

Indirect answers as potential solutions to decision problems

Jon Scott Stevens & Anton Benz

Center for General Linguistics
Berlin, Germany

Sebastian Reuße, Ronja Laarmann-Quante & Ralf Klabunde

Ruhr University Bochum
Bochum, Germany

Abstract

We present a game-theoretic model of an exchange between a sales agent—an expert with access to a database of information—and a customer who poses yes/no questions to the sales agent in order to help resolve a decision problem. We first provide a game-theoretic description of such an exchange, whereby the sales agent selects an answer to the customer’s question by reasoning about a space of plausible underlying decision problems. We propose a model of both answer generation and interpretation which specifies a solution to this game. The model appropriately selects indirect answers and implicatures for a particular class of yes/no questions. Implicatures can be drawn even when the speaker and hearer have partially misaligned preferences, as long as there is no incentive to lie.

1 Introduction

Indirect answers to yes/no questions come in different flavors: they may entail the direct answer, either semantically as in answer (a) in (1) or when combined with contextually shared world knowledge as in answer (b), or they may allow a direct answer to be inferred probabilistically (de Marneffe et al., 2009) as in answer (c).

- (1) Q: Does the apartment have a garden?
A: a. Apartments in this neighborhood never have gardens.
b. There’s no direct sunlight.
c. Gardens are pretty rare here.

Perhaps more surprising are felicitous answers whose denotation does not entail (semantically, contextually or probabilistically) a direct answer. An example of this is given in (2) in the form of an exchange between a customer looking to rent an apartment and a real estate agent tasked with helping her find the right one.

- (2) CUSTOMER:
Does the apartment have a garden available?
REAL ESTATE AGENT:
It has a beautiful balcony.

Although there is nothing about the semantics of the real estate agent’s response that directly suggests whether there is a garden, the real estate agent’s answer is felicitous under the shared assumption that customers who are interested in a garden might also have their needs met by a balcony. For instance, the customer may want an apartment with a place to grow flowers, in which case a balcony could substitute. The real estate agent’s answer implicates that the answer to the customer’s question is ‘no’, but that the attribute supplied (that there is a balcony) serves as a substitute.

Indirect answers (and indirect speech acts in general, see e.g. Briggs and Scheutz, 2013) can reflect the answerer’s ability to make inferences about the questioner’s *plan* (Allen and Perrault, 1980; Green and Carberry, 1999), that is, how the current question fits into the questioner’s method of accomplishing some goal. In (2), the real estate agent may guess that that goal is to find an apartment with a place to grow flowers, in which case a ‘no’ answer might prompt the follow-up, “does it have a balcony, then?” By looking ahead into this plan, the real estate agent can more efficiently help the customer accomplish her goal.

Also, Benz et al. (2011) suggest that recommender systems (such as a sales agent recommending objects to a customer) can exploit conceptual similarity, i.e. that such a system does better to suggest semantically related alternatives than simply to admit that the customer’s needs cannot be perfectly met. For example, a customer asking for a red sofa might be recommended an orange sofa instead of simply being told that there are no red sofas on hand, or being recommended a white one.

Neither relevance to a plan nor conceptual similarity can be the whole story, however. Firstly, in a dialogue like the one in (2), the sales agent doesn’t know the underlying reason for the customer’s question (i.e. the customer’s plan), and in fact, a number of potential reasons are plausible. Perhaps the customer wants a place to relax in the sun, or perhaps the customer specifically requires a garden. The real estate agent in this case must be able to reason probabilistically about the space of possible intents, such that (a) and (b) are possible answers in (3), but crucially not (c).¹

¹The awkwardness of bringing the customer’s attention to the basement is mitigated in cases where a direct an-

- (3) Q: Does the apartment have a garden?
 A: a. It has a beautiful balcony.
 b. There is a park very close by.
 c. #It has a basement with a large storage area.

Secondly, semantic/conceptual similarity is not a sufficient constraint on indirect answers of this type. While such a constraint could indeed rule out the “basement” answer in (3)—‘garden’ and ‘balcony’ have many properties in common which are not shared with ‘basement’—there must be more to the story, as the following example shows.

- (4) Q: Is there an elementary school nearby?
 A: #There is a university nearby.

Both ‘elementary school’ and ‘university’ are educational institutions, arguably as semantically related as ‘garden’ and ‘balcony’, at least in terms of their basic attributes. A better generalization is that the answer in (4) is inappropriate because elementary schools and universities do not overlap with respect to the problems they solve. In other words, the likelihood that a close-by elementary school and a close-by university will both equally satisfy the customer is simply too low for ‘university’ to be considered as a substitute. They don’t solve the same problem for the customer.

In this paper we provide a formal model of the generation and interpretation of indirect answers of the type seen in (3). Under this model, answers are generated by reasoning about plausible motivations for asking the current question by representing a space of *decision problems* for the questioner (van Rooij, 2003; Benz and van Rooij, 2007). A speaker S (the real estate agent in our example) provides an answer to a question q , which was posed previously by a hearer H (the customer), and based on the answer to q , H chooses a resolution to her decision problem d . This process is modeled as a variant of a *signaling game* (Lewis, 1969) which generates an answer to a question together with a pragmatic interpretation and subsequent decision on the part of the hearer.

We do not assume that the goals of the speaker and the hearer are perfectly aligned. In fact, we assume that the speaker wants to steer the hearer toward a particular action (in this case continuing to consider what the salesperson is offering), and that the speaker is only cooperative in the sense that she has no incentive to tell an outright lie. By deriving the correct interpretations for the indirect answers in (3) from such a model, we show that it is possible for implicatures to arise in non-

answer is also supplied, especially with a particular intonation and some hedging, e.g. “no, there’s no garden, unfortunately... but there is a huge basement with lots of storage!” Although the offering of the basement attribute is directly related to the “no” answer to the customer’s question in this case, it can be seen as a separate speech act and not itself the answer to the question. We are interested here only in the case where (c) is taken as an indirect answer.

cooperative situations (see e.g. Asher and Lascarides, 2013, for a discussion of such situations), as long as honesty is enforced, either by reputation or other factors.

The remainder of this section introduces the notion of *decision problem* used in our analysis. Section 2 develops a game-theoretic description of a sales dialogue exchange, a solution to which can be calculated via an answer generation model which is given in Section 3 and an implicature calculation model which is given in Section 4. Section 5 derives the facts seen in (3) using the current model, and Section 6 concludes with a discussion of possibilities for further research.

Decision problems A decision problem is taken to be a tuple $\langle \Omega, A, U \rangle$, where Ω is a set of possible worlds (where the identity of the real world is unknown to the decider, which for our purposes is the customer), A is the set of possible actions from among which the agent must decide, and U is a utility function encoding the payoff for choosing a particular action in A given a world in Ω . The deciding agent must make inferences about the identity of the real world in order to choose the action from A which is the best candidate to maximize payoff. For current purposes we limit the space of decision problems to those that are in the real estate domain, as in (3), where there is a current “apartment under discussion”, whose attributes are represented in a database visible only to the real estate agent, and where a unit of dialogue consists of a question-answer sequence pertaining to an attribute of the current apartment under discussion.² We assume the following correspondences for this domain.

- A “world” ω corresponds to an apartment, represented as a matrix of attribute values (e.g. +balcony, –garden, etc.) that describe the apartment.
- A consists only of two possible actions: CONTINUE and REJECT, where to CONTINUE is to carry on discussing ω and where to REJECT is to ask to end the discussion of ω .
- U assigns a utility of 1 to CONTINUE and 0 to REJECT in worlds in which CONTINUE is preferred, and assigns 0 to CONTINUE and 1 to REJECT in worlds where REJECT is preferred.

For our purposes, all possible decision problems share these constraints, such that two decision problems d

²This simplified dialogue structure is based on observations from a series of simulated sales dialogues which we conducted between research assistants (trained to use database software and told to play the part of a sales agent) and undergraduate subjects who were instructed to find an apartment in Berlin for a hypothetical friend given some pre-supplied preferences.

These dialogues provided the inspiration for the data in (1)–(4); we observed that our “sales agent” readily used indirect strategies like those seen in (a) and (b) in (3), and that our “customers” had no problem interpreting them.

ω	$U(\cdot, \text{CONT.})$	$U(\cdot, \text{REJECT})$
[+garden, +balcony]	1	0
[+garden, -balcony]	1	0
[-garden, +balcony]	1	0
[-garden, -balcony]	0	1

Table 1: A decision problem $d = \text{'}\omega \text{ has a place to grow flowers'}$

and d' differ only in which subset of the space of possible apartments determines Ω and in the binary values assigned by the utility function U . A possible utility function for a decision problem is represented in Table 1. In plain English, the decision problem represented in Table 1 corresponds to the decision on the part of H between continuing to discuss vs. rejecting the current apartment under discussion (ω) given the requirement that H must have a place to grow flowers in her new apartment.

An easier way to represent such a decision problem is as the set of worlds in which $U(\cdot, \text{CONTINUE}) = 1$. For the problem in Table 1 this is the set $\{[+garden, +balcony], [+garden, -balcony], [-garden, +balcony]\}$. Taking propositions to be sets of worlds, this is equivalent to the proposition, ' ω has a garden *or* ω has a balcony', which is in turn equivalent to the proposition (under some contextual restrictions), ' ω has a place to grow flowers.'

As mentioned above, the sales agent in (3) does not have direct access to the customer's decision problem d , and thus must reason about likely candidates for d when evaluating the felicity of an indirect answer. Therefore the sales agent must represent a space of plausible decision problems. This provides a way of encoding the sales agent's prior world knowledge—she must know in advance that the customer may want to grow flowers, relax outside, etc. Therefore, where the attributes +garden and +balcony are strongly related in virtue of belonging to at least two plausible decision problems (' ω has a place to grow flowers' and ' ω has a place to relax outside'), there is no such relation between +garden and +basement insofar as the agent cannot imagine a plausible underlying decision problem corresponding to ' ω has a garden *or* ω has a basement.' This should rule out the answer "it has a basement with a large storage area" as a possible indirect answer in (3).

2 Dialogue game

The possible space of indirect answers in a sales dialogue exchange like in (3) can be derived by first representing such an exchange as a signaling game \mathcal{G} , one branch of which is represented in Fig. 1, equal to the tuple $\langle \{S, H\}, \Omega, \mathcal{D}, \Delta, \mathcal{Q}, M, \llbracket \cdot \rrbracket, A, U_S, \mathcal{C}, U_H \rangle$ where:

- S and H are the speaker (i.e. sales agent) and hearer (customer), respectively.
- Ω is the set of possible worlds, where a world is




Figure 1: A branch of \mathcal{G} (ω_x is unknown to H , and d_x is unknown to S)

conceived of as an attribute value matrix exhaustively specifying the attributes of a single possible database object.

- \mathcal{D} is the set of shared plausible decision problems, each represented as the set of worlds (i.e. proposition) in which the best decision for the hearer is to continue discussing the current database object ω . ($\mathcal{D} \subset \mathcal{P}(\Omega)$.)
- Δ , a function from $\Omega \times \mathcal{D}$ to the interval $[0, 1]$, is a probability distribution over worlds and decision problems. We assume that Δ is flat, i.e. worlds and decision problems are *a priori* equiprobable, and that Δ provides prior probability terms for Bayesian posterior probabilities which determine expected utility for the speaker and hearer.
- \mathcal{Q} is the set of possible attribute queries, e.g. questions of the form 'what is the value of attribute α in ω ?', where each question is conceived of as a set of possible answers (Hamblin, 1973), or a set of sets of worlds (a set of worlds being a proposition). ($\mathcal{Q} \subset \mathcal{P}(\mathcal{P}(\Omega))$.)
- M is a language of possible *messages*, in this case taken to be the set of possible answers to an attribute query.
- $\llbracket \cdot \rrbracket$ is a denotation function, from M to $\mathcal{P}(\Omega)$.
- A is the set of possible hearer actions, equal to the set $\{\text{CONTINUE}, \text{REJECT}\}$.
- U_S is a function from $\Omega \times M \times A$ to the interval $[0, 1]$, specifying speaker utility.
- \mathcal{C} is a function from M to the interval $[0, 1]$ corresponding to the *cost* of sending a message in M . We assume a higher cost for longer messages and a nominal cost for not providing (the semantic equivalent of) a literal yes/no answer to the hearer's question q (i.e. a member of q).

- U_H is a function from $\Omega \times \mathcal{D} \times A$ to $\{0, 1\}$, specifying hearer utility.

U_S and U_H represent imperfectly aligned preferences on the part of the speaker and hearer, such that: (i) the hearer’s utility is positive only if she continues discussing an apartment which solves her decision problem or rejects one which doesn’t, and (ii) the speaker’s utility is positive only if the hearer chooses to continue. This reflects the fact that in many sales dialogues, the sales agent has strong incentive to sell a particular object, e.g. if it is expensive and she works on commission. Honesty is strictly enforced, encoding a strong role for reputation in the possible answers given by the sales agent. (After all, outright lying to your customers tends to be a bad business decision.) Thus, the speaker’s utility function is positive only if her utterance is true. The cost term $\mathcal{C}(m)$, taken to encode both a higher cost for increased message length and a nominal cost for non-literal answers, is subtracted from base values of 1 and 0. The utility functions for the hearer and speaker, respectively, are as follows.

$$\begin{aligned} U_H(\omega, d, a) &= 1 \text{ if } \omega \in d \ \& \ a = \text{CONTINUE} \\ &= 1 \text{ if } \omega \notin d \ \& \ a = \text{REJECT} \\ &= 0 \text{ otherwise} \end{aligned} \quad (1)$$

$$\begin{aligned} U_S(\omega, m, a) &= 1 - \mathcal{C}(m) \text{ if } \omega \in \llbracket m \rrbracket \ \& \ a = \text{CONT.} \\ &= -\mathcal{C}(m) \text{ otherwise} \end{aligned} \quad (2)$$

U_S does not depend on the hearer’s decision problem d directly, but it depends on the hearer’s action, which is in turn dependent on d . Therefore, S will indeed need to reason probabilistically about d (which is unknown to S) in order to choose an optimal message. The probability of any particular d can be inferred on the basis of the question q . Rather than supplying the model an externally determined probability distribution over decision problems, we assume that all decision problems are *a priori* equiprobable, and that S can infer a conditional probability function over decision problems, $P(\cdot|q)$, via Bayesian reasoning. Bayes’ theorem specifies $P(d|q)$ as the product of $P(q|d)$ and the fraction $P(d)/P(q)$. Assuming $P(d)$ and $P(q)$ to be constants, $P(d)/P(q)$ serves only to normalize the values of $P(q|d)$ for all d in \mathcal{D} . We take $P(q|d)$ to be the probability of randomly selecting q from the set of simple attribute queries (i.e. the subset of \mathcal{Q}) which has the property that at least one answer in that set, if true, would solve d (i.e. make one action dominant over the other). This can be formulated as follows, where the numerator returns 1 iff q contains an answer which solves d and 0 otherwise (the “int” function transforms boolean values into 0 or 1), and where the denominator is the size of the set of all questions in \mathcal{Q} that contain such an answer.

$$P(q|d) = \frac{\text{int}(\exists \phi \in q. \phi \subseteq d)}{|\{q \subset \mathcal{Q} | \exists \phi \in q. \phi \subseteq d\}|} \quad (3)$$

In addition to representing a probability for each possible d , the speaker must also have a belief for each possible d about which action a type- d hearer will take given the content of the speaker’s message—this is the hearer’s *strategy* for selecting an action. By first assuming some fixed strategy for the hearer, the speaker can determine which message has the best chance of leading to an outcome which maximizes the speaker’s own utility. Since the hearer may choose an action at random in some situations, the hearer’s strategy is represented as a probability distribution over actions, $\mathcal{H}(a|d, m)$. We can now specify an expected utility function for the speaker, which returns the weighted average, for all possible underlying decision problems, of the expected payout to the speaker given that decision problem and the hearer strategy \mathcal{H} .

$$\begin{aligned} EU_S(\omega, m|q, \mathcal{H}) &= \\ \sum_{d \in \mathcal{D}} P(d|q) \cdot \sum_{a \in A} \mathcal{H}(a|d, m) \cdot U_S(\omega, m, a) \end{aligned} \quad (4)$$

Similarly, expected utility for the hearer is calculated by assuming a fixed strategy for the speaker. The posterior probability $P(\omega|m, \mathcal{S})$ assigns zero probability to any world in which the speaker would not send m assuming some fixed speaker strategy \mathcal{S} , where $\mathcal{S}(\omega, q)$ outputs a message.

$$EU_H(d, a|m, \mathcal{S}) = \sum_{\omega \in \Omega} P(\omega|m, \mathcal{S}) \cdot U_H(\omega, d, a) \quad (5)$$

The optimal behavior in a dialogue exchange like the one in (3) is specified by an equilibrium in \mathcal{G} , which is a pair of strategies $\langle \mathcal{S}, \mathcal{H} \rangle$ such that each player’s expected utility is maximized by playing their own strategy while assuming the other player’s strategy to be fixed. (In other words, no single player does better by unilaterally deviating from $\langle \mathcal{S}, \mathcal{H} \rangle$.)

We now propose an answer generation procedure for the speaker (sales agent) which specifies a strategy \mathcal{S} which is part of an equilibrium in this game.³ This generation model is shown to correctly predict constraints on indirect answers for a fragment of sales dialogue.

3 Indirect answer generation

Given the game \mathcal{G} introduced in the previous section, an optimal answer for the sales agent in a dialogue exchange of this type is one that maximizes the odds that the customer will be prompted to choose the action CONTINUE. Given the utility structure for \mathcal{G} , a rational customer will choose CONTINUE if the denotation of the hearer’s message is a subset of d . (A rational customer assumes the message to be true, knowing there is no incentive to lie in this situation.) If the

³We include no formal proof here due to space constraints, but it can be shown that the speaker and hearer strategies given in the following two sections correspond to a perfect Bayesian equilibrium (Harsanyi, 1968; Fudenberg and Tirole, 1991).

customer’s underlying decision problem d were known, the speaker’s problem would reduce to that of finding the least costly true message for which this holds. Of course d is not known, and so probabilistic reasoning must be incorporated into the speaker’s strategy. To this end, we first define a set \mathcal{D}'_m of “compatible decision problems” given a message m .

$$\mathcal{D}'_m = \{d \in \mathcal{D} \mid \llbracket m \rrbracket \subseteq d\} \quad (6)$$

The speaker does best by maximizing the *probability of compatibility* (P_{comp}) between a given message m and whichever value of d holds for the hearer.

$$P_{comp}(q, m) = \sum_{d \in \mathcal{D}'_m} P(d|q) \quad (7)$$

The optimal answer for the speaker, then, is a true message which maximizes P_{comp} and minimizes cost.

We assume that the cost function $\mathcal{C}(m)$ grows with the size of the message such that the speaker prefers messages which convey a single attribute of the database object under discussion. Without such an assumption, the optimal message would always be to list all possible solutions to the hearer’s underlying decision problem, rather than choosing one alternative over another, a strategy which seems to be rare in real dialogue situations. Although relatively short conjunctive answers to (3) such as “it has a beautiful balcony, and there is a park nearby” are not infelicitous, we consider for simplicity’s sake only a set $M' \subset M$ of messages which convey a single attribute.

Also, recall that $\mathcal{C}(m)$ encodes a nominal (i.e. tie-breaking) cost for indirect answers such that, if all options are otherwise equal, the speaker prefers simply to provide a literal yes/no answer. This cashes out the intuition that, if the speaker is guaranteed to lose utility by responding to q , that is, if the object under discussion has no chance of being desirable to the customer given her decision problem, the speaker wishes to appear cooperative by providing a direct answer, e.g. “unfortunately there is no garden” over an irrelevant response, e.g. “my sister paints portraits of bees”, which would otherwise yield the same utility for the speaker. Like the enforcement of honesty, this could be seen as a byproduct of reputation, or instead seen as a reflex of coherence requirements or discourse obligations which are introduced by the question (Traum and Allen, 1994).

Putting it all together, the optimal speaker strategy in the game \mathcal{G} is obtained via the following answer generation procedure, for which we first give an informal specification.

1. Let M'_{TRUE} be the subset of M' which excludes all false messages.
2. Obtain the set of messages in M'_{TRUE} that maximize the probability P_{comp} that m is compatible with (i.e. is a subset of) the hearer’s underlying decision problem.

3. Eliminate from that set any messages for which there is a lower cost alternative, where a message has lower cost iff it directly answers the hearer’s question q (i.e. if $\llbracket m \rrbracket \in q$).

4. Output a random message from that new set.

Formally, this can be represented with the following algorithm for a speaker strategy $\mathcal{S}(\omega, q)$, which outputs a message.

1. Let $M'_{TRUE} = \{m \in M' \mid \omega \in \llbracket m \rrbracket\}$
2. Let $\mu = \arg \max_{m \in M'_{TRUE}} P_{comp}(q, m)$
3. Let $\mu' = \{m \in \mu \mid \exists m' \in \mu. \llbracket m' \rrbracket \in q \rightarrow \llbracket m \rrbracket \in q\}$
4. Output some member of μ'

4 Implicature calculation

One *prima facie* peculiarity with the speaker’s strategy $\mathcal{S}(\omega, q)$ is that it filters potential messages by maximizing the likelihood that they will guarantee a CONTINUE action, and does not consider the possibility that a message will make the hearer indifferent between CONTINUE and REJECT, which could in some instances benefit the speaker. For example, one might argue that the answer “there is a basement” for (3) is better for the speaker than the direct answer “there is no garden” under \mathcal{G} , because the former is guaranteed not to address d at all, and thus would make the hearer indifferent, leading to a 0.5 probability of the desired CONTINUE outcome, whereas the latter could result in a guaranteed REJECT outcome if the hearer’s decision problem is simply ‘ ω has a garden.’ In other words, one might argue that a non-sequitur answer is better than one that might prompt a negative reaction from the hearer, even considering any nominal costs for not directly answering the current question.

This proves not to be a problem, however, because a truly rational speaker will take into account the implicatures that the hearer will draw from her message. For example, if the speaker answers the question, “does the apartment have a garden?” with “it has a basement”, the hearer knows that the speaker would have been better off saying “yes” to her question if she could have done so truthfully. Therefore, that answer must be false in the current world. This implicature (that the apartment in fact does not have a garden) makes the “basement” answer equivalent to a “no” answer, except that it bears an increased cost for being a non-literal answer, i.e. for failing to provide a direct ‘yes’ or ‘no’ answer.

This is encoded in the hearer’s expected utility function for \mathcal{G} via $P(\cdot|m, \mathcal{S})$: if the hearer’s beliefs are reasonable, then she will assign zero probability to worlds in which m is not a possible output of $\mathcal{S}(\omega, q)$, thereby drawing the implicature that any messages that would otherwise be better for the speaker are false in ω . This should be made part of the hearer strategy \mathcal{H}

which specifies the space of hearer-optimal responses to m , which in turn determines EU_S , and with it the speaker’s optimal message. Because the speaker considers \mathcal{H} , the speaker knows that an alternative speaker strategy \mathcal{S}' which attempts to trick the hearer with non-sequiturs, is necessarily less optimal than \mathcal{S} .

The aforementioned implicatures⁴, which can serve to provide a direct answer to the hearer’s question, can be calculated by reverse engineering the speaker’s strategy and assuming the falsity of messages that would be more optimal than the observed one if true. This can be accomplished by simply assuming the falsity of any message which has a higher value for P_{comp} than the message that was actually sent. This yields the following algorithm, which we’ll call IMPL, which takes a message m as input and outputs a proposition.

1. Let $\beta = \{\llbracket m' \rrbracket \in M' \mid P_{comp}(q, m') > P_{comp}(q, m)\}$
2. Output $\Omega \setminus \cup \beta$

This outputs only the implicatures drawn from m ; the complete pragmatic interpretation assigned to m by the hearer is $\llbracket m \rrbracket \cap \text{IMPL}(m)$. The hearer’s strategy, then, can be specified as follows.

$$\begin{aligned} \mathcal{H}(\text{CONT.} \mid d, m) &= 1 \text{ iff } \llbracket m \rrbracket \cap \text{IMPL}(m) \subseteq d \\ &= 0 \text{ iff } \llbracket m \rrbracket \cap \text{IMPL}(m) \cap d = \emptyset \\ &= 1/2 \text{ otherwise} \\ \mathcal{H}(\text{REJECT} \mid d, m) &= 1 - \mathcal{H}(\text{CONT.} \mid d, m) \end{aligned} \quad (8)$$

5 Example

We now use the answer generation and implicature calculation procedures given above to derive the facts in (3), reproduced below as (5), given a fragment of world knowledge.

- (5) H: Does the apartment have a garden?
S: a. It has a beautiful balcony.
b. There is a park very close by.
c. #It has a basement with a large storage area.

Although a decision problem is formally represented as the set of worlds in which the decision problem is solved, any decision problem consistent with the sales agent’s world knowledge can also be represented as a complex preference statement, e.g. ‘ ω has a balcony or ω has a garden.’ While conjunctive decision problems are logically possible, we only consider disjunctive ones, i.e. decision problems that can be phrased as ‘ ω has value x for attribute α or ω has value y for attribute β .’ Accordingly, we use a short-hand set notation, such that $\{+\alpha, +\beta\}$ means the proposition ‘ ω is

⁴We emphasize that no claims are made about the generalizability of the current model to other kinds of implicatures, e.g. those which arise in purely cooperative dialogue situations.

$+\alpha$ or ω is $+\beta$.’ Using this notational shortcut, we begin to build a fragment of world knowledge with which to derive example (5).

Consider a fragment of a context for example (5) where there are only four apartment attributes represented in the database: (i) whether there is a garden available, (ii) whether there is a balcony, (iii) whether there is a park nearby⁵, and (iv) whether there is a basement storage area available. To abbreviate, we use ‘B’ for *balcony* and ‘K’ (as in German *Keller* ‘basement’) for *basement*. Table 2 shows the possible worlds.

The space of possible questions is: $\mathcal{Q} = \{\text{‘What is the value for attribute } \alpha \text{ in } \omega\text{?’}\}$, where $\alpha \in \{\text{garden, balcony, park, basement}\}$, and the current question under discussion is $q = \text{‘What is the value for attribute } \textit{garden} \text{ in } \omega\text{?’}$, equivalent to the set containing: (i) the set of worlds in which ω has a garden, and (ii) the set of worlds in which ω does not have a garden.

Table 3 shows the decision problems deemed to be reasonable in this fragment, along with their conditional probabilities. We consider the following possibilities: the customer either wants a garden, balcony, park or basement specifically, or else a place to grow flowers, a place nearby to go for a walk outside, or just a place to relax outside.

Table 4 specifies a space of possible utterances, all specifying a $+/-$ value for a single attribute. Table 5 shows binary truth values for whether m is in d for all m/d combinations, as well as the conditional probabilities for each d , the value of $P_{comp}(q, m)$ for each message, and whether each m is a literal answer (that is, whether the denotation of m is in q). Putting it all together, we obtain the following dominance hierarchy of best messages. The speaker should use the best message that also happens to be true.

$$m_G \succ m_B, m_P \succ m_{-G}$$

In plain English, we have obtained the following strategy for our sales agent for this particular dialogue exchange.

1. If ω has a garden say, “there is a garden.”
2. Else, if ω has a balcony say, “there is a balcony”, or if ω has a park nearby say, “there is a park nearby.”
3. Else, say, “there is no garden.”

Finally, we can use the hearer’s representation of the speaker’s strategy to derive the indirect meaning carried by the speaker’s answer.⁶

⁵It is a simplification to treat this as a binary variable; in actuality, the database would contain a distance value to the nearest park, with the definition of “nearby” left to the judgment of the interlocutors.

⁶Note that the implicature algorithm in Section 4 assumes that the hearer only considers P_{comp} , and not the cost for the speaker. This allows the hearer to derive correct implica-

Ω	AVM	Ω	AVM
ω_{GBPK}	[+garden, +balcony, +park, +basement]	ω_{GBP}	[+garden, +balcony, +park, -basement]
ω_{GBK}	[+garden, +balcony, -park, +basement]	ω_{GPK}	[+garden, -balcony, +park, +basement]
ω_{BPK}	[-garden, +balcony, +park, +basement]	ω_{GB}	[+garden, +balcony, -park, -basement]
ω_{GP}	[+garden, -balcony, +park, -basement]	ω_{GK}	[+garden, -balcony, -park, +basement]
ω_{BP}	[-garden, +balcony, +park, -basement]	ω_{BK}	[-garden, +balcony, -park, +basement]
ω_{PK}	[-garden, -balcony, +park, +basement]	ω_G	[+garden, -balcony, -park, -basement]
ω_B	[-garden, +balcony, -park, -basement]	ω_P	[-garden, -balcony, +park, -basement]
ω_K	[-garden, -balcony, -park, +basement]	ω_\emptyset	[-garden, -balcony, -park, -basement]

Table 2: Worlds

\mathcal{D}	Attributes	Plain English	$P(\cdot q)$
d_G	{+garden}	'Access to a garden'	$\frac{9}{14}$
d_B	{+balcony}	'A balcony'	0
d_P	{+park}	'A nearby park'	0
d_K	{+basement}	'A basement'	0
d_F	{+garden, +balcony}	'A place to grow flowers'	$\frac{3}{14}$
d_W	{+garden, +park}	'A place to walk outside'	$\frac{3}{14}$
d_R	{+garden, +balcony, +park}	'A place to relax outside'	$\frac{2}{14}$

Table 3: Plausible decision problems

M'	English	$\llbracket \cdot \rrbracket$
m_G	"There is a garden"	$\{\omega_{GBPK}, \omega_{GBP}, \omega_{GBK}, \omega_{GPK}, \omega_{GB}, \omega_{GP}, \omega_{GK}, \omega_G\}$
m_B	"There is a balcony"	$\{\omega_{GBPK}, \omega_{GBP}, \omega_{GBK}, \omega_{BPK}, \omega_{GB}, \omega_{BP}, \omega_{BK}, \omega_B\}$
m_P	"There is a park nearby"	$\{\omega_{GBPK}, \omega_{GBP}, \omega_{BPK}, \omega_{GPK}, \omega_{PK}, \omega_{BP}, \omega_{GP}, \omega_P\}$
m_K	"There is a basement area"	$\{\omega_{GBPK}, \omega_{BPK}, \omega_{GBK}, \omega_{GPK}, \omega_{PK}, \omega_{GK}, \omega_{BK}, \omega_K\}$
m_{-G}	"There is no garden"	$\{\omega_{BPK}, \omega_{BP}, \omega_{BK}, \omega_{PK}, \omega_B, \omega_P, \omega_K, \omega_\emptyset\}$
m_{-B}	"There is no balcony"	$\{\omega_{GPK}, \omega_{GP}, \omega_{GK}, \omega_{PK}, \omega_G, \omega_P, \omega_K, \omega_\emptyset\}$
m_{-P}	"There is no park nearby"	$\{\omega_{GBK}, \omega_{GB}, \omega_{GK}, \omega_{BK}, \omega_G, \omega_B, \omega_K, \omega_\emptyset\}$
m_{-K}	"There is no basement area"	$\{\omega_{GBP}, \omega_{GB}, \omega_{GP}, \omega_{BP}, \omega_G, \omega_B, \omega_P, \omega_\emptyset\}$

Table 4: Messages

	m_G	m_B	m_P	m_K	m_{-G}	m_{-B}	m_{-P}	m_{-K}	$P(\cdot q)$
$\llbracket \cdot \rrbracket \subseteq d_G$	1	0	0	0	0	0	0	0	$\frac{9}{14}$
$\llbracket \cdot \rrbracket \subseteq d_B$	0	1	0	0	0	0	0	0	0
$\llbracket \cdot \rrbracket \subseteq d_P$	0	0	1	0	0	0	0	0	0
$\llbracket \cdot \rrbracket \subseteq d_K$	0	0	0	1	0	0	0	0	0
$\llbracket \cdot \rrbracket \subseteq d_F$	1	1	0	0	0	0	0	0	$\frac{3}{14}$
$\llbracket \cdot \rrbracket \subseteq d_W$	1	0	1	0	0	0	0	0	$\frac{3}{14}$
$\llbracket \cdot \rrbracket \subseteq d_R$	1	1	1	0	0	0	0	0	$\frac{2}{14}$
P_{comp}	1	$\frac{5}{14}$	$\frac{5}{14}$	0	0	0	0	0	
$\llbracket m \rrbracket \in q$	1	0	0	0	1	0	0	0	

Table 5: Optimality of messages if true. An answer is sub-optimal if there is a true answer within a group to its left, as indicated by dashed lines. Within each grouping, a message is optimal only if either (i) it is a literal answer, or (ii) there are no literal answer alternatives that could be used.

1. If the speaker answers either “there is a balcony” or “there is a park nearby”, then there is no garden.
2. If the speaker answers, “there is no garden”, then there is no garden, balcony, or park nearby.

While the first implicature is clear from example (5), the second seems disputable. Is it really the case that a “no” answer implicates that there are no possible substitute solutions to the hearer’s problem? One gets the intuition that this is only the case under strong common knowledge assumptions about how willing the sales agent is to query the database for multiple attributes to find alternatives. This willingness undoubtedly depends on personality traits which must be attributed to the sales agent by the customer (see Walker et al., 1997, and related work), for example laziness. If there is the possibility of a lazy sales agent, for example, the hearer will be less ready to draw the second implicature, because she cannot be certain that the sales agent has checked the database to see whether there is a balcony or a park nearby. But the first implicature is a safe bet in any case, because the customer can be sure that the sales agent has checked to see whether there is a garden, since that attribute was the target of the customer’s question. This intuition could be cashed out within the current framework as an effect of uncertainty on the hearer’s part about the cost function \mathcal{C} . The first implicature, but not the second, is calculable under any reasonable value for $\mathcal{C}(m)$. Further investigation of such effects must be left to future research.

6 Discussion

We have presented a game-theoretic description of a yes/no question-answer exchange between a sales agent and a customer in which the sales agent (speaker) must consider the customer’s (hearer’s) underlying decision problem which motivated her question before supplying an answer. We have proposed speaker and hearer strategies designed to find equilibria in this game. The resulting model has three key properties. First, the speaker has motivation to produce indirect answers insofar as those answers serve as potential alternative solutions to the hearer’s underlying problem. Second, the hearer can infer a direct answer to her question from an indirect one, even if no entailment relationship exists between the speaker’s response and a direct yes/no answer. Third, these inferences are possible even when the speaker and hearer have partially misaligned goals.

The partial misalignment of preferences in the model represents a move beyond traditional Gricean accounts

tures even for “off-equilibrium” messages. For example, it is never optimal, due to cost, for the speaker to answer, “it has a basement”, but if the speaker did so for some reason, the hearer would still reason that there is no balcony or park or garden. If the hearer considered cost as well, then she would reason that both “there is a garden” and “there is no garden” are false—a logical contradiction.

of implicature into cases where the speaker has some incentive to be non-cooperative (what Asher and Lascarides, 2013, call “strategic conversation”). Under our model, implicatures arise in non-cooperative situations as long as honesty is enforced, either through reputation or through other means. In a sales dialogue like the one studied here, the sales agent wants the customer to choose the action CONTINUE regardless of whether the object being sold is truly optimal for the customer, and yet if she cannot lie, the sales agent behaves as if she is fully cooperative. The reason for this is that, if the salesperson’s goals are known by the customer, then the customer will draw implicatures from any indirect answers by assuming the falsity of any answers that would have been more optimal given those goals. Misleading irrelevant answers become no better than answers which directly prompt an unwanted action from the customer—the customer is too smart to be swindled.

This work is intended as a starting point for a more general inquiry into such phenomena in dialogue. Further research is required to assess the generalizability of the current approach to different dialogue situations, as well as the validity of our assumptions regarding how world knowledge is represented in the dialogue model. For example, we currently posit that the interlocutors have access to a discrete space of plausible decision problems (\mathcal{D}), such that extremely unlikely question motivations (e.g. $d = \omega$ has a place for my cat, who only likes balconies and basements, to take naps’) are not considered. It is important to determine whether this aspect of our approach is fully justified, and, if so, how such a discrete space might be built and represented from prior experience.

Finally, future research will determine whether such considerations can be practically implemented within an automated dialogue system. Namely, while the algorithm in Section 3 can be used to select from among a finite space of possible answers to a yes/no question, the output relies crucially on the space of possible decision problems. It remains to be assessed whether a richer space could be empirically obtained, and whether such a space would yield realistic answers to a wider variety of questions in a sales dialogue.

References

- Allen, J. F. and Perrault, C. R. (1980). Analyzing intention in utterances. *Artificial intelligence*, 15(3):143–178.
- Asher, N. and Lascarides, A. (2013). Strategic conversation. *Semantics and Pragmatics*, 6(2):1–62.
- Benz, A., Bertomeu, N., and Strelakova, A. (2011). A decision-theoretic approach to finding optimal responses to over-constrained queries in a conceptual search space. In *SemDial 2011: Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue*, pages 37–46.

- Benz, A. and van Rooij, R. (2007). Optimal assertions, and what they implicate. a uniform game theoretic approach. *Topoi*, 26(1):63–78.
- Briggs, G. and Scheutz, M. (2013). A hybrid architectural approach to understanding and appropriately generating indirect speech acts. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, (forthcoming).
- de Marneffe, M.-C., Grimm, S., and Potts, C. (2009). Not a simple yes or no: Uncertainty in indirect answers. In *Proceedings of the SIGDIAL 2009 Conference: The 10th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 136–143. Association for Computational Linguistics.
- Fudenberg, D. and Tirole, J. (1991). Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2):236–260.
- Green, N. and Carberry, S. (1999). Interpreting and generating indirect answers. *Computational Linguistics*, 25(3):389–435.
- Grice, H. P. (1975). Logic and conversation. In Cole, P. and Morgan, J., editors, *Syntax and Semantics vol. 3*. Academic Press, New York.
- Hamblin, C. (1973). Questions in Montague english. *Foundations of Language*, 10:41–53.
- Harsanyi, J. C. (1968). Games of incomplete information played by ‘Bayesian’ players, part ii. *Management Science*, 14(5):320–334.
- Lewis, D. (1969). *Convention: A Philosophical Study*. Cambridge University Press, Cambridge.
- Traum, D. R. and Allen, J. F. (1994). Discourse obligations in dialogue processing. In *Proceedings of the 32nd annual meeting on Association for Computational Linguistics*, pages 1–8.
- van Rooij, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy*, 26(6):727–763.
- Walker, M. A., Cahn, J. E., and Whittaker, S. J. (1997). Improvising linguistic style: Social and affective bases for agent personality. In *Proceedings of the first international conference on autonomous agents*, pages 96–105. ACM.