Why do you ask? The informational dynamics of questions and answers

Robert D. Hawkins, Noah D. Goodman
Stanford University

Asking questions is one of our most efficient and reliable means of learning about the world. Yet we do not often pose these questions to an impartial oracle; we ask cooperative social partners, in dialogue. In this paper, we aim to reconcile formal models of optimal question asking and answering with classic effects of social context. We begin from the observation that question-answer dialogue is motivated by a two-sided asymmetry in beliefs: questioners have a private goal but lack goal-relevant information about the world, and answerers have private information but lack knowledge about the questioner's goal. We formalize this problem in a computational framework and derive pragmatic questioner and answerer behavior from recursive social reasoning. Critically, we predict that pragmatic answerers go beyond the literal meaning of the question to be informative with respect to inferred goals, and that pragmatic questioners may therefore select questions partly to signal their goals. We evaluate our pragmatic models against asocial models in two ways. First, we present computational simulations accounting for three classic answerer effects in psycholinguistics. We then introduce the Hidden Goal paradigm for experimentally eliciting questioner and answerer behavior in scenarios where there is uncertainty about the questioner's goal. We report data from three experiments in this paradigm and show how our core computational framework can be composed with more sophisticated question semantics, hierarchical goal spaces, and a persistent state over which extended dialogue can unfold. We find that social inference is needed to account for critical aspects of the data.

Keywords: questions and answers, semantics and pragmatics, social reasoning, social learning, language, computational modeling, active learning

Asking questions is one of our most valuable methods of gathering information and learning from others. At a very young age, children begin asking questions in order to solve problems and test lay theories (Legare, Mills, Souza, Plummer, & Yasskin, 2013; Chouinard, 2007; Callanan & Oakes, 1992). Robots that ask questions to clarify commands (Deits et al., 2013) or obtain assistance with a task (Fong, Thorpe, & Baur, 2003) make for better collaborators. Students who ask better questions in the classroom tend to be more successful (Graesser & Person, 1994), and scientists who design

their surveys to ask better questions can obtain more informative results (Clark & Schober, 1992).

How, then, do people decide which question to ask in a given situation? One particularly influential rational account formalizes human inquiry as an *optimal experiment design* (OED) problem (Oaksford & Chater, 1994; Nelson, 2005; Myung & Pitt, 2009; Gureckis & Markant, 2012; Coenen, Nelson, & Gureckis, 2018), where people choose a question to maximize the knowledge they expect to gain from knowing its answer.

A central component in the OED framework is the distribution over possible answers questioners expect to receive after asking a particular question. This distribution is typically assumed to be determined by the ground truth of the task. For instance, in a category learning task (Markant & Gureckis, 2014), the possible answers to a query about the category membership of a

This report is based in part on work presented at the 37th Conference of the Cognitive Science Society. Correspondence should be addressed to Robert X.D. Hawkins, e-mail: rxdh@stanford.edu

particular object are simply the possible categories.

(1) Q: Which category is this?

A: It's a *dax*.

In other words, an OED questioner expects answers to come from an impartial oracle that responds truthfully but literally with the requested information. This answering strategy has been implemented in sophisticated semantic parsers that extract the logical form of the question and either execute it on the world state or efficiently search a database for an entity that satisfies it (Rothe, Lake, & Gureckis, 2017).

While this 'oracle' assumption has been sufficient to capture a wide variety of inquiry phenomena in active learning domains, including causal learning (Bramley, Lagnado, & Speekenbrink, 2015) and sequential information search (Ruggeri & Lombrozo, 2015), it is violated precisely under the social learning conditions where inquiry may be most prevalent and useful: asking linguistic questions to other social agents in cooperative dialogue. Answerers regularly go *beyond* the literal meaning of a question in helpful, context-dependent ways.

A simple example is an indirect speech act, where questioners typically will not be satisfied with a literal yes/no answer to their yes/no question (Clark, 1979):

(2) Q: Do you know the time?

A: It's a quarter past 3.

But answerers will go beyond the literal meaning even for *direct* questions. A corpus study found that 14% of responses to direct yes/no questions used a full declarative statement instead of a simple "yes" or "no" (de Marneffe, Grimm, & Potts, 2009):

(3) Q: Is it in Dallas?

A: Uh, it's in Lewisville.

Questions like "where are you?" permit answers at many degrees of specificity: *the United States*, *my apartment*, and *by the big tree* are each perfectly appropriate in some context and highly inappropriate in others (Potts, 2012). Or, if a child asks "What's that?" while pointing at a common household object, a parent's response will likely be much different than if one of their adult friends asked the same question. Identification questions like "who is X?" can be resolved in many ways (Boër & Lycan, 1975; Gerbrandy, 2000;

Aloni, 2005). For example, if an undergraduate asked "Who is Noam Chomsky?" in an introductory course, it would be appropriate to respond "The influential scholar who wrote *Aspects of the Theory of Syntax*" or "The father of modern linguistics." If a potential donor asked the same question at a fundraiser that Chomsky was attending, though, it would be more appropriate to answer by pointing at him in the crowd.

Specific *wh*-questions like "Who passed the examination?" may also receive different answers in different contexts. If an instructor asks their teaching assistants this question in a meeting, they likely want an *exhaustive* list, while a group of students talking amongst themselves may only care about a *selective* list of relevant entities (Schulz & Van Rooij, 2006). In some domains, radically underspecified questions are the rule rather than the exception. Teachers, doctors, technical troubleshooters, concierges, and others in the service industry are regularly faced with the challenge of responding helpfully even when the questioner themselves may not be clear about precisely what information they need:

(4) Q: What's wrong with my answer to #1?

A: Well, let me remind you how to compute a derivative...

These cases present two immediate challenges for computational models of inquiry. First, if people do not simply respond literally to questions, how *do* they choose answers? What makes an answer more or less helpful? To extend rational models of question asking to social settings, we must first formulate a social and context-sensitive model of question *answering* that captures the relevant psycholinguistic phenomena. Second, if questions are not just distributions over their possible (literal) answers, as OED models have assumed, then what is the meaning of a question? This problem has been addressed extensively in a largely independent literature in linguistics focusing on the semantics of questions and answers (Groenendijk & Stokhof, 1984; Ginzburg, 1995; Van Rooy, 2003; Ciardelli, Groenendijk, & Roelofsen, 2013). These formal accounts have successfully worked through a number of thorny semantic problems, but a psychological account of question *pragmatics* based on these insights has remained elusive.

This paper presents a unified computational framework that synthesizes and extends accounts of inquiry in linguistics and psychology to address these challenges. We begin with the idea that the functional conditions for inquiry in social contexts emerge from a *two-sided* asymmetry in beliefs between social agents. One agent, the questioner, has a private *goal* but lacks goal-relevant information about the world. Another agent, the answerer, has private information about the *world* but critically lacks foreknowledge of the questioner's goals. If either of these conditions were not satisfied, there would be no need for a question: the questioner would either already have the information they need, or the answerer could provide the relevant information without being asked. This formulation suggests a dependence between the two agents and a theoretical view where question asking and answering behavior is derived from deeply intertwined *social inferences*.

We make this view precise below in a recursive probabilistic model of language understanding, and show that existing OED models of question asking are a special case at lower levels of recursion. By adding additional layers of recursive reasoning, we derive a pragmatic answerer that uses the question to infer the questioner's underlying goal and provides an answer that attempts to address this goal. A pragmatic questioner, in turn, can then select questions that more clearly signal her goals and intentions. This model supports an alternative view of the function of a question: not as a request for information, but as a signal about private goals (De Ruiter, 2012).

The rest of this paper is structured as follows. After reviewing previous experimental and formal modeling efforts we lay out the details of our computational framework, formalizing different hypotheses in a series of questioner and answerer models. We then evaluate these models in two ways: through a set of computational experiments capturing three classic *answerer* sensitivity effects, and through three novel behavioral experiments addressing a longstanding paucity of data on social reasoning in *questioner* behavior.

Our behavioral experiments introduce and develop the *Hidden Goal* paradigm, a novel family of cooperative communication tasks inducing uncertainty for the answerer over the questioner's private goal. In this paradigm, different questioner and answerer models can be rigorously distinguished through both quali-

tative predictions falsifying previous models as well as quantitative model comparisons based on relative fit. Additionally, each experiment demonstrates how our core computational framework can be composed and extended with additional elements to capture more sophisticated behaviors, including more complex question semantics (Exp. 1), hierarchically structured goals (Exp. 2), and multi-round dialogue (Exp. 3).

Evidence for goal-sensitivity in question answering

While we do not attempt to overview all work on questions and answers in psycholinguistics, a number of studies have provided direct empirical evidence that answerers are both sensitive to a questioner's goals and attempt to be informative with respect to those goals. In Clark (1979), for instance, researchers called liquor merchants and asked "Does a fifth of Jim Beam cost more than \$5?" Before asking the question, they provided some context for their call by saying either "I want to buy some bourbon" (the *uninformative* condition) or "I've got \$5 to spend" (the *five dollar* condition). These contexts were designed to establish different speaker goals. In the uninformative condition, where the goal was to buy whiskey, the exact price would be relevant information; in the five dollar condition, where the goal was simply to find out whether or not the whiskey was affordable, a yes/no answer would suffice. If merchants inferred these goals from the context signal and responded with respect to these goals, we would expect different types of answers in the two conditions. Indeed, merchants gave a literal yes/no answer more often in the latter condition than the former, where an exact price was more common.

Goals can also be inferred from non-linguistic environmental cues. For instance, Der Henst, Carles, and Sperber (2002) investigated answers to questions like "Do you have the time?" that permit answers with different degrees of approximation to the true time (see also Gibbs Jr & Bryant, 2008). In a baseline condition, participants approached in public typically rounded their answers to the nearest 5 or 10 minute interval even when they were wearing a digital watch. When the experimenter mentioned that they had an appointment at a given time, however, answers became more precise as the true time approached the appointment time, indicating sensitivity to implicit time-related goals like being on time.

These studies suggest that *goal-relevance* is a critical feature of question-answering: answerers are sensitive to questioner goals and go beyond the literal meaning of a question to help the questioner achieve them. To our knowledge, this empirical work has focused entirely on the answerer, with the experimenter holding the question constant and investigating the effect of different contexts. To date, empirical evidence about how social context affects the *questioner's* choice has been scarce. Next, we turn to a set of existing formal proposals about how goals may be incorporated into models of optimal question asking and answering.

Formal approaches to goal-sensitivity

The earliest attempts to construct artificial agents who could answer natural-language questions in a goal-sensitive manner took place within the artificial intelligence community (Simmons, 1965; Lehnert, 1977; Allen & Perrault, 1980; Green & Carberry, 1994; Mollá & Vicedo, 2007). Many of these classic systems were essentially domain-restricted interfaces for databases, with applications in customer service, technical support, health care, or financial consulting. Other systems, however, built on insights from early work in cognitive science to represent questioner intentions as schemas or plans. For instance, the system designed by Allen and Perrault (1980) used a set of logical rules to generate a space of possible plans the questioner may be pursuing, searched this space using a set of heuristics, identified obstacles to fulfilling likely plans, and finally took a set of actions to address the obstacles to help the questioner. Because these systems relied heavily on problem-specific rules and heuristics, it has not been straightforward to generalize them into a broader computational framework for inquiry.

A more promising thread of formal work on goal-sensitivity in inquiry has been developed in the field of linguistics (see Ciardelli, Groenendijk, & Roelofsen, 2019, for a thorough review of this literature). It is a fundamental problem in formal semantics to precisely define the meaning of questions and answers. Like the OED framework in psychology, early approaches in linguistics simply identified the meaning of a question with a set of possible answers. (Hamblin, 1973; Karttunen, 1977). Seminal work by Groenendijk and Stokhof (1984), however, pointed out several limitations of these early approaches and formulated an in-

fluent alternative where questions have a single exhaustive answer in each world. Because each possible world has a unique answer, questions induce a *partition* or *coarse-graining* over the space of possible worlds, where all worlds in a particular cell of the partition have the same answer. This partition, then, is taken to be the formal meaning of a question.

While subsequent developments of partition-based question semantics have grappled with increasingly subtle phenomena, the meanings they assign are by definition context-independent. A question like “Where are you?” will always induce the same partition, demanding the same answers. Several accounts have addressed this limitation by arguing that the ‘goodness’ of an answer for resolving a question ought to be defined relative to the questioner’s current goals (e.g. Ginzburg, 1995; Asher & Lascarides, 1998). The most relevant formalization of this idea is due to Van Rooy (2003), who elegantly incorporated goals using a decision-theoretic framework. In this framework, questioners face a *decision problem* under uncertainty, and may ask a question to gain information before acting.

It can be shown that a rational questioner should choose the question yielding the maximal expected increase in payoff for the decision problem they are facing, similar to OED accounts. An answerer agent can similarly be formulated, who attempts to select the maximally useful answer *with respect to that decision problem*. Because both questioner and answerer expect one another to behave rationally, the questioner can assume their partner will take their decision-problem into account when interpreting the question. Van Rooy (2003) argues further that the literal meaning of a question is underspecified: it only resolves to a particular partition given the contextual parameter of a specific decision problem.

We take this proposal as a starting point for approaching context-sensitivity in our social inference framework below. Our account, however, seeks to address a critical limitation the earlier proposal faces as a cognitive model of inquiry: how does the answerer know what the decision problem is? Van Rooy (2003) and subsequent accounts in this tradition (e.g. Benz, 2006) implicitly assume that the decision problem is *already mutual knowledge shared by both questioner and answerer*: “On the assumption that the answerer knows what the agent’s decision problem is, the

coarser-grained, and thus partial, answer is [...] preferred.” (p. 738).

In our view, this assumption undermines the two-sided functional conditions motivating inquiry. A cognitive account of answerer goal-sensitivity must provide an explicit mechanism for how the answerer learns the goal. We will suggest that the necessary mechanism is (Bayesian) social inference: the answerer inverts a model of the questioner to reason about their likely goals, conditioned on their question utterance (and other context). Our framework can thus be viewed as introducing the cognitive mechanisms needed to extend Van Rooy’s account to the more general case where answerers have a priori uncertainty over the questioner’s goal¹. This extension has a number of broad consequences and forces a re-evaluation of the meaning of a question, which we will return to in the General Discussion. In addition, by formulating our model within the Rational Speech Act framework we expose connections to recent advances in computational psycholinguistics.

A Rational Speech Act framework for question and answer pragmatics

Overview

Broadly speaking, an agent must solve two problems to participate in a communicative exchange. First, they must *listen*, interpreting the meaning of their partner’s utterance and updating their own beliefs accordingly. Then they must *speak*, producing an utterance that informatively addresses the current goal or topic of conversation. At either stage, the utterance may be a statement, a question, or another more exotic kind of speech act.

The listening problem and the speaking problem have been addressed separately in prior work within the Rational Speech Act (RSA) framework (Frank & Goodman, 2012; Goodman & Stuhlmüller, 2013; Goodman & Frank, 2016), in which language understanding is formalized as recursive Bayesian inference. Pragmatic speaker agents choose utterances to minimize their partner’s surprisal about the true state of the world – a basic epistemic goal – and pragmatic listener agents interpret utterances by inverting the speaker model. This basic framework has been used to account for a diverse selection of pragmatic phenomena including scalar implicature (Goodman & Stuhlmüller, 2013), manner impli-

cature (Bergen, Levy, & Goodman, 2016), generic language (Tessler & Goodman, 2018), and interpretation of context-sensitive adjectives like “tall” or “cheap” (Lassiter & Goodman, 2015).

RSA models have also recently been extended to incorporate the critical notion of *goal-relevance* into the speaker’s utility (Roberts, 1996; Wilson & Sperber, 2012). Instead of attempting to be maximally informative about the *full* state of the world, a goal-relevant speaker ignores irrelevant dimensions and only attempts to be informative about relevant aspects. The projection function that collapses irrelevant dimensions is often called a Question Under Discussion (QUD) and is incorporated as a parameter of the speaker that a pragmatic listener may have uncertainty over and make inferences about. RSA models with relevance have been used to capture non-literal language use like metaphor, irony, and hyperbole. In the case of hyperbole, for instance, a speaker who says “It took a million years to get a table” may be intending to be informative about the topic of their affective state (e.g. their frustration) rather than the exact time it took to get a table (Kao, Wu, Bergen, & Goodman, 2014).

Although the social interplay between questioner and answerer may rely on that of the *speaker* and *listener* agents in previous RSA models, the minimal form of dialogue introduced by a question-answer exchange presents fundamentally new challenges. Here, we build upon these recent modeling advances in several ways. First, while the role of a speaker in prior models is limited to production, and the role of listener is limited to interpretation, a question-answer exchange requires unifying the listening problem and speaking problem in a single agent. Addressing this problem in a rational model requires deriving the appropriate contingency between linguistic input and output. Different kinds of input shift an agent’s beliefs in different ways, motivating different kinds of responses.

Second, the utilities used for previous RSA speakers cannot straightforwardly capture the utility of uttering a question, since questions don’t provide direct information about the world to the listener. While a declarative

¹The only previous work we are aware of that considers how the answerer might determine the questioner’s decision problem from a space of alternatives is by Stevens, Benz, Reuße, Laarmann-Quante, and Klabunde (2014) which treats a limited case under a game-theoretic perspective.

Figure 1. Relationships between the questioner and answerer models we consider.

statement in our framework provides evidence about the true state of the world, as in previous models, we conjecture that a question is intended to be informative about the *speaker's goals*. When it is the listener's turn to speak, they can then generate statements that better address the goal-relevant aspects of the world. In the next section, we review the basic RSA framework and explore several different proposals for how agents interpret and respond to questions within this framework.

Preliminaries

Our computational framework begins with the sets of world states, goals, and utterances (both questions and answers). For concreteness, we use the scenario from Clark (1979) as an example throughout this section:

$\mathcal{W}$: a set of possible states of the world, such as the true price of a bottle of whiskey $\{\$1, \$2, \dots, \$10\}$. This may include arbitrarily many features of the world, such as the current weather, the location of nearby cafés, or the kinds of payment that the store accepts.

$\mathcal{G}$: a set of possible underlying questioner goals, such as learning the actual price of the whiskey or learning whether or not one can afford to buy the whiskey at all. Following previous RSA models, we formalize the key notion of a goal $g \in \mathcal{G}$ as a projection function $g : \mathcal{W} \rightarrow \mathcal{W}$ that maps a world state w to a particular feature or set of features that the questioner cares about, which we

denote $\hat{w} = g(w)$. For example, when world states are prices, a customer's goal of learning the actual price of the whiskey is given by the identity function $g_=(w) = w$, while the goal of learning whether it is affordable is given by

$$g_<(w) = \begin{cases} 1 & w < 5 \\ 0 & \text{otherwise} \end{cases}$$

such that all worlds where the agent can afford the whiskey are mapped onto one state, and all other worlds are mapped onto a different state.

$\mathcal{Q}$: a set of possible question utterances, such as “Does a fifth of Jim Beam cost more than \$5?” and alternatives like “Do you accept credit cards?” We take the literal meaning $\llbracket q \rrbracket$ of a question to be of the same type as a goal, such that a particular question utterance corresponds to a particular goal projection². For example, the literal meaning of the question “Does a fifth of Jim Beam cost more than \$5?” is a function that maps all worlds w such that $w > \$5$ to one value and all other worlds to a different value. Note that the correspondence between questions and goals is not necessarily symmetric: while every question corresponds to a goal, we do not expect that every possible goal corresponds to an easily articulated question utterance.

²This is equivalent to the partition semantics of Groenendijk and Stokhof (1984), as can be seen by considering the pre-image of such a projection: $q^{-1}(\hat{w})$.

$\mathcal{A}$: a set of possible declarative utterances, such as “It costs \$7.” or “Yes, it costs more than \$5.” We take the literal meanings of these declarative utterances $\llbracket a \rrbracket$ to be truth-conditional: a map from world states to booleans. For instance, the utterance “It costs \$7.” would be true in the world where $w = \$7$ and false otherwise.

Each of these sets is equipped with a prior, which we take to be shared in common ground. This prior in principle allows us to explicitly model the knowledge that certain world states, goals, or utterances are *a priori* more or less likely. Unless otherwise stated, we will assume uniform priors. These sets and their priors are the inputs we will vary to model different scenarios and experiments; the core computational machinery, described next, will be constant throughout the paper. Our models of questioners depend on the answerer they imagine responses to come from, hence we first introduce a family of answerer models.

A_0 : Purely informative answerer

We begin with the simplest possible answerer, A_0 , who attempts to be informative about the state of the world without considering relevance to his partner’s goals. Formally, A_0 hears a question utterance $q \in \mathcal{Q}$ and has private knowledge about the true world state $w \in \mathcal{W}$. He then uses a soft-max decision rule to choose an utterance $a \in \mathcal{A}$ proportional to the epistemic utility of that utterance to his communication partner:

$$P_{A_0}(a|q, w) \propto \exp\{U_{A_0}(a; w)\}$$

For simplicity, we omit the implicit uniform prior over answers $P(a)$, which strictly serves to specify the space of answers considered.

We define the utility of the utterance using the information-theoretic measure of surprisal: the increase in an imagined interlocutor’s I ’s certainty about the true state of the world after hearing a . This surprisal is traded off against the cost $C(a)$ of producing the answer utterance:

$$U_{A_0}(a; w) = \alpha_A \ln P_I(w|a) - \beta C(a) \quad (1)$$

where α_A and β are parameters weighting the importance of each term.

A_0 assumes that I uses Bayesian inference to update their beliefs about worlds, conditioning on the literal meaning of the declarative utterance a being true:

$$P_I(w|a) \propto \delta_{\llbracket a \rrbracket(w)} P(w)$$

where δ_e is the delta function returning 1 when e evaluates to true and 0 otherwise, and $P(w)$ is the (shared) prior over states of the world.

This formulation of A_0 is equivalent to the pragmatic speaker S_1 for declarative utterances in previous RSA models (Goodman & Frank, 2016). It thus serves as a useful baseline, more charitable than a purely random speaker: if asked about a fifth of Jim Beam and knowing that it costs \$8, A_0 would prefer “It costs \$8” to “It costs more than \$5” or (falsely) “It costs \$4” simply because the former leads to an interpreter placing more probability on the true state of the world ($w = \$8$). Furthermore, because A_0 chooses utterances using the same core utilities as the more sophisticated models we consider, it serves as a “lesioned” speaker to evaluate the necessity of the relevance component.

A_1 : Relevantly Informative Answerer

While A_0 is informative, he is not *relevant*. When asked about the weather, he is just as likely to inform his partner about what he ate for breakfast as he is to say “it’s sunny.” How do we incorporate relevance into the answerer’s utility?

Following Kao, Wu, et al. (2014), we formalize speaker relevance by means of a projection function. Instead of trying to increase the listener’s certainty about the most fine-grained true world state, a relevant speaker only cares about the listener’s certainty in a coarser space in which irrelevant features have been collapsed together. That coarser space is the image $\widehat{\mathcal{W}}$ of the goal-projection g . Crucially, a belief distribution over worlds, $P(w|\cdot)$, will project under g to a belief distribution on this coarser space³:

$$\widehat{P}^g(\hat{w}|\cdot) = \sum_{w \in \mathcal{W}} \delta_{\hat{w}=g(w)} P(w|\cdot).$$

Thus, we may modify the utility from Eq. 1 to define a relevantly informative answerer A_1 :

$$U_{A_1}(a; g, w) = \alpha_A \ln \widehat{P}^g(g(w)|a) - \beta C(a).$$

³For coarse goal-image spaces that are not discrete, it may be useful to replace the delta in this projection with a smoother kernel.

In the context of answering a question, the simplest assumption about the goal projection is that it is the literal meaning of the question:

$$P_{A_1}(a|q, w) \propto \exp\{U_{A_1}(a; q, w)\}$$

This implements a literal answerer: A_1 directly interprets a question as an (epistemic) goal, then produces an informative utterance that reduces the questioner’s uncertainty under that goal projection. We note, however, that A_1 is not otherwise constrained in their set of possible answers by the question. For example, A_1 is capable of giving non-yes/no answers to polar questions to the extent that they resolve the literal goal raised by the question⁴.

A_1 may be sufficient for simple fact-based questions, like “Who was the 16th president of the U.S.?” given a rich enough semantic parser to supply the semantics of q (Berant, Chou, Frostig, & Liang, 2013). But the context-sensitivity and indirectness displayed in everyday communication suggests that deeper social reasoning may be at play. To explore more sophisticated answerers that reason about context and their partner’s underlying goals, we must first address the problem of how a speaker chooses between questions.

Q_i : Questioners

If the utility of uttering a declarative statement is imparting information about the true state of the world, what is the utility of uttering a question? We begin by assuming that a questioner aims to *learn information relevant to a private goal*. In order to choose a question that results in useful information, the questioner reasons about how her dialogue partner would respond in different possible states of the world. She then selects among questions proportional to the extent that the answers they would elicit reduce her own uncertainty about goal-relevant information.

More formally, a questioner agent takes a goal $g \in \mathcal{G}$ as input and returns a distribution over question utterances $q \in \mathcal{Q}$ by solving a simple planning problem involving two components: (1) which answers are likely to come back after asking this question, and (2) how much would be learned from each of those answers. Critically, in order to solve either of these parts, the questioner must have in mind an imagined answerer who hears her question and responds appropriately:

thus, we have a family of questioners Q_i , each using the corresponding answerer A_i :

$$\begin{aligned} P_{Q_i}(q|g) &\propto \exp\{U_{Q_i}(q; g)\} \\ U_{Q_i}(q; g) &= \alpha_Q \mathbb{E}_{P(w)} \left[\mathbb{E}_{P_{A_i}(a|q, w)} [IG^g(a, q)] \right] \end{aligned}$$

The expectation describes a questioner who first imagines what state the answerer might know the world to be in, then imagines what answer the answerer would thusly provide. The value of this answer depends on its *information gain* $IG^g(a, q)$: the gap between the questioner’s beliefs about the g -relevant state of the world before and after hearing an answer.⁵ Drawing again upon information theory, we formalize the notion of information gain using the Kullback-Leibler divergence between the g -projected world prior and posterior after hearing a :

$$IG^g(a, q) = D_{KL}(\widehat{P}^g(w|q, a) \| \widehat{P}^g(w))$$

In order to anticipate her posterior beliefs after hearing an answer, the questioner must again make use of an imagined answerer, but now imagining herself as a listener in the future and inverting that model to infer the true world the answerer was trying to communicate:

$$P(w|q, a) \propto P_{A_i}(a|q, w)P(w).$$

The above equations specify a questioner for each answerer model. Q_0 , who plans their question by reasoning about the likely responses of A_0 , does not prefer any utterance over any other, because they will all lead to

⁴Note that having the utility only depend on informativeness under the goal projection can lead to behavior where the goal-relevant answerer is happy to say something literally false as long as it maps to a truthful resolution of the coarser goal. Sometimes this is desired (as when we round the time to the nearest hour), but other times it is unintuitive. In principle, this problem could be solved by introducing the notion of a *goal hierarchy* where the pure identity goal always has some small weight (see Exp. 2 for a formalization of this notion). In our simulations, however, we use a “hard” version of the same idea that removes literally false utterances from the possible answer set and then renormalizes the answer prior before computing the utility.

⁵As with our answerer, we omit the implicit uniform prior over questions $P(q)$. A cost term $C(q)$ would be appropriate in the questioner’s utility, but for the cases we consider in this paper, the questions of interest have the same length, so we also omit cost for simplicity.

the same distribution of (informative but not necessarily relevant) answers. Q_1 , on the other hand, expects A_1 to directly interpret her question as a goal and respond informatively and relevantly, thus allowing for some questions to be better than others. It therefore instantiates a version of the expected information gain measure of usefulness proposed by Optimal Experiment Design (OED) models (Coenen et al., 2018). Critically, in contrast to classical views (e.g. Groenendijk & Stokhof, 1984) the questioner’s behavior is not necessarily governed fully by the semantics of the question she asks, but by what she actually expects her partner to say after hearing it. They coincide when reasoning about the simple A_1 answerer, but we next turn to a more socially perspicacious answerer where they diverge.

A_2 : Goal-Sensitive Answerer

Now that we have defined the utility of asking a question, we can construct our final pragmatic answerer: A_2 . The speaker component of this answerer attempts to informatively and relevantly address the questioner’s goal, like A_1 . The listener component, however, accounts for the generative process of an asker with an underlying goal rather than taking that goal to be identical to the literal meaning of the question. That is, A_2 assumes that the question was produced by a rational agent Q_1 , who is trying to address a private goal $g \in \mathcal{G}$, and then attempts to be relevantly informative with respect to his posterior over *underlying* goals given the question $P(g|q)$:

$$U_{A_2}(a; q, w) = \alpha_A \mathbb{E}_{P(g|q)} \left[\ln \widehat{P}_I^g(w|a) \right] - \beta C(a) \quad (2)$$

Reasoning backwards from questions to goals is a simple Bayesian inversion of Q_1 using the prior on goals:

$$P(g|q) \propto P_{Q_1}(q|g) \cdot P(g)$$

Here we have defined A_2 recursively in terms of Q_1 . This recursion could give rise to an arbitrarily deep chain of nested reasoning where A_i infers a posterior over underlying goals by inverting Q_{i-1} , and so on, terminating in the base case A_1 , which uses the literal question semantics. In this paper, we consider models up to a depth of $i = 2$; beyond this depth there are no qualitative changes in the model’s behavior for the cases we examine⁶.

The framework specified above has three free parameters: separate α parameters for the questioner and answerer, and an answerer cost weight β . (In principle, this parameterization could be loosened by allowing different parameters at different levels of recursion, or tightened by requiring the questioner and answerer to use the same α .) This concludes our specification of the model space, giving a set of three answerers and three corresponding questioners (see Fig. 1). In the remainder we consider evidence that allow us to select among these models.⁷

Case studies in answerer sensitivity

In this section, we illustrate our *answerer* models through several classic goal-sensitivity effects reported by Clark (1979).⁸ These case studies center around three empirical observations about the factors determining answerer behavior in dialogue. First, and most obviously, answers should be sensitive to the surface form of the question utterance: “What time do you close tonight?” should elicit a different distribution of responses than “What is the price of a fifth of Jim Beam?”. Second, answers should be sensitive to context: the *same* question utterance elicits different responses across contexts where different goals are more or less likely. Third, and most subtly, answers should be sensitive to the cue provided by the question utterance itself toward the questioner’s underlying goals.

Sensitivity to question utterance

The minimal requirement for any answerer model is that it behaves differently in response to different questions. This can also be considered a special case of the

⁶Recent systematic model exploration by Frank, Emilsen, Peloquin, Goodman, and Potts (2017) suggests that recursive depth trades off with the soft-max optimality parameter α . Because these two parameters are not generally identifiable, we fix the maximal level of recursion and achieve the behavior of higher levels of recursion by allowing α to vary.

⁷We have implemented these models in WebPPL, a probabilistic programming language (Goodman & Stuhlmüller, electronic). Reproducible experimental materials as well as code for all reported analyses and model simulations is available at https://github.com/hawkrobe/Q_and_A

⁸These examples are implemented and runnable at forestdb.org/models/questions-answers.html

minimal requirement for a dialogue agent: that their behavior at time t is in some way dependent on their partner's at $t - 1$. Consider the problem faced by a liquor store cashier answering one of the phone calls reported by Clark (1979): they may be asked one of two questions: q_t ("What time do you close tonight?") and q_p ("What is the price of a fifth of Jim Beam?"). How do our answerer models respond to each of these questions?

To model this situation, we take a world $w \in \mathcal{W}$ to be a tuple containing the closing time and the price of Jim Beam at the store in question. Thus, $\mathcal{W}$ is the set of all possible tuples $w = (t, p)$, where t is a time in the (simplified) set $\{9\text{pm}, 10\text{pm}\}$ and p is a price in $\{\$4, \$5, \$6\}$. Similarly, we consider a simplified set of answers $\mathcal{A}$: an answerer may either state a time a_{t_i} ("We close at 9:00.") or the price of Jim Beam a_{p_i} ("A fifth costs \$5."). These utterances evaluate to true in a particular world w if and only the stated dimension has the stated level.

Our informative answerer A_0 does *not* reply differently to the two questions. In either case, he slightly prefers to inform the questioner about the price using a_{p_i} because he reasons that a literal interpreter I , after hearing such an utterance, would be left with uncertainty over only two possible worlds (where the closing time is either 9:00 or 10:00) rather than the three possible worlds consistent with an utterance stating the true time (where the price is either \$4, \$5, or \$6).

Both A_1 and A_2 , on the other hand, appropriately give different answers to the two questions, though for different reasons. A_1 gives different answers because the two questions simply have different literal meanings: q_t projects a world $w_i = (t_i, p_i)$ to its first element, $q_t(w_i) = t_i$, whereas q_p projects w_i to its second element p_i . In this example, because we take the space of goals $\mathcal{G}$ to contain these same two projections, A_2 makes essentially the same inference. He reasons it's more likely, via Bayes' Rule, that a questioner agent Q_1 who chose to ask q_t would have the goal of learning the closing time than the price, and updates his beliefs accordingly.

These answerers then choose an utterance that *relevantly* addresses the goal they inferred. a_{p_i} , for instance, would leave the questioner with perfect information about p_i but gives no information at all about t_i . Thus, it would be a preferred response to q_p and a dispreferred response to q_t . This shows the desired

sensitivity to question.

Sensitivity to context

Next, we show how our models can provide different—sometimes over- or under-informative—answers to the same explicit question, depending on context. For this illustration, we consider the results of Experiment 4 from Clark (1979). Recall that liquor merchants were more likely to give over-informative answers (specifying exact price) to the question "Does a fifth of Jim Beam cost more than \$5?" in the uninformative context ("I want to buy some bourbon") than in the five dollar context ("I've got \$5 to spend"). Which of our answerer models, if any, shows a similar sensitivity to context?

Our space of possible worlds consists of possible prices for the whiskey: $\mathcal{W} = \{\$1, \dots, \$10\}$. We consider two possible goals with equal prior probability: $g_=$$, learning the actual price of whiskey, and $g_>$, learning whether or not the agent can afford to buy the whiskey at all (i.e. whether the price is greater or less than the amount the agent has in their pocket). The set of answers $\mathcal{A}$ includes exact prices $a \in \{\$1, \dots, \$10\}$ as well as "yes" and "no." Although it does not change the qualitative effects of interest, we assume that bare "yes/no" answers are slightly less costly to produce than more informative "price" answers.

We model the context sentence as a statement which the answerer must interpret before interpreting the question; we assume that it has the effect of shifting the answerer's beliefs about likely questioner goals $P(g)$. When the context is "I'd like to buy some whiskey," we assume no shift from the goal prior: $P(g_=|c) = P(g_>|c) = 0.5$. When it is "I only have \$5 to spend," we assume a shift toward the comparison goal: $P(g_>|c) = .99$.⁹

⁹This can be seen as a case of "inquisitive content" in a declarative sentence (Ciardelli et al., 2019): the sentence is literally informative about how much money the speaker has but it has the same inquisitive effect of shifting the listener's expectations about the speaker's goals as a question does. This could be derived via further pragmatic inference or via an explicitly inquisitive semantics. While our framework is compatible with either approach (see General Discussion), this is the only example of such a "hybrid" declarative utterance that arises in the paper, so to avoid complicating our presentation, we stipulate the effect and leave the formal deriva-

Now, A_0 is unable to show sensitivity to context for the same reason that he did not show sensitivity to the question utterance in our first simulation: he does not adjust his beliefs on the basis of what he hears. Thus, he always prefers to say the exact price (e.g. “The whiskey costs \$6”) because it leaves less uncertainty about the true state of the world.

Because A_1 uses the literal meaning of the question, which is context-independent and not tied to the questioner’s goals, he does *not* give different responses in the two contexts. In either case, he addresses the goal of finding out whether the price is more than \$5, so his answer distribution is only determined by the relative costs of giving a yes/no answer and an exact answer: they provide equivalent information about whether the true price is greater than or less than \$5.

Finally, A_2 is context-dependent in the same way as Clark (1979) observed. When the question is prefaced with “I only have \$5 to spend,” which shifts the answerer’s beliefs about goals strongly toward $g_>$ (equivalent to the literal meaning of the question), then A_2 behaves like A_1 . However, when the question is prefaced with “I’d like to buy some whiskey,” which leaves in place the answerer’s uniform prior over goals, then the *exact price* answer is favored more strongly. This answer is much more informative in the case when the questioner’s goal is to know the exact price, and no worse in the other case. The expected value of the non-literal answer is thus higher, and the answerer responds proportionally.

This example demonstrates that decoupling the inferred goal from the explicit meaning of the question is necessary to provide the model freedom to reflect context, though the exact shift in goals depending on context is assumed rather than derived.

Sensitivity to relationship between goal and question utterance

The subtlest effects of answerer sensitivity go beyond the overt question utterance and context alone; given the same context, answerers may respond to one question according to its literal meaning but another using more indirect or overinformative statements. One of the key predictions of A_2 is that these effects are due to different inferences about the questioner’s *underlying goals* on the basis of the questioner’s choice of utterance itself.

For our final case study, we consider Experiment 5

from Clark (1979), which called restaurants and asked one of four yes/no questions about which *credit cards* the restaurant accepted. We focus on three¹⁰:

1. “Do you accept MasterCard?”
2. “Do you accept American Express?”
3. “Do you accept credit cards?”

Clark (1979) found that (1) and (2) were nearly always answered literally, with a ‘yes’ or ‘no’, while (3) was significantly more likely to be answered with full information (a list of all credit cards accepted). Which, if any, of our answerer models show this pattern of responses?

We formalize the scenario as follows. The set of possible worlds $\mathcal{W}$ is given by the power set $\mathcal{P}(C)$, where $C = \{\text{Visa, MasterCard, American Express, Diner’s Club, and Carte Blanche}\}$, the set of five possible credit cards. The prior distribution over worlds, $P(w)$, can in principle take into account the empirical likelihoods of restaurants accepting different cards, but we treat them as uniform to make the dynamics of our simulations clearer.

The question space $\mathcal{Q}$ contains the five basic questions (“Do you accept c ?” with $c \in C$ one of the five cards) in addition to “Do you accept credit cards?” For the literal semantics of the five basic questions, indexed by c , we use the function $q_c(w) = c \in w$, which projects worlds to a boolean corresponding to whether they accept card c or not. The latter question $q_{any}(w) = |w| > 0$ projects to a boolean corresponding to whether any cards are accepted at all.

The answer space $\mathcal{A}$ includes lists of cards (e.g. “the cards we accept are Visa and MasterCard”), which are

tion for future work.

¹⁰The fourth question was “Do you accept any kinds of credit cards?” We view Clark’s results for this question as an example of M-implicature: by using a more costly utterance with the same literal meaning as question (3), the meaning becomes marked. Rational speech act models capture M-implicature by introducing *lexical uncertainty* (Bergen, Goodman, & Levy, 2012), where both speaker and listener begin their inference with uncertainty over the literal meaning of utterances. Because this is the only example of M-implicature that arises in this paper, and this is not the critical result from the experiment, we exclude it from our discussion.

Figure 2. Illustration of how A_2 uses the questioner’s utterance as a cue to their underlying goal and responds relevantly, thus producing the sensitivity reported by Clark (1979). We set the optimality parameter $\alpha_A = \alpha_Q = 50$ and cost $\beta = 5$. The middle column is the marginalized goal posterior showing the expected number of cards; right column is expected probability of exhaustive listing (vs. literal yes/no answer) marginalized over all possible worlds.

interpreted exhaustively, and ‘yes’ and ‘no.’ We take the denotations of these polar responses to depend on the literal meaning of the previous question: $a_{\text{yes}}(w; q) = q(w)$ and $a_{\text{no}}(w; q) = \neg q(w)$.

Finally, we take the possible goals, $\mathcal{G}$, of the questioner to be finding out whether the restaurant accepts at least one of a set of cards of interest (i.e. “does the restaurant take any card in my wallet?”). This family of goals is parametrized by $C_q \in \mathcal{W}$, the set of cards that the questioner is actually interested in: $g_{C_q}(w) = |C_q \cap w| > 0$.

Because A_0 does not consider the question meaning, “yes” and “no” are not well-defined: the true list of cards is the most informative answer to any question, as it provides the exact identity of the true world. A_1 , on the other hand, gives either a literal yes/no answer or exhaustive answer in proportion to their prior probability, regardless of which question is asked. Both an-

swers provide complete information under the question projection.

By reasoning about the questioner’s underlying goals, A_2 displays the pattern that Clark (1979) observed: a higher probability of giving a literal yes/no response to “MasterCard?” and “American Express?” than to “Credit cards?”. To understand *why* A_2 behaves this way, we walk through the chain of pragmatic reasoning (Fig. 2).

The key observation lies in the evidence provided by the question about the questioner’s likely goal: $P(g|q)$. Consider two cases, beginning with the questioner’s perspective (left-most column of Fig. 2). First, suppose Q_1 is only interested in an American Express card (goal g_{AmEx}). Any answer A_1 gives to q_{AmEx} (the question “American Express?”) would fully resolve her goal. “Yes” would yield certainty in accepting it and “no” in not accepting it. On the other hand, some answers to

q_{any} (“Credit cards?”) would not be resolving: if A_1 responded “yes”, there would still be substantial uncertainty over whether the shop takes American Express in particular (for example, A_1 could informatively say “yes” if they only accepted Visa). A questioner with goal $g_{\{AmEx\}}$ is thus more likely to ask q_{AmEx} .

Now, suppose instead that Q_1 has a long list of cards (say, all five: g_{all}). A yes/no answer for any one card is not guaranteed to resolve her goal: if she asked q_{AmEx} and got a “no”, she is still uncertain whether another card on her list would be accepted. On the other hand, q_{any} is expected to be highly informative under A_1 ; if she gets a “no”, she is certain that no card on her list is accepted and if she gets a “yes”, she is certain that one is. A questioner with goal g_{all} is thus more likely to ask q_{any} .

This asymmetry in question behavior given different goals provides A_2 information about the goal given a question: he can invert this generative model to obtain a posterior $P(g|q)$ over goals (the center column of Fig. 2 summarizes this goal distribution). By Bayes’ rule, a singleton goal is more likely if the questioner asks q_{AmEx} and a longer list of cards is more likely if she asks q_{any} . The latter goal is precisely the case where an exhaustive answer is most informative (left-most column of Fig. 2), thus leading to the empirical result: answerers are more likely to give indirect, exhaustive answers when asked “Do you accept credit cards?”

Discussion

In these case studies, we examined three empirical answerer-sensitivity effects that serve as desiderata for how a dialogue agent ought to interpret and respond to questions. They should be sensitive to the question utterance, the context, and the relationship between the questioner’s goal and the utterance they produce to fulfill that goal.

A simple informative speaker A_0 was insensitive to all of these factors: it lacked a listener component that could induce any dependence on its partner’s utterance. An informative answerer A_1 who also attempted to be *relevant* with respect to the meaning of the question was sensitive to the question utterance, but it was unable to adapt its responses to different contexts or give more or less literal answers to different questions. Finally, a more sophisticated answerer A_2 , who inferred the questioner’s underlying goals and attempted to be directly

relevant to those goals, qualitatively captured all three effects.

One concern with these results is that the answer distribution depends in each case on a stipulated space of alternative questions. We chose simple, minimal sets to make the dynamics clear, but how robust are the results to the inclusion of additional alternatives? To address this concern, we ran two additional analyses. First, for our simulation evaluating sensitivity to context, we previously only included a single possible question (to isolate the effect of the context sentence). Does this context effect disappear if a question directly corresponding to the other goal were available? We re-ran our model after supplementing the set of possible questions with the alternative “How much does a fifth of Jim Beam cost?” which we assigned a literal meaning that is identical to the exact price goal, $g_{=}$. We found that the existence of this alternative does slightly reduce the answerer’s preference for responding with the exact price (from 62% to 60%, because A_2 reasons that if Q_1 really cared about the exact price, they would have asked this other question). However, this alternative has the same effect under both contexts, leaving the qualitative context effect unchanged.

A similar question arises in the credit card scenario: if the most direct question, “What cards do you accept?”, were a potential alternative, would it interfere with the answerer’s sensitivity to the questions they were asked? To explore this case, we assigned this question a literal meaning that projects the world to the exact list of cards that are accepted and re-computed A_2 ’s behavior. Again, we found that the qualitative effect of interest was robust to this change. Because Q_1 will consider this a good question regardless of their goal, failing to ask it is not particularly informative about the goal: the more informative signal, for A_2 , about the goal is provided by the *relative differences* in preferences for the other questions, which roughly puts us back in the case when we did not include this question in the set of possibilities. We conclude that, at least in the most obviously worrisome cases, our model’s qualitative pattern of behavior does not change with larger spaces of question alternatives, although the effect size may differ.

These case studies focus on answerer behavior and demonstrate how social reasoning explains prior experimental findings in a unified formal framework. How-

ever, our framework also generates novel predictions not addressed by existing experimental data. Most conspicuously, there is a relative lack of data on how people *choose questions* in similar social scenarios—when the questioner knows her partner wants to be helpful but may not know ahead of time exactly what goal she is pursuing. We therefore designed a new experimental paradigm—the *Hidden Goal* paradigm—to simultaneously probe questioner and answerer behavior under such conditions. This paradigm allows us to rigorously compare our different models, and to empirically justify our modeling assumptions.

The Hidden Goal paradigm

The functional motivations for asking (epistemic) questions lie in two sources of informational asymmetry. Agents have both (1) private goals and (2) private knowledge about the world. If the questioner already knew the relevant information about the world, there would be no need to ask about it. Similarly, if a cooperative answerer already knew the questioner’s goal, they would just provide the necessary information without needing to hear a question. Thus, studying how people ask and answer questions in dialogue requires an experimental paradigm that can induce and manipulate these two forms of uncertainty.

Prominent paradigms for studying question asking in the lab, including *Battleship* (Rothe, Lake, & Gureckis, 2016, 2018) or *20 Questions*-style games (Siegler, 1977; Cohen & Lake, 2016; Nelson, Divjak, Gudmundsdottir, Martignon, & Meder, 2014; Ruggeri, Lombrozo, Griffiths, & Xu, 2015), have been ideal for investigating how people select between different possible queries. By withholding from the questioner the ground truth location of the ships in *Battleship*, or the identity of the object that their partner is thinking of in *20 Questions*, they effectively induce an asymmetry in knowledge about the world.

However, by design, the answerer in these paradigms knows the questioner’s true goal with certainty (e.g. to “find all battleships”) and is constrained to respond truthfully but otherwise uncooperatively. This simplification—fixing a single goal and limiting the expected behavior of the answerer—is desirable for measuring question preferences over large, unbounded spaces (Cohen & Lake, 2016; Rothe et al., 2017) or probing developmental change in basic informational

search (Ruggeri et al., 2015; Ruggeri, Lombrozo, Griffiths, & Xu, 2016), but severely limits the pragmatic interplay that is characteristic of real-world dialogue.

We propose the *Hidden Goal* paradigm to address the challenge of creating the functional conditions for question and answer pragmatics in a controlled lab setting. There are two key features that distinguish this paradigm from those used in prior work. First, while preserving the asymmetry in world knowledge between questioner and answerer used in prior work, it also explicitly introduces an asymmetry in access to *goals*. We privately assign one of several possible goals to the questioner and withhold information about the identity of this goal from the answerer. Second, it establishes a *joint task* shared by both participants—the answerer is only rewarded when the questioner meets their goal—thus creating the need for cooperation in combining these different sources of private knowledge.

In the remainder of the paper, we present three behavioral experiments in this paradigm that provide a more rigorous evaluation of both questioner and answerer models in our framework. Each experiment additionally demonstrates how the core computational mechanisms of our theory can be elegantly composed with other linguistic and cognitive primitives to capture more sophisticated question and answer behaviors. In Exp. 1, we extend our model with an underlying formal account of definite reference in questions (e.g. “where is *the* dog?”) which uses graded knowledge of concept typicality. In Exp. 2, we show how the model extends to hierarchically structured spaces of goals. And finally, in Exp. 3, we show how the same framework straightforwardly extends to multi-round dialogue, where the answer to one question affects the next.

Experiment 1:

Definite reference and conceptual knowledge

As a first quantitative test of our models, we designed a cooperative guessing game that relies on knowledge of ontological relations between concepts. On each trial, four objects are hidden behind four doors. One player – the *questioner* – is assigned one of the objects to find as a private goal but does not know which object is behind which door. The other player – the *answerer* – knows each object’s true location but does not know which one the questioner is looking for (Fig. 3A). To succeed in the face of this knowledge asymmetry,

Figure 3. Task display and procedure for Experiment 1. (A) Questioner is privately assigned a goal object to find, but does not know which objects are behind which doors; answerer is privately shown the locations of objects but does not know which object the questioner is trying to find. Possible questions are shared. (B) Example of a single trial proceeding through 4 phases: goal assignment, question selection, answer selection, and location guess.

the questioner must first ask a question (i.e. *where is the...?*), then the answerer must respond helpfully so that the questioner can guess the location of their goal object more accurately (Fig. 3B).

Critically, we manipulated the sets of goals, questions, and answers across trials to create a number of different scenarios where our models make distinguishable predictions. To do so, we exploited the hierarchical conceptual structure provided by nominal referring expressions. For example, a questioner could ask about the “Dalmatian,” “dog,” “pet,” or “animal” and all could plausibly describe an image of a Dalmatian in the answerer’s view (R. Brown, 1958; Graf, Degen, Hawkins, & Goodman, 2016). However, among these labels, only the super-ordinate label “animal” could describe an image of a whale. These conceptual relationships interact with common knowledge of alternative questions and answers to license different pragmatic inferences.

Methods

Participants. 200 participants from Amazon’s Mechanical Turk were recruited and paired into teams of two. We excluded 50 participants due to a server

crash that terminated the task before completion, and excluded two additional games because participants reported that they were not native English speakers. This left 74 unique completed games (148 participants).

Stimuli. To support quantitative model comparison, we needed to elicit graded behavior under a variety of different scenarios. We thus constructed twelve items by crossing four conceptual domains (*animals*, *plants*, *places*, and *artifacts*) with three concept hierarchy structures (*nested*, *overlapping*, and *equivocal*; see Figure 4). Each item consisted of four objects and four labels, such that different labels referred to different subsets of objects at different levels of abstraction. The four objects served as possible goals and answers, while the labels served as possible questions¹¹. The different concept hierarchy structures were intended to serve as conditions driving qualitatively different model predictions, while the conceptual domains were intended to support item-level generalizability within each condition (i.e. to ensure that a particular idiosyncratic choice of images or labels is not driving the effect).

¹¹A document containing these stimuli for all items is available online at <https://osf.io/cjs85/>

Figure 4. Example of each condition of conceptual hierarchy used in our experiment. Example stimuli are shown for the *animals* domain, but the other conceptual domains were structurally identical.

Procedure. After passing a short comprehension quiz on the game instructions, participants were randomly assigned to the role of questioner or answerer and paired with another participant in a shared virtual environment using the infrastructure described in Hawkins (2015). Some aspects of each participant’s environment were denoted as private (i.e. the questioner’s goal and the answerer’s knowledge of object locations) while others were explicitly flagged as common ground (i.e. the sets of possible questions and objects). Fixing these sets and making them explicit as common ground allowed us to avoid ad hoc assumptions about these sets in our model formulation, and also reduced the memory load on participants.

We constructed the trial sequence such that each participant provided one response for each of the twelve items, in a random order. Each trial proceeded through four phases. First, the questioner was randomly assigned one of the four possible goals by a spinning wheel animation. They then used a drag-and-drop interface to select one of the four questions to send. Upon seeing this question appear, the answerer responded with a location (“The X is behind door Y”) by dragging the object they wanted to reveal into a box. Each of these drag-and-drop actions was rendered for the other participant in real time to emphasize that they were playing with a real human partner. Finally, after receiving the answer, the questioner clicked on one of the gates to make their guess. Both players received feedback about the outcome of the guess: the questioner saw the true locations of the objects, and the answerer saw the questioner’s true private goal.

Qualitative results

Before proceeding to quantitative model comparisons, we highlight several key patterns in the behavioral data where models make qualitatively different predictions. This allows us to *falsify* certain models on the basis of failing to predict these qualitative patterns (Palminteri, Wyart, & Koechlin, 2017). Because these predictions are based purely on the hierarchy structures, we treat the four domains as distinct items in our analyses but continue using examples from the “animal” domain for readability.

Answerer behavior. The most straightforward qualitative prediction distinguishing our answerer models concerns the *nested* condition. This condition was characterized by four successively more specific labels, ranging from the super-ordinate “animal” to the subordinate “Dalmatian.” At each layer of conceptual nesting, one of the objects in context was excluded: “animal” is true of all four objects, “pet” is true of three, and so on. The aggregated answerer response distribution across domains is shown for all possible questions in Fig. 5A.

Some answerer behavior is straightforwardly predicted by both A_1 and A_2 : when asked “Where is the *Dalmatian*?”, 100% of participants revealed the location of the Dalmatian. Other questions, however, lead to sharp qualitative differences between models. Suppose the questioner asks “Where is the *animal*?”. Neither A_0 nor A_1 have a preference over which of the animals to provide information about, because all four satisfy the literal meaning of the question. Human participants, on the other hand, displayed a strong preference across domains: 94% provided the location of the “whale,” the

Figure 5. Qualitative model predictions and results for (A) answerers in nested condition and (B) questioners in overlapping condition. Highlighted rows indicate clearest qualitative differences. Response counts are aggregated across the conceptual domains.

object that isn't queried by any of the other questions.

We statistically tested the non-uniformity of this distribution while accounting for domain-level variability by running a log-linear mixed-effects model with random effects for domain¹². Specifically, we used a Poisson linking function to predict the (log) count data for each response level, as implemented in lme4 (Bates, Mächler, Bolker, & Walker, 2014). Our null model contained only a fixed intercept, and random intercepts for each domain, capturing the qualitative prediction of A_0 and A_1 that all response levels share the same frequency. We compared this to an alternative model adding an additional regressor that was dummy-coded 1 for “whale” and 0 for the other levels, allowing flexibility to break the symmetry. Because these models are nested, we used a likelihood-ratio test comparison accounting for additional degrees of freedom. We found that the intercept-only model was strongly rejected in favor of the non-uniform model, $\chi^2(3) = 181, p < 0.001$.

This non-uniform pattern is only consistent with the predictions of A_2 . To understand why, consider the inference that can be made about the underlying goal from hearing the utterance “Where is the animal?”. A_2 reasons that if the questioner were looking for the Dalmatian, they would be more likely to ask “Where is the Dalmatian?”. If looking for the other dog, the poodle, they would have asked “Where is the dog?” More generally, they would have been more likely to ask a more direct question under any goal except the “whale,” and the fact that they *didn't* ask one of these more direct questions is evidence that they didn't have these goals.

¹²Because of the small number of responses per participant, we were unable to also estimate participant-level random effects. We considered controlling for domain more important, however — although there high agreement in responses across domains overall, the ‘place’ domain was a notable outlier due to a confusing choice of images (see Appendix Table 1).

Having inferred that the most likely goal is the whale, A_2 gives the most informative, relevant answer: the location of the whale.

Questioner behavior. Next, we turn to a subtler qualitative pattern in the “overlapping” condition, which was designed to distinguish between our *questioner* models. The aggregated counts of participants asking each possible question in this condition are shown in Fig. 5B. One goal in particular leads to qualitatively diverging predictions across models: suppose the questioner is assigned the “house cat” as their goal object (the row highlighted in green). Q_0 is equally likely to ask any of the four questions in this scenario, since A_0 ignores their utterance anyway. Q_1 , in turn, is evenly split between the two parents in the tree – “Pet?” and “Cat?” — it reasons that A_1 would have a 50% chance of responding with the house cat’s location in either case. Empirically, we observed that 64% of questioners asked “Cat?” while only 25% asked “Pet?”

We again used a mixed-effects regression model to account for domain-level variability when testing whether these probabilities deviated from uniform. Our simplest model only contains a fixed intercept, forcing all responses to be generated from the same parameter. Our second model, representing Q_1 ’s qualitative predictions, adds a categorical predictor coded 1 for “Cat?” and “Pet?” and -1 for the other two responses, thus forcing the former two responses to share a single parameter. Our final model added an orthogonal contrast allowing “Cat?” and “Pet?” to be different (i.e. 1 for “Cat?”, -1 for “Pet?”). All of these models contained domain-level random effects for each fixed-effect regressor. We found in a nested likelihood-ratio test that the second model was significantly preferred to the intercept-only model, $\chi^2(3) = 49, p < 0.001$, but that the third model allowing the “Cat?” and “Pet?” to differ was significantly preferred over the model forcing them to be the same, $\chi^2(4) = 26, p < 0.001$. This allows us to falsify Q_0 and Q_1 because they fail to capture this qualitative asymmetry.

Q_2 , on the other hand, successfully predicts this asymmetry (see Fig. 5B) because it takes into account the goal inferences A_2 would make after hearing her question. Once A_2 hears “Cat?” he will infer, using the same Gricean logic described above, that because the more specific “Lion?” is an available alternative for the questioner, she would have *directly* asked

about it if it were her goal. Because she didn’t, he infers that she must be interested the other cat (her true goal). “Pet?” on the other hand, gives no clue to her goals and is therefore less preferred. This critical condition thus provides qualitative evidence for Q_2 : in some cases, questioners consider the answerer’s inference about their goals when selecting their question.

To summarize, our non-pragmatic models cannot account for key qualitative phenomena while our pragmatic models (A_2 and Q_2) make the correct predictions. Next, we show that our pragmatic models also provide better overall quantitative fits to the raw behavioral data. Before describing the results of our model comparison, we formalize the task setup in our modeling framework.

Quantitative model comparison

Setup. Corresponding to the structure of our experiment design, we take the set of possible worlds $\mathcal{W}$ on a given trial to be the set of possible permutations of the four objects to the four locations (yielding $4! = 24$ possibilities). Because we explicitly instructed questioners to find the location of a particular object, we then take the space of goals $\mathcal{G}$ to be the set of goal projections

$$g_o(w) = \text{loc}(o, w)$$

where $\text{loc}(o, w)$ is a simple function looking up the location of object o in world w . Similarly, the answer space $\mathcal{A}$ is the set of constructions “The o is behind gate i .” which evaluate to true in worlds where o is at location i and false otherwise. We assume uniform prior probability over worlds, goals, questions, and answers.

Specifying the formal semantics of the question space $\mathcal{Q}$ draws upon ideas from linguistics. The denotation of a noun phrase containing a definite article is only defined if there is a unique, salient object that the noun picks out in the world; pragmatically, the use of a definite article *presupposes* the uniqueness or identifiability of some object in shared context (Lewis, 1979; Clark, Schreuder, & Buttrick, 1983; Roberts, 2003). Thus we take the meaning of each question q to be:

$$\llbracket \text{Where is the } x? \rrbracket = q_x(w) = i_x$$

where i_x is the *identifiable* object of category x . In contexts where there are multiple objects of category x , the definite article is interpreted by sampling an object from a prior identifiability distribution over objects

in the labelled category: $i_x \sim \mathcal{I}(o|x)$. This distribution may depend on distinctive visual properties, the typicality of the object for the category (Rosch, 1975), world knowledge about which labels people tend to use to refer to different objects, and other information in common ground (Clark & Marshall, 1981). To estimate the combined weight of these influences on interpretation, we estimated the empirical identifiability prior from an independent elicitation task.

Eliciting empirical identifiability priors. In prior work, there have been two approaches to setting model-internal prior distributions in a principled manner. The first is to treat these internal distributions as additional free parameters and jointly infer them from the behavioral data along with other parameters of interest (e.g. Tessler & Goodman, 2018; Tauber, Navarro, Perfors, & Steyvers, 2017). The second is to use an independent task to empirically elicit these priors from a separate population of participants, and then use this (parameter-free) distribution in the model (e.g. Kao, Wu, et al., 2014). Here, we take a hybrid approach. We elicit empirical identifiability priors from an external sample, as in the second approach, but then infer the extent to which these empirical priors are actually necessary for explaining participants’ question and answer behavior in our model comparison.

We thus designed a simple elicitation task using the stimuli from our communication task. On each trial, we presented participants with a set of four pictures and a label X and were asked to “Click *the X*.” Each participant was given one trial for each of the twelve items used above; to obtain the label X , we sampled one of the four possible category nouns that could be used in questions. In this way, no participant saw the same set of objects more than once. Of 192 participants recruited from Amazon’s Mechanical Turk, 13 participants were excluded after reporting a non-English native language, and 15 more were excluded due to self-reported confusion over the instructions. This left data for 164 participants. To form an empirical identifiability distribution $\mathcal{I}(o|x)$ that can be used in the model, we computed the proportion of participants that clicked on each of the objects o for each label x .

Before formally testing the importance of these empirical priors in our model comparison, we conducted a simple test of the null hypothesis that participants use a *uniform* prior over objects consistent with the noun

when interpreting the definite article. We performed Pearson χ^2 goodness of fit tests on the response distribution for each label: a total of 32 tests. We found that 69% of these tests rejected the null hypothesis of uniformity at a (Bonferroni-corrected) significance level of $\alpha = 0.05/32 = 0.002$ indicating that the identifiability distribution often deviates from uniform: not all members of a category are equally likely to be identified as the referent of the definite description.

Results of model comparison. For a rigorous model comparison among our different questioner and answerer models, we conducted a Bayesian data analysis (Lee & Wagenmakers, 2014). First, because utterances did not differ in cost, we remove the cost term of the model for simplicity by setting $\beta = 0$. Additionally, to test whether a model using our empirically-measured identifiability prior provides a better fit to the data than a uniform prior, we add an additional parameter $w \in [0, 1]$ which interpolates between them. In other words, we assume participants use some mixture between the uniform prior (over objects consistent with the noun) and empirical priors, but we have uncertainty over the mixture weight. We place uninformative priors over these parameters, $\alpha_A \sim \text{unif}(0, 40)$, $\alpha_Q \sim \text{unif}(0, 40)$, $w \sim \text{unif}(0, 1)$.

We estimated the marginal likelihood of our raw answerer data under A_0 , A_1 , and A_2 , and of our questioner data under Q_0 , Q_1 , and Q_2 . Additionally, for each of these models, we constructed a lesioned variant that did not use the empirical identifiability priors (i.e. $w = 0$). Our model comparison thus includes six questioner models and six answerer models. To estimate marginal likelihoods, we used annealed importance sampling (AIS; Neal, 2001) and averaged estimates across 16 chains of 10,000 MCMC steps each to obtain stable estimates for each model.

Results of our model comparison are shown in Table 1. We can immediately rule out A_0 and Q_0 , which predict a uniform distribution of responses within each condition and performed many orders of magnitude worse than the other models. Among the remaining full answerer models, we found overwhelming evidence for A_2 relative to A_1 ($\text{BF} \approx 4 \times 10^{92}$). For our full questioner models, we found substantial evidence for Q_2 relative to Q_1 ($\text{BF} \approx 4 \times 10^7$).

Next, having established support for A_2 and Q_2 , we examine their lesioned variants to test whether the in-

		uniform identifiability	empirical identifiability
answerer models	A_0	-1231.0	-1231.0
	A_1	-848.5	-767.6
	A_2	-587.9	-554.3
questioner models	Q_0	-1025.9	-1025.9
	Q_1	-344.2	-329.3
	Q_2	-342.9	-311.8

Table 1

Results of quantitative model comparison for Exp. 1; marginal likelihoods shown on log scale.

roduction of empirical identifiability distributions improved our predictions. Again using Bayes Factors, which automatically penalize our full model for its additional flexibility, we found strong evidence against the $w = 0$ model for both the answerer ($\text{BF} \approx 4 \times 10^{14}$) and questioner (3.2×10^{13}), suggesting that participants account for non-uniform knowledge about salience and conceptual typicality when interpreting the definite article.

Finally, to examine parameter posteriors and predictives for each model, we performed inference using MCMC. We used a burn-in of 4000 steps and lag of 12 to obtain 500 posterior samples for each model. We find that the posterior predictives for the best-fitting answerer and questioner models explain a substantial proportion of the item-wise variance across domain and hierarchy type: $R^2 = .92$ and $R^2 = .98$, respectively (see Fig. 6). The MAP parameters are $\alpha_A = 36, \alpha_Q = 6, w = 0.4$ for the best performing questioner model and $\alpha_A = 3.5, \alpha_Q = 21, w = 0.6$ for the best-performing answerer model.

Discussion

In Exp. 1, we presented both answerer and questioner models with the *qualitative* and *quantitative* challenge of accounting for trial-level variance in a Hidden Goal task. We found strong converging support for A_2 relative to our other answerer models and, for the first time, tested the predictions of our questioner models. While

both Q_1 and Q_2 provided good predictive fits to our question-asking data, only Q_2 could account for asymmetries in the critical *overlapping* condition and was therefore also preferred in a Bayesian model comparison. We also found that incorporating realistic identifiability priors to handle the definite article provided a substantially better fit to our data.

There are two main reasons why it was important to incorporate these priors into our model. First, failures to account for non-uniform factors like conceptual typicality could potentially invalidate the conclusions of our model comparison. For example, we expect a house cat to be a much more typical exemplar of the category ‘cat’ than a lion. Hence, participants may not be choosing the house cat in this item for pragmatic reasons; they may just be expecting the definite description to identify the more typical member of the category. If this were the case, including common knowledge about typicality should improve the predictions of A_1 and Q_1 , thus potentially ‘explaining away’ any benefit of deeper pragmatics. After accounting for this potential confound, we still found strong evidence for A_2 and Q_2 over simpler models.

Second, without the empirical priors, A_2 and Q_2 fail to capture graded responses across different domains. This is especially striking in the ‘equivocal’ condition where models using a uniform prior predict no preference in questioner behavior: if asking about the goldfish, which is both a “pet” and a “fish,” then the symmetry of the alternatives provides no reason to favor one label over the other. Instead, we found a high degree of variability across domains. For example, when questioners were trying to find a pet goldfish, 85% of participants preferred the label ‘fish’ to the label ‘pet,’ even though both are technically true. In the ‘artifacts’ domain, on the other hand, questioners had no preference between asking about the ‘seat’ or ‘metal thing’ when the objects (metal chairs) fell into both categories.

To understand how our elicited identifiability distribution helped to explain these data, consider that a Dalmatian is a far more typical and salient member of the ‘pet’ category than a goldfish (Rosch, 1975). Hence a questioner considering this knowledge as common ground would know that asking “Where is the pet?” would tend to elicit information about the Dalmatian, not the goldfish. If they were looking for information about the goldfish, they would then prefer to ask about

Figure 6. Posterior predictives for each model (excluding A_0 and Q_0) given data from Experiment 1 compared to empirical data.

the ‘fish’ category, where a goldfish is a highly typical example. In these cases, we expect that typicality strongly guides interpretation of the embedded definite article in questions.

Our results in the overlapping condition also raise a broader question about depth of recursive reasoning in RSA models. The predicted asymmetry between “cat?” and “pet?” that we used to qualitatively distinguish our Q_1 and Q_2 models is a *second-order implicature*: it only emerges for models with chains of more than two embedded agents (i.e. Q_2 but not Q_1). Previous attempts to test such second-order implicatures outside the question-answer setting have found mixed evidence. For example, Franke and Degen (2016) used a hierarchical analysis to conclude that only one participant (out of 50) was accurately described as a second-order speaker in an ad-hoc implicature task (see also Frank et al., 2017). We expect our data also represent a heterogeneous population of questioners and answerers, but it remains an open question for future work

to investigate which aspects of the task facilitated this higher incidence of second-order inferences. One hypothesis is that the use of natural images and familiar conceptual relationships removed certain cognitive demands imposed by the *ad hoc* binary properties used in prior reference games. A different possibility is that second-order reasoning is commonly needed for asking questions (but not for implicatures) and thus average language users exhibit this competence more in the question-answer setting.

Our model for Exp. 1 also served as a first demonstration of the “composability” of our computational framework with other formal proposals in linguistics and cognitive science. Our model framework derives question and answer pragmatics from choices about literal question and answer semantics, world states, and goals. This allowed us to neatly accommodate definite reference and typicality effects in questions by importing proposals about definiteness developed in linguistics and philosophy. In other words, the core compu-

tational framework itself did not need to be extended, only applied to the appropriate space of questions. In Exp. 2, we design a more targeted experiment to test our pragmatic questioner Q_2 and show that our model can similarly accommodate more structured *goals*.

Experiment 2: Structured goals

The construct of a goal has a long history in psychology (G. A. Miller, Galanter, & Pribram, 1960; Schank & Abelson, 1977; Austin & Vancouver, 1996). The simplest spaces of goals, like those used in Exp. 1, are discrete sets of projections on the world state. For example, in the Stroop task commonly used to study cognitive mechanisms of goal-directed behavior, the goal is either to (1) read the presented word or (2) name the color in which they are written (E. K. Miller & Cohen, 2001). However, for more complex behaviors, it is often useful to represent goals in *hierarchically structured* spaces (Badre, 2008; Botvinick, 2008): we often need to achieve multiple subgoals in pursuit of a larger, more abstract goal. These more structured representations have also been key to explaining third-party goal inferences in Bayesian models of theory of mind (Baker, Saxe, & Tenenbaum, 2009). For example, if we know an agent takes the same route home every day, but one day we observe them veering off of this route toward the grocery store, we simply infer that they intend to run this errand as a subgoal of their ultimate goal of getting home.

The hierarchical structure of goals has also long been recognized in linguistic theories of questions in discourse (Kuppevelt, 1995; Roberts, 1996; Büring, 2003; Groenendijk, 1999; Rojas-Esponda, 2014). For instance, if a traveler asks *Are there window seat tickets to Berlin at 7:00?* and there are no such tickets available, it would be more helpful for the travel agent to respond *the next available tickets are for 10:00* than *no* (Rojas-Esponda, 2013). This has been understood as a recognition by the agent that the specific flight and seat preference mentioned by the traveler are subgoals in a hierarchy below the broader goal of getting to Berlin.

Structured goals of this type are naturally accommodated by our framework. To show this, we designed a simple Hidden Goal collection game inspired by recent experiments in multi-task reinforcement learning (Oh, Singh, Lee, & Kohli, 2017; Andreas, Klein, & Levine, 2017). Furthermore, this task provides an additional,

stronger test of the pragmatic questioner model Q_2 .

As in Exp. 1, participants were either assigned the role of questioner or answerer, the questioner was given a private goal, and the answerer was given private information that was necessary for the questioner to fulfill this goal. Unlike Exp. 1, however, each goal was a *combination* of multiple objects that must be collected rather than a single object. More specifically, the questioner was privately assigned a hand of playing cards to find (Fig. 7A, left). The cards were scattered around a grid, and only the answerer had visibility of where the cards were located (Fig. 7A, right). The possible hands of cards the questioner might be assigned to collect were shared in common ground.

These multi-object goals induce a hierarchy where the ultimate goal of collecting the full set was decomposed into subgoals of finding each individual object in the set. A pragmatic answerer, then, ought to infer and relevantly respond not only to the questioner’s current subgoal but also their likely ultimate goal. A pragmatic questioner, in turn, ought to be sensitive to which questions best signal the ultimate goal, beyond any particular subgoal.

Methods

Participants. We recruited 62 participants from Amazon Mechanical Turk to play a communication game. We excluded 4 participants who disconnected before completion, 10 participants who failed pre-specified catch trials (see below), and 3 additional participants who self-reported confusion about the instructions. This left 45 participants in our final sample, each of whom provided responses in both the questioner role and the answerer role.

Stimuli. Stimuli were constructed from a standard deck of 52 playing cards (4 suits $\times$ 13 ranks). For each trial, we randomly sampled between 5 and 9 cards from the deck. These cards were randomly assigned to locations in the grid.

We constructed the space of goal “combos” on each trial in two different ways, leading to two conditions where our questioner models made qualitatively different predictions. In the *baseline* condition, we sampled two disjoint sets containing two cards each. In this case, both models predict that the questioner should have no preference over which card to ask about. In the *overlapping* condition, we again sampled two sets contain-

Figure 7. (A) Task displays and (B) procedure for Experiment 2.

ing two cards each but ensured that one card appears in both sets. Here, our pragmatic model predicts that the questioner will avoid asking about this card; although it is a necessary subgoal, it would provide an ambiguous signal to a pragmatic answerer about the ultimate goal.

Finally, we constructed catch trials using single-object goals. Because there is only one card that must be found to complete the hand, questioners ought to ask about this card and answerers hearing this question ought to (at least) reveal the requested card. If the questioner asked about a different card, or the answerer failed to reveal the card that was asked about, we took this as a grounds for exclusion due to inattentiveness.

Procedure. Each trial proceeded through the same four phases as Exp. 1. Questioners were first assigned a “combo” to complete out of a shared list of possible alternatives and used dropdown menus to choose a suit and rank to ask about (e.g. “Where is the king of hearts?”). Upon seeing this question, the answerer selected either one or two cards to reveal to the questioner and the questioner was then shown the locations of these cards.

If all of the cards in their combo were revealed, participants advanced to the next trial; if not, the questioner was prompted to try again. It was always possible to complete the combo in a single exchange. To incen-

tivize performance, we gave a bonus when the exact cards in the combo were revealed on their first try.

The trial sequence was constructed such that each participant saw each of the three trial types (baseline, overlapping, and catch trials) once as questioner and once as answerer for a total of six trials. We grouped trials into two blocks, one for each role, with the sequence of conditions randomized within each block. The starting role was counterbalanced across participants.

Unlike Exp. 1, which used a real-time multi-player methodology to ensure participants knew they were playing with a real partner, we decided in this experiment that the benefit of such minimal interaction was not worth the loss in control over the input given to participants. Instead, participants were paired with an artificial agent that played the opposite role. When in the questioner role, this agent selected question stimulus with maximum probability according to Q_2 . This choice ensured that we densely sampled answer responses in the cases where the answerer models made diverging predictions. In the answerer role, the artificial agent similarly responded according to A_2 , which was intended to minimize participant frustration.

Using our models directly in the task admittedly presents a potential for ‘contamination’ where participants imitate or learn from the model’s behavior. We

addressed this concern in two ways. First, the experiment is limited to only six trials, minimizing the potential for learning. Second, because we counterbalanced the order in which roles were assigned, we will empirically test whether patterns of responses differed for participants who had already seen their partner respond in a role.

Predictions and results

Model setup. Exp. 2 was designed to provide a strong *qualitative* test of the questioner models in a scenario with more structured goals. We formalize the notion of a goal hierarchy by extending the goal space $\mathcal{G}$ from a set of possible goal projections $\{g_i\}$ to a set of ordered pairs $\{g_{ij}\} = \{(g_i, g_i^{(j)})\}$ where g_i represents an ultimate goal and $g_i^{(j)}$ represents a subgoal of g_i . In the case of Exp. 2, g_i represents the goal of finding an entire combo i — the projection from the world w to the list of locations of all cards in the combo — while $g_i^{(j)}$ represents the simpler goal of finding just the single j th card in the i th combo.

The primary computational consequences of introducing a hierarchical goal space emerge from the pragmatic answerer A_2 . Upon hearing a question, A_2 must jointly infer the ordered pair $g_{ij} = (g_i, g_i^{(j)})$ containing both the questioner’s ultimate goal and their local subgoal:

$$P(g_{ij}|q) \propto P_Q(q|g_{ij})P(g_{ij})$$

If the subgoal $g_i^{(j)}$ depends only on the ultimate goal g_i , then the prior factors as $P(g_{ij}) = P(g_i) \cdot P(g_i^{(j)}|g_i)$. To compute the likelihood $P_Q(q|g_{ij})$ of a particular question q under a tuple of goals, we assume the questioner simply sums the utilities of asking q under the component goals:

$$U_Q(q; g_{ij}) = U_Q(q; g_i) + U_Q(q; g_i^{(j)})$$

That is, a question will be preferred if it is expected to be informative for both the subgoal *and* ultimate goal.

Given this posterior over the goal tuple g_{ij} , we can generalize Equation 2 by assuming the answerer also has a combined utility. That is, she attempts to be informative with respect to each goal in the tuple:

$$U_{A_2}(a; q, w) = \alpha_A \mathbb{E}_{P(g_{ij}|q)} \left[\sum_{g \in \{g_i, g_i^{(j)}\}} \ln \widehat{P}_I^g(w|a) \right] - \beta C(a)$$

Information relevant to the ultimate goal would already provide information relevant to the subgoal; adding the subgoal explicitly in this utility increases the desire of the answerer to address the current subgoal compared to others (that are still part of the ultimate goal). While this presentation is written in terms of ordered pairs, it straightforwardly extends to deeper goal hierarchies represented as n -tuples. The simpler goal spaces considered earlier in the paper are special cases of this more general formulation, where the goal tuples are singletons ($n = 1$).

Given this generalization of the model to more structured goal spaces, we can derive predictions for Exp. 2. The world space $\mathcal{W}$ consisted of all assignments of cards to locations. Questions in the question space $\mathcal{Q}$ were of the same form as in Exp. 1 (“Where is the X?”), but because X was always constrained to be a unique card, issues of multiple identifiability do not arise. Finally, the answer space $\mathcal{A}$ consisted of any combination of one or two card locations, which evaluate to true in worlds where the card or cards are in fact at the location that was mentioned, and false otherwise.

Answerer results. Although the different conditions of the experiment were primarily designed as a stronger qualitative test of the questioner model, it also provides another opportunity to test the answerer models. To begin, consider the simple catch trials, where each goal contained only one card. Here, both A_1 and A_2 make the same prediction: that participants will reveal the single card that was asked about with high probability and, with lower probability, they may be “overinformative” by revealing both the card that was asked about along with some other card. Indeed, we found that 83% of participants responded literally with the exact card asked about, while the remaining 17% gave overinformative responses (Fig. 8A). The precise probability the models put on such overinformative responses depends on the optimality parameter α and a cost parameter penalizing longer answers; at a cost penalty of 1, the empirical response distribution is obtained with moderate values of $\alpha = 3$ for A_2 and $\alpha = 3.5$ for A_1 .

In the *baseline* and *overlapping* conditions, however, A_1 and A_2 make diverging predictions. A_1 treats these conditions the same as the catch trial, since it attempts to be directly informative with respect to the goal inherent in the question, and all questions are about a

Figure 8. Results for Exp. 2. (A) Answerers are more likely to go beyond the literal question and reveal additional cards in the baseline and overlap conditions. (B) Questioners are more likely to ask about the non-overlapping card relative to baseline. Dotted line is chance, and error bars are bootstrapped 95% CIs.

single card. A_2 , on the other hand, goes beyond the literal question to infer that the questioner asked about just one subgoal of one of the ultimate goals. When asked about a card that appears in one goal combo but not the other, A_2 infers with high probability which of the two ultimate goals is most likely the hidden goal. It is thus much more likely to respond with *both cards* than about only the single card that was asked about (or any other combination of cards). We found that only 23% of participants gave literal answers in each of these conditions, which is significantly different than the proportion (83%) in catch trials, $\chi^2(1) = 35, p < 0.001$, Fig. 8A. The empirical response distributions for these conditions are thus consistent with the predictions of A_2 and strongly falsify A_1 .

Questioner results. Next, we turn to questioner behavior, the primary focus of Exp. 2. The Q_1 and Q_2 models make qualitatively different predictions about which questions participants chose in the *overlap* and *baseline* conditions. Q_1 predicts no difference in behavior across these two conditions. Because it reasons that A_1 will always provide a literal or (indiscriminately) overinformative response, as we saw in the previous section, then asking about either card is equally useful for completing the target combo.

The pragmatic questioner Q_2 , on the other hand, recognizes that certain questions may provide better sig-

nals about the ultimate goal than others. In particular, for the *overlap* condition, Q_2 reasons that if they asked about the card appearing in both combos, A_2 would not be able to infer which of the two combos is their ultimate goal. If they instead asked about the unique card, however, A_2 would infer their true goal with high probability and subsequently be more likely to reveal the needed information. Thus, Q_2 predicts that participants will prefer to ask about the unique card. The *baseline* condition serves as a control where both cards are unique, hence both questioner models predict no preference over the two cards.

We found that 77% of participants asked about the unique card in the overlap condition, which is significantly different from the null probability of 50% predicted by Q_1 , $\chi^2(1) = 11.3, p < 0.001$. Meanwhile, the proportion of participants (48%) who asked about the card in the same position in baseline trials did not differ from chance: the Bayes Factor for binomial proportions (Morey, Romeijn, & Rouder, 2016; Morey, Rouder, & Jamil, 2015) shows weak evidence in favor of the null, $BF = 0.37$, Fig. 8B. This qualitative pattern of responses is consistent with the predictions of Q_2 but cannot be explained by Q_1 .

It was possible, however, that this pattern was driven by participants who had already seen their partner ask about the unique card on a previous trial and were copy-

ing their behavior. To address this concern, we examined these response patterns for the subset of participants who were assigned the role of the questioner on the first half of trials and thus had no opportunity to learn from their partner: 18 of these 21 participants (87%) asked about the unique card, which remained significantly different from chance: $\chi^2(1) = 9.3, p = 0.002$.

Discussion

Exp. 2 was motivated by two primary aims: to provide a stronger evaluation for the questioner models, and to demonstrate how our core framework generalizes to more structured spaces of goals. To meet these challenges, we designed a Hidden Goal task where questioners were privately assigned a set of *multiple* objects to find, inducing a hierarchy of subgoals. By manipulating whether objects overlapped or differed across sets in the space of possible goals, we constructed scenarios where our different models made sharply different predictions. We found that answerers reasoned beyond the literal subgoal prompted by the question to respond helpfully to the inferred ultimate goal in the hierarchy. Questioners, in turn, systematically chose to ask about the objects that would most clearly signal their ultimate goal to their partner, even though these different questions would have equal expected information gain under OED models.

The minimal goal hierarchy studied in Exp. 2 is a key first step toward formally addressing more structured spaces of goals in cognitive models of question asking and answering. However, handling more complex goal structures will require composing our core computational framework with additional mechanisms. In particular, deciding how to ask and answer questions in the presence of sequential dependences in the goal space likely requires a *planning* mechanism over a persistent state. While introducing such a planning mechanism into our model is outside the scope of this paper¹³, our final experiment aims to set a foundation for future work by (1) introducing a notion of persistent state into the model and (2) investigating several pragmatic phenomena that emerge from iterating our model forward in time in a step-wise manner.

Experiment 3: Extended dialogue

In our computational framework, an answerer updates their beliefs about the questioner’s goals after hearing a question, and a questioner updates their beliefs about the world state after hearing an answer. In principle, this step can be repeated multiple times until the questioner’s goal is resolved, allowing for an extended sequence of questions and answers: the questioner’s posterior over worlds after hearing an answer to their first question becomes the prior for selecting their second, and so on.

In Exp. 3, we developed a final task in the Hidden Goal paradigm to investigate how the pragmatics of asking and answering questions unfolded over multiple exchanges. This task was designed with several desiderata in mind. First, we needed a larger and more natural space of goals that wouldn’t always be possible to satisfy in a single round. This would allow us to test our models’ predictions about how the answer to one question may influence the choice of a subsequent question, or how the answerer’s beliefs about the questioner’s goal may shift after hearing additional questions. Second, although A_2 ’s internal goal inference mechanism was implicit in the success of our model predictions in Exps. 1 and 2, we sought explicit validation of this model component by directly eliciting answerers’ beliefs about the questioner’s goal.

The *Don’t Explode* task is a cooperative communication game inspired by Battleship and Minesweeper (see Fig. 9). Participants are presented with a shared 3×3 grid and told that some tiles are safe but others contain bombs. One player—the questioner—is privately assigned a particular goal pattern they must complete by revealing tiles. The other player—the answerer—is given a map showing where the bombs are. As in our previous experiments in this paradigm, the answerer has uncertainty about the questioner’s goal, and the questioner has uncertainty about the map. To succeed, then, the questioner must ask questions (e.g. *Is tile X safe?*) and the answerer choose what information to reveal in response.

¹³OED policies based on global planning, compared to locally ‘greedy’ policies, are not yet well-understood even in the simpler non-pragmatic case of sequentially querying an oracle, though see Nelson, Meder, and Jones (2018) for recent theoretical progress.

Figure 9. (A-B) Task design and procedure for Experiment 3. Answerers are shown a private map with bomb locations, while questioners are assigned a private goal. Both participants see a common board over which extended dialogues can take place. (C) Examples of initial board states in different conditions.

Methods

Participants. We recruited 124 participants from Amazon Mechanical Turk, 21 of whom were excluded for disconnecting before the end of the task. This left 104 participants in our final sample, yielding a dataset of 1774 question trials and 1801 answer trials. All participants were self-reported native English speakers.

Stimuli. Each trial was defined by three stimulus choices: (1) a choice of goal, (2) an underlying ‘map’ assigning bombs to locations, and (3) an ‘initial state’ indicating which tiles were shown as safe to both players at the outset of the trial. By manipulating the state of the initial grid, we generated a large number of items where our models make different predictions.

Goals were sampled from a space of two ultimate objectives: complete any row or complete any column (see Fig. 9A). Because these goals are logical disjunctions, they could be satisfied in as many ways as there are rows or columns, inducing a hierarchical goal space

with each row/column as a possible subgoal the questioner could pursue. This goal space was intended to be intuitive for participants, given its resemblance to the win conditions of well-known games like Tic-Tac-Toe or Connect Four.

Underlying bomb maps were sampled from the space of all assignments of bombs to locations, subject to the constraint that it was possible to win given either goal. We included this constraint so that the arrangement of bombs did not give the answerer any cues to the questioner’s goal. Finally, the initial state of the shared board was chosen from four different conditions to create a variety of different scenarios (see Fig. 9C for examples). Cells in the shared board were labeled with coordinates (A1, A2, . . .) for easy reference.

- In the *empty* condition, no tiles were included in the initial state.
- In the *pragmatic* condition, a single tile was included in the initial state. This tile was chosen

from a completely safe row or column, depending on the goal.

- The *blocked* condition was similar to the pragmatic condition, except a bomb was placed in the row or column partially completed by the initially revealed cell.
- The *random* condition randomly sampled a subset of safe tiles from the map, under the constraint that they did not fully complete a row or column.

Procedure. The task procedure was adapted from that of Exps. 1 and 2 to support extended exchanges. First, questioners were shown their assigned goal and selected a cell label from a dropdown menu to ask a question of the form “Is X safe?”. Answerers selected a yes/no answer to this question and were presented with the option to either send more information or to continue with only this yes/no answer. If they chose to send more information, we showed dropdown menus with all possible cells, allowing for responses like “no, X is not safe, but Y is.” Finally, this answer was shown to the questioner. If the new information fulfilled the goal, participants advanced to the next trial. If the goal was not fulfilled, the questioner was prompted to ask another question, and the cycle repeated. Participants received a bonus on each trial related to the number of exchanges taken before completing the goal: one cent was subtracted from a total of \$0.05 for each additional question.

We elicited the answerer’s beliefs about the hidden goal after they saw the questioner’s message but before they made their response. Three response options (“rows,” “columns,” or “not sure”) were provided to the prompt “Which goal do you think your partner has?” No feedback was provided for this response. For the other responses, answerers were prevented from sending blatantly false information (e.g. saying there was a bomb when there was no bomb), and questioners were prevented from asking about tiles that had already been revealed.

Participants alternated between the roles of questioner and answerer for a sequence of 20 trials. To acquaint participants how the task worked, all trial sequences began with two practice trials, followed by four random trials. Practice trials were constructed by including two of the three cells of a goal-relevant column

or row in the initial state. It should have been clear from a basic understanding of the instructions that the questioner needed to ask about the third cell and the answerer ought to reveal it. After these initial six trials, we constructed a scrambled sequence where each condition appeared once in each role and these critical trials were interspersed with more random trials. As in Exp. 2, participants were paired with an artificial agent playing the opposite role. This agent acted according to our Q_2 and A_2 models, with the optimality parameter set high ($\alpha = 100$).

Results

Model setup. The key innovation required to account for extended sequences is an explicit representation s_t of the *discourse state* at a given time. In general, the discourse state could contain prior utterances and any non-linguistic observations about the agent’s partner or context. In this experiment all information was updated and maintained visually, so it is sufficient to take as s_t the literal state of the game board shared in common ground at time t .

Each state is then associated with a conditional belief distribution over the world, $P(w|s_t)$. In our experiment, computing this distribution simply entailed removing any world states from the prior that are incompatible with the current revealed state and re-normalizing. Now, when a questioner is selecting a question in state s_t , rather than taking expected information gain relative to the initial prior $P(w)$, they use their current beliefs $P(w|s_t)$ as the baseline. Moreover, since the discourse state s_t is in common ground, we assume this alternative prior is used throughout internal levels of recursion, e.g. that A_1 also attempts to be informative with respect to this prior.

Beyond introducing state to handle extended exchanges, the task details for this experiment were formalized similarly to the previous experiments. The world prior was uniform across all 2^9 assignments of bombs to tiles. The question space contained the 9 options available to participants in the experiment, where “Is A1 safe?” is represented as a projection from the full world to the status of the queried tile. The answer space contained both the literal answers “yes, it is safe” and “no, it is not safe,” as well as longer answers that additionally mentioned the state of another tile, e.g. “yes, it is safe but A3 is not.”

Figure 10. Goal inferences shift with additional questions. Curves are the proportions of empirical answerer responses to “Which goal do you think your partner has?” after hearing the k th question. Error ribbons are bootstrapped 95% CIs.

Goals were again represented hierarchically, with a joint representation containing rows vs. columns at the super-ordinate level and *which* row or column at the subordinate level. The super-ordinate row (column) goals projected worlds to arrays of three booleans representing which rows (columns) contained bombs and which did not. The subgoal representing a particular row (column) projected worlds to a single boolean representing whether that particular row (column) contained bombs.

Hidden goal inferences. A primary internal mechanism driving pragmatic behavior in our models is A_2 ’s Bayesian inference about the questioner’s goal conditioned on the question. While the overall success of A_2 in predicting answerer behavior in Exps. 1 and 2 provided indirect evidence for this mechanism, the elicitation step included in our design provides the opportunity for a more direct test of the answerer’s beliefs about the goal. Do participants’ inferences about the questioner’s true goal match those internally predicted by the model?

We examined participants’ goal inferences in the *empty* condition, which was designed to measure the changes in answerer’s beliefs about the hidden goal over repeated exchanges. A_2 predicts that the first question should provide no information about the hidden goal: Q_1 is equally likely to ask about any cell under either goal. In the context of the information revealed in the first exchange, however, the A_2 model predicts

that the second question ought to license stronger inferences. We compared these shifts in the empty condition against those in the *pragmatic* condition, which served as a baseline. In the pragmatic condition, the A_2 model predicts that sufficient context is already available from the initial state to make a strong goal inference.

We aggregated data across the two possible goals by classifying responses as ‘correct’ if they matched the true goal and ‘incorrect’ if they differed. We found that after the first question in the “empty” condition, participants were roughly at chance between the response options, only guessing the true goal with 39% accuracy¹⁴. After the second question, however, they were able to infer the true goal with 82% accuracy (see Fig. 10). The distribution over the response variables changed significantly from the first to second question, $\chi^2(2) = 31, p < 0.001$.

Furthermore, answerers’ goal inferences after the second question in the empty condition were close to those observed after the *first* question in the baseline pragmatic condition (82%), where we initialized participants with similar information in common ground. These results provide a direct validation of the internal mechanism of goal inference. Answerers systematically use context to make accurate inferences about their partner’s goals, whether that context is provided in the initial common ground or by earlier question-answer exchanges.

Questioner behavior across multiple exchanges.

Just as the answerer’s beliefs about the hidden goal shifted with each question, the questioner’s beliefs about the world ought to shift with each answer. In particular, this shifting belief state provides a mechanism for the answer to one question to affect the selection of a follow-up question. The *blocked* condition was designed precisely to empirically elicit such sequences, providing a key test of our model’s ability to handle the updating state. In this condition, questioners are faced with a setback: there was a hidden bomb in the row or column that would provide the shortest path to goal completion. Answerers in this condition therefore needed to decide how to most helpfully reveal this setback, while questioners needed to decide how to use this initial answer to proceed in future exchanges.

¹⁴Apparently some proportion of participants simply guessed instead of using the “not sure” option, which received 31% of responses.

Figure 11. Decision tree of questions and answers in ‘blocked’ condition as dialogue unfolds. Response probabilities are empirical estimates; line width is proportional to response probability.

To illustrate our findings, Fig. 11 shows the distribution of responses in an example blocked trial. The initial common ground revealed that the top-left tile was safe, and the questioner’s private goal was to complete any row. All questioners began by asking about either A2 or A3, the other tiles in the top row. If they asked about A2, the most frequent empirical answer (62%) was to confirm that A2 was safe, and then give additional information that A3 was *not* safe. If they asked about A3, the most frequent empirical answer (52%) was to confirm that A3 contained a bomb, and then give additional information about a safe tile in a row that was not blocked (e.g. C1 or C2).

In both cases, answerers went beyond the literal question to provide additional goal-relevant information. When it was the questioner’s turn to ask a second question, they were thus faced with a much different epistemic state. This updated state suggested pursuing a different subgoal (i.e. a different row) than they were pursuing at the outset. *Which* subgoal they pursued on their second question differed depended both on their

answerers	A_1	-1890
	A_2	-1380
questioners	Q_1	-1136
	Q_2	-1135

Table 2

Results of quantitative model comparison for Exp. 3; maximum likelihoods shown on log scale.

first question and the answerer’s response. The Q_2 and A_2 models successfully capture these dependencies between successive exchanges by updating their representations of the state.

Quantitative model comparison

While the primary aim of Exp. 3 was to demonstrate how our model framework can be straightforwardly extended to repeated question-answer exchanges, these

data also provided an opportunity to conduct an additional model comparison over a wide variety of items.

All models have exactly three parameters: a cost weight β and optimality parameters for the questioner and answerer, $\{\alpha_Q, \alpha_A\}$. Because evaluating marginal likelihoods over this parameter space was computationally intractable, we use the maximum likelihood as a point estimate of the marginal likelihood. This is equivalent to an AIC or BIC model comparison approach because the number of parameters is matched across the candidate models. We thus evaluated the maximum likelihood of our data, excluding the first six practice trials of each game, under the following parameter ranges:

$$\begin{aligned}\alpha_Q &\sim \text{unif}(0, 59) \\ \alpha_A &\sim \text{unif}(0, 19) \\ \beta &\sim \text{unif}(0, 3)\end{aligned}$$

Because Exp. 3 was not designed to distinguish our questioner models (unlike Exp. 1-2), we expected that Q_1 and Q_2 would make similar predictions for most trials. Indeed, the difference in the maximum log-likelihood for the two models was only $d = 0.68$, indicating a rough parity in fit. These predictions, however, provided a good overall fit to human responses (see Fig. 12).

For answerers, on the other hand, we found that the pragmatic model A_2 was able to fit the answerer data substantially better than A_1 (see Table 2). A_2 was maximized at $\alpha_Q = 59; \alpha_A = 2.7, \beta = 0.9$ while A_1 attained its maximum fit at $\alpha_Q = 1, \alpha_A = 11, \beta = 0.1$. The posterior predictives for A_2 are shown in Fig. 12, also indicating an excellent overall fit to human responses.

Discussion

In Exp. 3, we considered sequences of questions and answers that extended beyond a single exchange. We equipped our models with a persistent representation of the state and tested their predictions in a multi-round Hidden Goal task where the content of earlier exchanges may influence later exchanges. First, we validated the goal inference mechanism in the A_2 model by directly eliciting the answerer’s beliefs about the questioner’s goal. The A_2 model successfully accounted for the dynamics of how these beliefs shift over successive questions, and was also substantially preferred over the A_1 model in a quantitative model comparison. Next,

Figure 12. Posterior predictives for each model, binned to intervals of length 1/10.

we examined cases where the answer to one question affects the questioner’s subsequent choice of question. Given a persistent representation of state, both the Q_1 and Q_2 models accurately predicted the dynamics of questions asked over the sequential exchange.

These results show that extending our framework with a state representation is sufficient to explain some key features of sequential exchanges. However, we expect that *planning* will be necessary for integrating the pragmatics of question asking and answering into longer dialogue sequences, as it has been for other domains (e.g. Bramley et al., 2015). Our proposed extension maintaining a (belief) state could be used to define a Markov Decision Process where different question or answer selections lead to anticipated transitions in beliefs about goals or the world, respectively. A related approach was recently taken by Khani, Goodman, and Liang (2018), who used a forward-looking recursive utility to define an action policy in a sequential language game. In this task, pragmatic agents took turns saying one- or two-word utterances, informing their partner about the world state, but were not permitted to ask questions. An important step for future work is thus to work out an appropriate planning mechanism for how question utterances are expected to shift beliefs about goals.

Another challenge for future work concerns the structure of goals. The hierarchical ‘row’ and ‘column’ goals in Exp. 3 were more complex than the multi-object goals in Exp. 2¹⁵. However, both of these

¹⁵In fact, while the goals in Exp. 2 are technically hierar-

experiments satisfied the assumption that subgoals are *conditionally independent* given their parent in the hierarchy. This assumption is clearly violated in certain real-world cases. For instance, sequential dependencies among subgoals are introduced when a process unfolds in time. Certain questions may only be useful once earlier subgoals have been satisfied, leading to more or less optimal sequences of questions. Conversely, answerers may take advantage of schematic knowledge of this sequential structure to anticipate and address later subgoals (Jackendoff, 2007).

General discussion

Asking questions is one of our most efficient and reliable means of learning about the world. Yet we do not often pose these questions to an impartial oracle; we ask other agents, in dialogue. In this paper, we jointly consider the decision problems faced by socially aware questioner and answerer agents. Questioners must plan over the possible answers they are likely to receive from their partner and seek to maximize the future reduction in their own uncertainty. Answerers, in turn, seek to informatively and relevantly address the underlying goals most likely to have motivated the question.

We provided several lines of evidence supporting this theory—formalized in our A_2 and Q_2 models—over asocial questioners and answerers. In a series of answerer simulations, we showed that only A_2 could qualitatively account for classic answerer-sensitivity effects from Clark (1979): answers depend critically upon the question utterance, the context in which a question is asked, and the relationship between the questioner’s underlying goal and their question. We then found strong qualitative and quantitative evidence for both A_2 and Q_2 from three dialogue experiments in our *Hidden Goal* paradigm, which further demonstrated the extensibility of our core framework.

A central theoretical advance of the proposed framework is the synthesis of current (OED) approaches to optimal question asking, in active learning, with the generic social reasoning mechanisms proposed by the Rational Speech Act approach to language use. In addition to embedding the former approaches into a social context, this work moves RSA models beyond the production or interpretation of single utterances (in context) to consider the dynamics of simple dialogs where both listening and speaking—linguistic input and

output—must be integrated within the same agent. In the remainder of the discussion we address several broader issues raised by our framework and note a number of avenues for future work.

What is the meaning of a question?

While our framework draws extensively from previous theoretical accounts of question and answer semantics, the division of labor between semantics and pragmatics in our proposed model leads to a subtly different perspective on the meaning of a question. These differences stem from two psychologically meaningful distinctions at the core of our model. The first disentangles the literal meaning of a question from any particular set of possible answers and instead associates it with a specific goal that it raises or signals for an answerer. The second disentangles the space of possible latent goals in the questioner’s mind from the one associated with the interrogative signal they choose to provide. Together, these distinctions allow us to *derive* context-sensitivity at the pragmatic level via social reasoning that grounds out in context-independent meanings.

This idea differs somewhat from the context-sensitive literal meanings suggested by Van Rooy (2003). As previously described, Van Rooy argued that the meaning of a question ought to be underspecified and resolved only through the contextual parameter of the questioner’s decision problem. Once we explicitly address the issue of how the answerer actually *infers* this decision problem from hearing the question (and from other context), it becomes unnecessary to additionally posit an underspecified literal question meaning.¹⁶ Because we take a goal to be the same type

chical, a model ignoring the subgoal and only optimizing for ultimate goals would be sufficient to fit the data. This is not the case in Exp. 3, where reasoning about the current subgoal is integral for accurate predictions. We introduced goal hierarchies into the model at an earlier juncture for expositional clarity and consistency with their treatment in Exp. 3.

¹⁶RSA models have handled vague or underspecified literal meanings by introducing a free variable into the meaning that is determined by the speaker and inferred by the pragmatic listener (e.g. Lassiter & Goodman, 2015; Bergen et al., 2016). An implementation of van Rooy’s proposal would require the free variable in the literal meaning to be derived from the questioner’s private goal itself. It is unclear how such a free variable could be inferred without circularity, or whether such a mechanism could lead to appropriate interac-

of object as a classic partition-based question meaning (Groenendijk & Stokhof, 1984), we can implement its meaning in a literal answerer who identifies the goal directly from the question.

In this way, the meaning of a question in our framework resonates with dynamic accounts of question meaning (e.g. Groenendijk, 1999), as well as recent proposals for an inquisitive semantics (Ciardelli et al., 2019), where the formal content of a question is simply the “issue” it raises into the context. Our proposal that the meaning of an interrogative sentence shifts the hearer’s beliefs about the questioner’s goals can be seen as a probabilistic analogue to the notion of *inquisitive content* as implemented in formal logic. In both cases, defining literal question semantics directly in terms of their inquisitive content rather than possible answers imposes a crucial *modularity* at the semantic level of the model. This modularity frees higher-level pragmatic considerations to handle the problem of how questioners actually think their partner will generate answers given a question, which may or may not correspond to the literal resolution conditions.

The connection to inquisitive semantics is an intriguing area for further exploration. For example, a fundamental insight of inquisitive semantics is that declarative sentences sometimes *also* have inquisitive content. While we have made the simplifying assumption that declarative utterances are purely informative, the RSA framework could also support this “hybrid” notion. For example, previous RSA models have explained the use of non-literal language like hyperbole, irony, and metaphor in declarative sentences by the same mechanism we invoked here: providing information about the speaker’s goal (Kao, Bergen, & Goodman, 2014; Kao & Goodman, 2015; Kao, Wu, et al., 2014). A subtler difference, however, lies in the correspondence between the projection function we have used to define goals and the “issue” formalism in inquisitive semantics, which is not as strict. Thus, the class of functions used to represent a goal in our framework may need to be expanded to properly handle embedded questions, entailment, and disjunction (Ciardelli et al., 2019).

The correspondence between questions and goals

A straightforward consequence of defining the meaning of a question to be the same kind of object as a goal, is that our Q_2 model predicts that people will ask

the question most directly corresponding to their private goal if it is accessible and easy to produce. After all, it would be the clearest signal. So why would a questioner ever ask a question that didn’t directly correspond to their underlying goal? Our dissociation of question meanings and latent goals raises the interesting possibility that the set of questions that is easy to verbally express is effectively a *subset* of the possible latent goals one might have: there may not exist a question utterance literally encoding every underlying goal an agent may need to address, thus the need for pragmatics.

More specifically, we were motivated by real-world cases where the interrogative sentence with a meaning more closely corresponding to the underlying goal would be longer or more confusing to produce. For instance, when trying to meet up in a city, people regularly get away with asking the bare question “Where are you?” when they could directly signal the relevant coarse-graining, e.g. “At the intersection of which two streets are you currently located, and which direction are you heading?” (Potts, 2012). While we leave a careful formal account of “Where are you?” in our framework for future work, it seems natural for the literal meaning to be the finest-grained goal of locating the hearer’s exact point in space, such that higher levels of pragmatic recursion coarsen it to the appropriate level. Pragmatic questioners will prefer the less costly utterance if they can expect the answerer to recover the intended coarser goal. Issues of cost may enter *answers* to “Where are you?” as well, since giving a precise GPS coordinate is presumably costlier than more coarse-grained and easily retrievable spatial referring expressions.

Empirically, it was challenging to naturally capture such cost differences for our controlled experiments; instead, we imposed this cost artificially by disallowing the most direct questions for some scenarios. Under this constraint, we were able to test the model’s preferences among alternative questions. This does not invalidate our model comparison, but it does leave open for future work the empirical challenge of designing a lab task that can naturally modulate the cost of the direct question. Other empirical paradigms for question-asking also tend to impose strict constraints on allowable questions for similar reasons. In games like Battle-

ship or 20 Questions, the truly optimal questions would be “Where are all of your ships?” or “Which object are you thinking of?”

These considerations may also shed new theoretical light on recent results by Rothe et al. (2018) finding that in games like Battleship, people often fail to ask maximally informative natural language questions (although they *can* accurately rank the informativity of pre-generated questions). By embedding questioning in a social, communicative context (see also Shafto, Goodman, & Frank, 2012; Shafto, Goodman, & Griffiths, 2014), our framework suggests that this behavior may in fact be resource-rational under the conditions where most questioning takes place (Griffiths, Lieder, & Goodman, 2015). If the maximally informative questions are longer and more syntactically or conceptually complex (e.g. “At what location is the top left part of the purple ship?”), they may also be more cognitively demanding to generate. When assuming a cooperative pragmatic answerer, however, a less costly question may be just as effective for signaling the goal and getting a helpful answer.

Computational challenges

How do our assumptions scale up to less constrained interactions? We suspect that the simple, discrete sets of goals used in our models above may need to be generalized to continuous spaces and broader mixtures of discourse topics (Blei, Ng, & Jordan, 2003; Griffiths, Steyvers, & Tenenbaum, 2007). The computational challenges associated with such complex goal spaces may contribute to the difficulties young children face in answering broad, sentence-focus questions like “what’s happening?” (Salomo, Lieven, & Tomasello, 2013).

As a broader point, our framework offers a rational analysis of the social inferences involved in asking and answering questions, which only adds to the many important algorithmic questions already raised by rational models of inquiry Coenen et al. (2018). For instance, we could use without modification the sophisticated and open-ended grammar of questions developed by Rothe et al. (2017) as the prior for the question set Q in our pragmatic framework, but this would turn the literal questioner agent Q_1 into an infinite distribution, leading to very difficult inference at higher levels of recursion. Similarly, taking an analytic expectation over all possible worlds quickly becomes intractable as the size of the

world space W grows. We used exact enumeration for all of our results for theoretical precision, but it is clear that an algorithmic account must address how these computations can be effectively approximated, perhaps via Monte Carlo integration or approximation by neural networks.

Applications to other question phenomena

While the applications throughout this paper have primarily focused on canonical *epistemic* questions in a cooperative context, our framework also provides a tool set for formalizing pragmatic accounts of other question phenomena. We briefly highlight four potential applications: rhetorical questions, pedagogical questions, politeness, and uncooperative interrogation.

First, rhetorical questions are characterized by being either completely obvious or completely unanswerable (*What’s going to happen to these kids when they grow up?*). A pragmatic answerer hearing such a question in our framework would reason that a questioner is unlikely to have chosen it with the epistemic goal of getting information and instead must have another goal in mind such as signaling common ground or expressing affect (Kao, Wu, et al., 2014). In this way, one could derive the assertive force of rhetorical questions from pragmatic principles (see Rhode, 2006).

Second, consider how we could account for pedagogical questions asked by parents or educators (Moradlou & Ginzburg, 2014; Gall, 1970). A math instructor asking a question like “what is $2 + 2$?” already knows the answer, so what utility could there be in asking? Instead of trying to reduce their own epistemic uncertainty over the result of the math problem, perhaps the teacher is instead trying to reduce their uncertainty over the child’s knowledge state. If we introduced such uncertainty about the answerer’s background knowledge in our model and gave the questioner the goal of determining this background knowledge, it would be rational to ask targeted questions that would expose the answerer’s knowledge.

Third, we return to indirect speech acts like “Do you know where I can find a restroom?” These indirect forms of questions are often analyzed as being motivated by *politeness*, since they pinpoint likely obstacles to answering (Francik & Clark, 1985) and give answerers an easy way to opt out (P. Brown & Levinson, 1987). The same politeness considerations may influence an-

swerer behavior as well. For instance, when people do give literal answers they may use more words than necessary to avoid seeming curt or unhelpful (see also Chaudhry & Loewenstein, 2019). Such politeness considerations have been formally incorporated into RSA models as a kind of goal that takes into account social and presentational utilities (Yoon, Tessler, Goodman, & Frank, 2016), which could be straightforwardly incorporated into a questioner model.

Finally, although we have assumed basic prosociality and transparency in our agents, some of the most intriguing question-answering scenarios flout this expectation. In politics or court of law, for instance, it is sometimes in the answerer's best interest to be minimally informative while still appearing to address the literal semantics of the question. Indeed, a key legal question is whether such "bad-faith" literal answers can constitute perjury (Solan & Tiersma, 2005). In other cases, questioners may want to conceal or deceive the answerer about their true goals while still obtaining information (e.g. when gathering information to buy a gift.) To understand behavior in these cases it may be necessary to adjust the utility to *minimize* the reduction in uncertainty, or to assume the answerer is further down the recursion hierarchy than we found to be empirically the case in our experiments.

Conclusion

Humans are experts at inferring the intentions of other agents from their actions (Tomasello, Carpenter, Call, Behne, & Moll, 2005; Baker et al., 2009). Given simple motion cues, for example, we are able to reliably discern high-level goals such as chasing, fighting, courting, or playing (Barrett, Todd, Miller, & Blythe, 2005; Heider & Simmel, 1944). A long tradition in psycholinguistics has shown that this expertise extends to speech acts. Behind every question lies a goal or intention. This could be an intention to obtain an explicit piece of information ("Where can I get a newspaper?"), signal some common ground ("Did you see the game last night?"), test the answerer's knowledge ("If I add these numbers together, what do I get?"), politely request the audience to take some action ("Could you pass the salt?"), or just to make open-ended small talk ("How was your weekend?"). These wildly different intentions seem to warrant different kinds of answers. By formalizing the recursive utilities of questioner and

answerer agents and the computational principles by which agents infer each other's intentions from verbal behavior, our framework provides a foundation for re-situating question asking and answering in its social context.

Acknowledgments

We thank Leon Bergen, Judith Degen, Andreas Stuhlmüller, Arianna Yuan, MH Tessler, Dan Lassiter, and Judith Fan for thoughtful conversations and comments. This work was supported by ONR grants N00014-13-1-0788 and N00014-13-10341. RXDH was supported by the Stanford Graduate Fellowship and the National Science Foundation Graduate Research Fellowship under Grant No. DGE-114747.

References

- Allen, J. F., & Perrault, C. R. (1980). Analyzing intention in utterances. *Artificial Intelligence*, 15(3), 143–178.
- Aloni, M. (2005). A formal treatment of the pragmatics of questions and attitudes. *Linguistics and Philosophy*, 28(5), 505–539.
- Andreas, J., Klein, D., & Levine, S. (2017). Modular multitask reinforcement learning with policy sketches. In *Proceedings of the 34th international conference on machine learning-volume 70* (pp. 166–175).
- Asher, N., & Lascarides, A. (1998). Questions in dialogue. *Linguistics and philosophy*, 21(3), 237–309.
- Austin, J. T., & Vancouver, J. B. (1996). Goal constructs in psychology: Structure, process, and content. *Psychological bulletin*, 120(3), 338.
- Badre, D. (2008). Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends in cognitive sciences*, 12(5), 193–200.
- Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, 113(3), 329–349.
- Barrett, H. C., Todd, P. M., Miller, G. F., & Blythe, P. W. (2005). Accurate judgments of intention from motion cues alone: A cross-cultural study. *Evolution and Human Behavior*, 26(4), 313–331.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting linear mixed-effects models using lme4. *arXiv preprint arXiv:1406.5823*.

- Benz, A. (2006). Utility and relevance of answers. In *Game theory and pragmatics* (pp. 195–219). Springer.
- Berant, J., Chou, A., Frostig, R., & Liang, P. (2013). Semantic parsing on freebase from question-answer pairs. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing* (pp. 1533–1544).
- Bergen, L., Goodman, N. D., & Levy, R. (2012). That's what she (could have) said: How alternative utterances affect language use. In *Proceedings of the 34th Annual Conference of the Cognitive Science Society*.
- Bergen, L., Levy, R., & Goodman, N. (2016). Pragmatic reasoning through semantic inference. *Semantics and Pragmatics*, 9(20).
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022.
- Boër, S. E., & Lycan, W. G. (1975). Knowing who. *Philosophical Studies*, 28(5), 299–344.
- Botvinick, M. M. (2008). Hierarchical models of behavior and prefrontal function. *Trends in cognitive sciences*, 12(5), 201–208.
- Bramley, N. R., Lagnado, D. A., & Speekenbrink, M. (2015). Conservative forgetful scholars: How people learn causal structure through sequences of interventions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(3), 708.
- Brown, P., & Levinson, S. C. (1987). *Politeness: Some universals in language usage* (Vol. 4). Cambridge university press.
- Brown, R. (1958). How shall a thing be called? *Psychological Review*, 65(1), 14.
- Büring, D. (2003). On d-trees, beans, and b-accent. *Linguistics and philosophy*, 26(5), 511–545.
- Callanan, M. A., & Oakes, L. M. (1992). Preschoolers' questions and parents' explanations: Causal thinking in everyday activity. *Cognitive Development*, 7(2), 213–233.
- Chaudhry, S. J., & Loewenstein, G. (2019). Thanking, apologizing, bragging, and blaming: Responsibility exchange theory and the currency of communication. *Psychological review*, 126(3), 313.
- Chouinard, M. M. (2007). Children's questions: A mechanism for cognitive development. *Monographs of the Society for Research in Child Development*, i–129.
- Ciardelli, I., Groenendijk, J., & Roelofsen, F. (2013). Inquisitive semantics: a new notion of meaning. *Language and Linguistics Compass*, 7(9), 459–476.
- Ciardelli, I., Groenendijk, J., & Roelofsen, F. (2019). *Inquisitive semantics*. Oxford University Press.
- Clark, H. H. (1979). Responding to indirect speech acts. *Cognitive Psychology*, 11(4), 430–477.
- Clark, H. H., & Marshall, C. (1981). Definite reference and mutual knowledge. In A. Joshi, B. Webber, & I. Sag (Eds.), *Elements of discourse understanding* (p. 10–63). Cambridge University Press.
- Clark, H. H., & Schober, M. F. (1992). Asking questions and influencing answers. In J. M. Tanur (Ed.), *Questions about questions: Inquiries into the cognitive bases of surveys* (pp. 15–48). New York, NY: Russell Sage Foundation.
- Clark, H. H., Schreuder, R., & Buttrick, S. (1983). Common ground at the understanding of demonstrative reference. *Journal of Verbal Learning and Verbal Behavior*, 22(2), 245–258.
- Coenen, A., Nelson, J. D., & Gureckis, T. M. (2018). Asking the right questions about the psychology of human inquiry: Nine open challenges. *Psychonomic Bulletin & Review*, 1–41.
- Cohen, A., & Lake, B. M. (2016). Searching large hypothesis spaces by asking questions. In *Cogsci*.
- Deits, R., Tellex, S., Thaker, P., Simeonov, D., Kollar, T., & Roy, N. (2013). Clarifying commands with information-theoretic human-robot dialog. *Journal of Human-Robot Interaction*, 2(2), 58–79.
- de Marneffe, M.-C., Grimm, S., & Potts, C. (2009). Not a simple yes or no: Uncertainty in indirect answers. In *Proceedings of the SIGDIAL 2009 conference* (pp. 136–143).
- Der Henst, V., Carles, L., & Sperber, D. (2002). Truthfulness and relevance in telling the time. *Mind & Language*, 17(5), 457–466.
- De Ruiter, J. P. (2012). *Questions: Formal, functional and interactional perspectives* (Vol. 12). Cambridge University Press.
- Fong, T., Thorpe, C., & Baur, C. (2003). Robot, asker of questions. *Robotics and Autonomous systems*, 42(3), 235–243.
- Francik, E. P., & Clark, H. H. (1985). How to make

- requests that overcome obstacles to compliance. *Journal of Memory and Language*, 24(5), 560–568.
- Frank, M. C., Emilsson, A. G., Peloquin, B., Goodman, N. D., & Potts, C. (2017). Rational speech act models of pragmatic reasoning in reference games. Retrieved from *osf.io/f9y6b*.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998.
- Franke, M., & Degen, J. (2016). Reasoning in reference games: Individual-vs. population-level probabilistic modeling. *PloS one*, 11(5), e0154854.
- Gall, M. D. (1970). The use of questions in teaching. *Review of educational research*, 40(5), 707–721.
- Gerbrandy, J. (2000). Identity in epistemic semantics. In P. Blackburn & J. Seligman (Eds.), *Logic, language and computation*, vol. iii (pp. 147–159). Stanford, CA: CSLI.
- Gibbs Jr, R. W., & Bryant, G. A. (2008). Striving for optimal relevance when answering questions. *Cognition*, 106(1), 345–369.
- Ginzburg, J. (1995). Resolving questions, i. *Linguistics and Philosophy*, 18(5), 459–527.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11), 818 - 829.
- Goodman, N. D., & Stuhlmüller, A. (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science*, 5(1), 173–184.
- Goodman, N. D., & Stuhlmüller, A. (electronic). *The design and implementation of probabilistic programming languages*. Retrieved 2015/1/16, from <http://dippl.org>
- Graesser, A. C., & Person, N. K. (1994). Question asking during tutoring. *American Educational Research Journal*, 31(1), 104–137.
- Graf, C., Degen, J., Hawkins, R. X., & Goodman, N. D. (2016). Animal, dog, or dalmatian? level of abstraction in nominal referring expressions. In *Proceedings of the 38th annual conference of the Cognitive Science Society*.
- Green, N., & Carberry, S. (1994). A hybrid reasoning model for indirect answers. In *Proceedings of the 32nd annual meeting on association for computational linguistics* (pp. 58–65).
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7(2), 217–229.
- Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, 114(2), 211.
- Groenendijk, J. (1999). The logic of interrogation: Classical version. In *Proceedings of salt* (Vol. 9, pp. 109–126).
- Groenendijk, J., & Stokhof, M. (1984). On the semantics of questions and the pragmatics of answers. *Varieties of formal semantics*, 3, 143–170.
- Gureckis, T. M., & Markant, D. B. (2012). Self-directed learning: a cognitive and computational perspective. *Perspectives on Psychological Science*, 7(5), 464–481.
- Hamblin, C. L. (1973). Questions in montague english. *Foundations of language*, 10(1), 41–53.
- Hawkins, R. X. D. (2015). Conducting real-time multiplayer experiments on the web. *Behavior Research Methods*, 47(4), 966–976.
- Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *The American Journal of Psychology*, 243–259.
- Jackendoff, R. (2007). *Language, consciousness, culture: Essays on mental structure* (Vol. 2007). MIT Press.
- Kao, J. T., Bergen, L., & Goodman, N. D. (2014). Formalizing the pragmatics of metaphor understanding. In *Proceedings of the 36th Annual Conference of the Cognitive Science Society*.
- Kao, J. T., & Goodman, N. D. (2015). Let's talk (ironically) about the weather: Modeling verbal irony. In D. C. Noelle et al. (Eds.), *Proceedings of the 37th annual conference of the Cognitive Science Society*.
- Kao, J. T., Wu, J. Y., Bergen, L., & Goodman, N. D. (2014). Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences*, 111(33), 12002–12007.
- Karttunen, L. (1977). Syntax and semantics of questions. *Linguistics and philosophy*, 1(1), 3–44.
- Khani, F., Goodman, N. D., & Liang, P. (2018). Planning, inference and pragmatics in sequential language games. *Transactions of the Association for*

- Computational Linguistics*, 6, 543–555.
- Kuppevelt, J. v. (1995). Discourse structure, topicality and questioning. *Journal of linguistics*, 31(01), 109–147.
- Lassiter, D., & Goodman, N. D. (2015). Adjectival vagueness in a bayesian model of interpretation. *Synthese*, 1–36.
- Lee, M. D., & Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge University Press.
- Legare, C. H., Mills, C. M., Souza, A. L., Plummer, L. E., & Yasskin, R. (2013). The use of questions as problem-solving strategies during early childhood. *Journal of Experimental Child Psychology*, 114(1), 63–76.
- Lehnert, W. (1977). Human and computational question answering. *Cognitive Science*, 1(1), 47–73.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1), 339–359.
- Markant, D. B., & Gureckis, T. M. (2014). Is it better to select or to receive? learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, 143(1), 94.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual review of neuroscience*, 24(1), 167–202.
- Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior*. Henry Holt and Co.
- Mollá, D., & Vicedo, J. L. (2007). Question answering in restricted domains: An overview. *Computational Linguistics*, 33(1), 41–61.
- Moradlou, S., & Ginzburg, J. (2014). Learning to understand questions. In V. Rieser & P. Muller (Eds.), *Proceedings of the 18th workshop on the semantics and pragmatics of dialogue*.
- Morey, R. D., Romeijn, J.-W., & Rouder, J. N. (2016). The philosophy of bayes factors and the quantification of statistical evidence. *Journal of Mathematical Psychology*, 72, 6–18.
- Morey, R. D., Rouder, J. N., & Jamil, T. (2015). Package BayesFactor [Computer software manual].
- Myung, J. I., & Pitt, M. A. (2009). Optimal experimental design for model discrimination. *Psychological Review*, 116(3), 499.
- Neal, R. M. (2001). Annealed importance sampling. *Statistics and computing*, 11(2), 125–139.
- Nelson, J. D. (2005). Finding useful questions: on bayesian diagnosticity, probability, impact, and information gain. *Psychological review*, 112(4), 979.
- Nelson, J. D., Divjak, B., Gudmundsdottir, G., Martignon, L. F., & Meder, B. (2014). Children's sequential information search is sensitive to environmental probabilities. *Cognition*, 130(1), 74–80.
- Nelson, J. D., Meder, B., & Jones, M. (2018, Dec). *Towards a theory of heuristic and optimal planning for sequential information search*. PsyArXiv. Retrieved from psyarxiv.com/bxdf4 doi: 10.31234/osf.io/bxdf4
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101(4), 608.
- Oh, J., Singh, S., Lee, H., & Kohli, P. (2017). Zero-shot task generalization with multi-task deep reinforcement learning. In *Proceedings of the 34th international conference on machine learning-volume 70* (pp. 2661–2670).
- Palminteri, S., Wyart, V., & Koehlin, E. (2017). The importance of falsification in computational cognitive modeling. *Trends in cognitive sciences*, 21(6), 425–433.
- Potts, C. (2012). Goal-driven answers in the cards dialogue corpus. In *Proceedings of the 30th west coast conference on formal linguistics* (pp. 1–20).
- Rhode, H. (2006). Rhetorical questions as redundant interrogatives. *San Diego Linguistics Papers*, 2(1), 134–168.
- Roberts, C. (1996). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Working Papers in Linguistics-Ohio State University Department of Linguistics*, 91–136.
- Roberts, C. (2003). Uniqueness in definite noun phrases. *Linguistics and Philosophy*, 26(3), 287–350.
- Rojas-Esponda, T. (2013). The road signs of communication. *Proceedings of Semdial (DialDam)*, 17, 131–139.
- Rojas-Esponda, T. (2014). A discourse model for überhaupt. *Semantics and Pragmatics*, 7, 1–1.
- Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychol-*

- ogy: *General*, 104(3), 192-233.
- Rothe, A., Lake, B. M., & Gureckis, T. (2017). Question asking as program generation. In *Advances in neural information processing systems* (pp. 1046–1055).
- Rothe, A., Lake, B. M., & Gureckis, T. M. (2016). Asking and evaluating natural language questions. In *Proceedings of the 38th annual conference of the cognitive science society*.
- Rothe, A., Lake, B. M., & Gureckis, T. M. (2018). Do people ask good questions? *Computational Brain & Behavior*, 1(1), 69–89.
- Ruggeri, A., & Lombrozo, T. (2015). Children adapt their questions to achieve efficient search. *Cognition*, 143, 203–216.
- Ruggeri, A., Lombrozo, T., Griffiths, T. L., & Xu, F. (2015). Children search for information as efficiently as adults, but seek additional confirmatory evidence. In *Proceedings of the 37th annual conference of the cognitive science society*.
- Ruggeri, A., Lombrozo, T., Griffiths, T. L., & Xu, F. (2016). Sources of developmental change in the efficiency of information search. *Developmental Psychology*, 52(12), 2159.
- Salomo, D., Lieven, E., & Tomasello, M. (2013). Children's ability to answer different types of questions. *Journal of Child Language*, 40(02), 469–491.
- Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals, and understanding : an inquiry into human knowledge structures*. Hillsdale, N.J.; New York: L. Erlbaum Associates ; distributed by the Halsted Press Division of John Wiley and Sons.
- Schulz, K., & Van Rooij, R. (2006). Pragmatic meaning and non-monotonic reasoning: The case of exhaustive interpretation. *Linguistics and Philosophy*, 29(2), 205–250.
- Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning from others: the consequences of psychological reasoning for human learning. *Perspectives on Psychological Science*, 7(4), 341–351.
- Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, 71, 55–89.
- Siegler, R. S. (1977). The twenty questions game as a form of problem solving. *Child Development*, 395–403.
- Simmons, R. F. (1965). Answering english questions by computer: a survey. *Communications of the ACM*, 8(1), 53–70.
- Solan, L. M., & Tiersma, P. M. (2005). *Speaking of crime: The language of criminal justice*. Chicago: The University of Chicago Press.
- Stevens, J. S., Benz, A., Reuße, S., Laarmann-Quante, R., & Klabunde, R. (2014). Indirect answers as potential solutions to decision problems. In *Proceedings of the 18th workshop on the semantics and pragmatics of dialogue* (pp. 145–153).
- Tauber, S., Navarro, D. J., Perfors, A., & Steyvers, M. (2017). Bayesian models of cognition revisited: Setting optimality aside and letting data drive psychological theory. *Psychological review*, 124(4), 410.
- Tessler, M. H., & Goodman, N. D. (2018). The language of generalization. *Psychological Review*.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28(05), 675–691.
- Van Rooy, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy*, 26(6), 727–763.
- Wilson, D., & Sperber, D. (2012). *Meaning and relevance*. Cambridge University Press.
- Yoon, E. J., Tessler, M. H., Goodman, N. D., & Frank, M. C. (2016). Talking with tact: Polite language as a balance between kindness and informativity. In *Proceedings of the 38th annual conference of the cognitive science society* (pp. 2771–2776).

	Questioners					Answerers			
	animal	place	plant	artifact		animal	place	plant	artifact
animal	1.00	0.91	0.94	0.94		1.00	0.78	0.92	0.97
place	0.91	1.00	0.96	0.95		0.78	1.00	0.78	0.78
plant	0.94	0.96	1.00	0.97		0.92	0.78	1.00	0.91
artifact	0.94	0.95	0.97	1.00		0.97	0.78	0.91	1.00

Table 3

Appendix: Inter-domain correlations on corresponding response rates