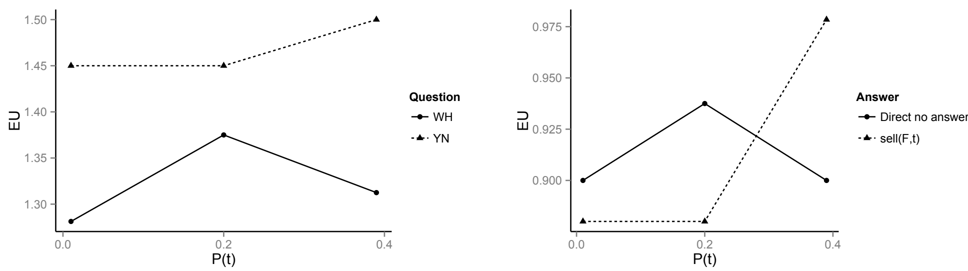

Figure 3: Values of $EU_Q(q|\gamma, \mathcal{S}^1_M)$ for a questioner with goal γ_t deciding between wh-question "Where can I get some turnips?" and YN-question "Does Bob's Market sell turnips?" (left) and values of $EU_M(m|\omega, q, \mathcal{S}^2, \mathcal{S}^2_Q)$ for an answerer having been asked the YN-question in a world state where $sell(B,t)$ is false but $sell(F,t)$ and $sell(B,h)$ are true (right), both plotted by prior probability of the questioner wanting to buy turnips ($P(t) = P(\gamma_t) + P(\gamma'_t)$). For illustration, the space of possible goals was limited to those relevant to (1), (5) and (6), with fixed probability 0.2 given to each singleton goal, and with cost terms $c = 0.1$ and $k = 0.02$.

Finally, $P(\gamma|q, \mathcal{S}_Q)$ is calculated via Bayes' rule, where the likelihood term $P(q|\gamma, \mathcal{S}_Q)$ is calculated by considering possible outputs of the assumed question selection procedure.

$$P(\gamma|q, \mathcal{S}_Q) \propto P(q|\gamma, \mathcal{S}_Q) \cdot P(\gamma) \tag{9}$$

$$P(q|\gamma, \mathcal{S}_Q) = \frac{\text{BOOL}([q \in \mathcal{S}_Q(\gamma)])}{|\mathcal{S}_Q(\gamma)|} \tag{10}$$

The graphs in Fig.3 show the expected utility of some possible questions and answers as a function of the prior probability of wanting to buy turnips (γ_t or γ'_t). We assumed that $P(\gamma_t) = P(\gamma'_t)$ and that $P(\gamma_B) = P(\gamma'_B)$, and reserved a fixed probability for the singletons, such that raising the value of γ_t/γ'_t lowers the value of γ_B/γ'_B, and vice versa. Already on the second iteration ($\mathcal{S}^2$, $\mathcal{S}^2_Q$ and $\mathcal{S}^3_M$) we observe the optimal behavior. A questioner with any preferred outcome (e.g., buying turnips from Bob's rather than Freshtown) does best to ask a YN-question about whether that outcome is possible ("Does Bob's Market sell turnips?"). If one particular goal is sufficiently probable, the answerer should provide either a 'yes' answer, if true, or an indirect answer addressing that goal where appropriate ("Freshtown sells turnips…"). Given an indirect answer, the questioner should infer that the preferred outcome is not possible and go for the runner-up (go to Freshtown).

In the end, this reduces the felicity of the answer "Bob's sells hand soap" for (5) and (6) to a threshold on the prior probability of γ_B. The expected utility of that answer will be the mirror image of the expected utility of "Freshtown sells turnips…" in Fig.3: When the prior probability of γ_t is higher than the prior probability of γ_B, "Bob's sells hand soap" is ruled out, but when the reverse is true, it is in fact the optimal answer. General world knowledge tells us that γ_B is a rather odd goal to have—hand soap does not often serve as a practical substitute for turnips—whereas γ_t seems uncontroversial. This is reflected in the prior probability distribution over goals. Thus, in an out-of-the-blue context, (1) is good while (5) is bad. But we can manipulate the prior by explicitly creating a context where γ_B is more likely, as in (6), in that case, it is acceptable to supply hand soap as an alternative to turnips.

References

Asher, N. and Lascarides, A. (2003). *Logics of Conversation*. Studies in Natural Language Processing. Cambridge University Press.

Benz, A. (2012). Errors in pragmatics. *Journal of Logic, Language and Information*, 21(1):97–116.

Benz, A., Jaeger, G., and van Rooij, R., editors (2006). *Game Theory and Pragmatics*. Palgrave MacMillan, New York.

Benz, A. and van Rooij, R. (2007). Optimal assertions, and what they implicate. A uniform game theoretic approach. *Topoi*, 26(1):63–78.

Büring, D. (2003). On D-trees, Beans, and B-accents. *Linguistics and Philosophy*, 26(5):511–545.

Clark, R. (2011). *Meaningful Games*. MIT Press, Cambridge, MA.

de Marneffe, M.-C., Grimm, S., and Potts, C. (2009). Not a simple yes or no. In *Proceedings of the SIGDIAL 2009 Conference: The 10th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 136–143.

Franke, M. (2011). Quantity implicatures, exhaustive interpretation, and rational conversation. *Semantic and Pragmatics*, 4:1–82.

Fudenberg, D. and Tirole, J. (1991). Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2):236–260.

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

Lai, C. (2012). *Rises all the way up: The interpretation of prosody, discourse attitudes and dialogue structure*. PhD thesis, University of Pennsylvania.

Parikh, P. (2010). *Language and Equilibrium*. MIT Press, Cambridge, MA.

Roberts, C. (1996). Information Structure: Towards an integrated formal theory of pragmatics. In Yoon, J. and Kathol, A., editors, *OSUWPL Volume 49: Papers in Semantics*. The Ohio State University.

Rooth, M. (1992). A theory of focus interpretation. *Natural Language Semantics*, 1:75–116.

Schwarzschild, R. (1999). Givenness, AvoidF, and other constraints on the placement of accent. *Natural Language Semantics*, 7:141–177.

Stevens, J., Benz, A., Reuße, S., and Klabunde, R. (2015). A strategic reasoning model for generating alternative answers. *Proceedings of the 53rd Annual Meeting of the ACL*, pages 534–542.

van Rooij, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy*, 26(6):727–763.

Wagner, M. (2012a). Contrastive topics decomposed. *Semantics and Pragmatics*, 5(8):1–54.

Wagner, M. (2012b). Focus and givenness: A unified approach. In Kučerová, I. and Neeleman, A., editors, *Information Structure: Contrasts and Positions*. Cambridge University Press, Cambridge.