Questions in Decision-Making Dialogues

Jerry R. Hobbs
Information Sciences Institute
University of Southern California
Marina del Rey, California
hobbs@isi.edu

1 Introduction

The research described here aims at investigating how group decisions emerge from the contributions of individual members and the role that questions play in that complex process. As such, the study is located at the interface of cognitive psychology and small-group sociology. Participants in the group make suggestions and ask questions arising from their own thoughts on the issues, under the influence of the group's interaction so far, and somehow in the course of the entire conversation a consensus arises about what actions the group as a whole is going to take. The aim of the research is to arrive at an account of this process, and of the role of questions in that process, that is as formal or computational as possible.

In Section 2 the data that is analyzed is described, viz., transcripts of three-person decision-making meetings. Section 3 describes the theoretical framework, specifically, defining mutual belief, shared plans, and the Conversational Record. In Section 4 seven different levels of processing are described as we go from individual fragments of utterances to a shared plan. In Section 5 questions are examined systematically, from both a syntactic and a pragmatic point of view, to see what role they play in these processes.

2 The Data

Five meetings of three people each were organized.[1] We audiotaped and videotaped the meetings, transcribed the audiotapes, and coded a significant portion of the transcripts for intonation.[2] The participants in the meetings were given fifteen minutes and told to construct a schedule for a visiting government contracts officer. They were given several constraints. The contracts officer was only going to be there for three hours, from 11 a.m. to 2 p.m. They had to schedule lunch during that time. They had to find out about his needs. Each of them had a presentation of a certain length he or she wanted to give, and

[1]The experimental data was collected, transcribed, and analyzed with the aid of Ann Podlozny.
[2]The intonation was coded by Rebecca Brown.

one or two demos; this information was initially unknown to the others. The total time all this would have taken was greater than the three hours allotted, so negotiation was necessary. The participants were artificial intelligence researchers at SRI International.

We wanted to be able to examine the effect of power relationships on the decision-making, so each group of three consisted of A, B, and C participants. A was told he or she had a $1,000,000 contract which was on track and only needed a report on recent progress. B had a $300,000 contract; while this specific project was not in trouble, the research area as a whole was. C had a $150,000 internal research and development project in a hot new area. In order to further distinguish the levels of power among the participants, they were assigned the A, B, or C roles in accordance with their actual seniority and management experience. Some of the groups had worked together closely before; others had not.

Although the specifics of the task they were given were artificial, the type of task was something all were familiar with and did often in their real work experience. All the participants fell very quickly into natural modes of interaction and problem-solving.

3 Theoretical Framework

Social life is built upon mutual belief (e.g., Lewis, 1979; Clark and Marshall, 1981). There are four principal properties of mutual belief. First, if a group of people mutually believe a proposition, each member of the group believes it. Second, the group mutually believes that it mutually believes it. Applying these two rules enables us to derive from mutual belief arbitrary embeddings of the form "A believes that B believes that ... P."

These two principles are necessary conditions on mutual belief. It is not possible to have a finite sufficient condition for mutual belief strictly in terms of belief. But our third principle is that we can successively approximate mutual belief by accumulatating more and more statements of embedded belief. We can imagine situations in which A believes P and B believes P, but they don't mutually believe P because each does not know that the other also believes it. We can also imagine situations in which A believes that B believes P, and B believes that A believes P, but they don't mutually believe P because, for example, B doesn't believe that A believes that B believes P. But the latter situation is much more rare than the former. As we embed beliefs more and more deeply in beliefs, mutual belief becomes more and more probable – the situations in which mutual belief does not hold are less frequent – and the inference of mutual belief is more and more justified.

The final principal property of mutual belief is the heuristic rule that co-presence defeasibly entails mutual belief. Two people who witness the same event, including a conversation they are engaged in, will normally mutually believe the facts of that event. Or at any rate, we need to assume so in order to get social life going. Failures of this rule are relatively rare, and when they occur and people become aware of them, they are generally repairable. In a question-answer exchange, the participants are co-present at the exchange, and we can therefore assume that they now mutually believe the information in the answer, or at least that that information was provided in the answer.

Joint action by a group of people requires a shared plan. Shared plans are characterized by three principal properties. First, the group itself has the top-level goal. For example, in our data, the group consisting of A, B and C have the goal that A, B and C have a successful meeting with the contracts officer. Second, defeasibly, the group mutually believes the subgoal structure of the plan, that is, the steps by which they intend to achieve their mutual goal. This rule is only defeasible because all the members may not be privy to all the details in the plan. Finally, if a member of the group is involved in an action in the plan, then the member is committed to performing his or her role in that action. The primary inference that can be drawn from commitment is that (unfortunately only defeasibly) the other members can rely on the action happening at the appropriate time. Thus, if we have a shared plan that I fall backwards and you catch me, I want to be assured that you are committed to your part of the plan. Much social discourse revolves around making this inference more reliable.

The Conversational Record is the set of mutual beliefs relevant to the conversation. This includes the content and interpretations of utterances. But as we will see in the analysis below, the content the Conversational Record is negotiated as the conversation proceeds. In a task-oriented conversation where the goal is to plan for some joint action, the shared plan that is being constructed is part of the Conversational Record. In our data it includes beliefs about their task of deciding on a schedule, but also about the process they are going through to accomplish this task. The latter includes at least a coarse-grained account of the structure of the conversation, e.g., knowledge of what is being discussed when. The Conversational Record is created by co-presence and by explicit negotiations and agreements arrived at the the course of the conversation, by the processes described in the next two sections.

The Conversational Record is an idealized social construct. An approximation of it resides in each participant's head. These may differ, but normally the differences don't matter, and when they do, the participants are able to do the necessary repair work to bring them into alignment.

The aim of the analysis here is to examine how, from individual contributions, including, crucially, questions, the participants are able to construct a shared plan, or rather, several shared plans—one for the visit itself, one for the task of deciding on a schedule for the visit, and occasionally in evidence, the shared plan for participating in our experiment. In Section 4 we examine how in general these shared plans emerge from the individual, frequently disfluent utterances, pointing out the role of questions in this process. In Section 5 we focus more specifically on questions to examine the forms and functions they take on.

4 Emergence of Shared Plans from Utterances

We can characterize the process of achieving a shared plan as taking place via seven successive levels of analysis, each indicating a different sort of structure.

1. The syntax of repairs
2. The syntax of repaired utterances
3. Co-construction of content segments
4. Global coherence
5. Local coherence
6. Effects of status
7. Decision points

We discuss all of these briefly in turn. Discovering the first three sorts of structure is necessary for identifying when an utterance is a question, and the third sort in particular raises a very interesting issue of what counts as a question. The last four sorts of structure provide arenas in which questions function. A question-answer exchange is one kind of a very important local coherence relation, and thus functions to effect the flow of the conversation. Questions can be used to redirect the global coherence structure of a dialog and to assert one's status in the group.

4.1 The Syntax of Self-Repairs

Very many utterances in our data, as in most conversational data, are disfluent, to the extent that someone unfamiliar with actual spoken dialogue may believe it is hopeless to give a rule-based, computational account of understanding something so messy. However, a great deal of research has been carried out discovering the kinds of disfluencies that occur, the strategies speakers have for repairing them, and the strategies hearers have for interpreting them.

The structure of the self-repairs in our data was very much in line with previous research on self-repairs, e.g., Levelt (1983). It can be expressed by a single rule with nine elements. All the elements are optional, although at least one should appear or it isn't a repair. No examples in our data instantiated the entire pattern, so it will be illustrated with a made-up example.

Initial Segment:	And
Pre-alignment Cue:	so
Mistake:	John
Post-alignment Cue:	said
Edit Signal:	um
Pre-alignment Cue:	so
Correction:	Bill
Post-alignment Cue:	said
Continuation:	we should leave.

The speaker is talking along—"And so". She makes a mistake—"John". She continues—"said"—before realizing she made a mistake. When she does, she utters an edit signal "um" to indicate that a correction is coming up. She repeats a segment before the mistake—"so"–so the hearer can make the correction to right place in what she said; this is the pre-alignment cue. Then she makes the correction—"Bill". Then she repeats what she had said before, after the mistake—"said"; this is the post-alignment cue. Finally she continues on with what she was trying to say to begin with—"we should leave".

In the actual data many of these slots are missing, as illustrated by the following four examples:

Pre-alignment Cue:	have
Mistake:	somebody
Edit Signal:	uh
Pre-alignment Cue:	have
Correction:	the secretaries
Continuation:	just bring us some sandwiches
Mistake:	with
Correction:	for
Pre-alignment Cue:	un'
Pre-alignment Cue:	un-
Continuation:	less
Initial Segment:	the
Edit Signal:	uh
Continuation:	intervening hour

Note that not all "repairs" involve a mistake.

In order to derive a logical form for the utterance, we need to undo the disfluency and convert the utterance into a grammatical sentence to which the rules of compositional semantics can be applied. The general pattern of self-repairs provides exactly what we need to do this. The corrected utterance is the initial segment, followed by one copy of the pre-alignment cue, followed by the correction, followed by one copy of the post-alignment cue, and then the continuation.

This repair rule has to be applied iteratively to many disfluent utterances. For example, the following sentence requires five such repairs, as indicated. The material deleted at each step is italicized.

> Basically *we* we need to, we need to give a convin', uh convincing sort of uh, uh discussion.
>
> Basically *we need to,* we need to give a convin', uh convincing sort of uh, uh

discussion.

Basically we need to give a *convin', uh* convincing sort of uh, uh discussion.

Basically we need to give a convincing sort of *uh,* uh discussion.

Basically we need to give a convincing sort of *uh* discussion.

Basically we need to give a convincing sort of discussion.

The end of this process is a grammatical sentence that can be interpreted by compositional semantics.[3] Self-repairs are not specific to questions, but in a computational treatment of questions, disfluencies must first be fixed if the standard rules of compositional semantics are to be applied to generate a logical form that can then be used for inference, and in particular for inferring function.

4.2 The Syntax of Repaired Utterances

To achieve a shared plan, the participants have to comprehend the others' contributions. The next step in this is a syntactic analysis of the sentences, making it possible to do compositional semantics, which in turn enables inference and interpretation more generally. So the question we address here is whether, after the disfluencies are repaired as in the previous section, the results are grammatical and can be interpreted. For example, can questions be identified as questions, and can it be determined what, on the surface, is being asked?

One fragment of 118 utterances was examined with respect to syntactic structure, after repairs were made. Of these 74 were complete sentences, and five more were complete sentences followed by a conjunction. 23 were sentences that were broken off in the middle someplace other than a conjunction. Thus, 102 of 118 utterances were sentences or on their way to being sentences.

Of the other 16 utterances, there were two finite verb phrases, one infinitive, five subordinate clauses or prepositional phrases, two predicate complements, five noun phrases, and one relative clause. These were ususally repairs or additions to the previous sentence.

Several observations are worth making. Sentence boundaries were not always signalled by clause boundary tones. There were frequent "restated" noun phrases, that syntactically could be thought of as appositives, as in

We don't require a lot of background, *stage-setting.*

He reads roughly speaking one page out of ten, wha' I would guess usually *the abstract and the last page out of ten.*

Topicalizations occur, as in

This guy, um, you don't want to stuff food down his mouth.

But *the demos* he likes to see kind of concrete evidence of thing, of of things actually working.

[3] I have ignored several complications here as this step is needed only to get the process of interpretation going and is not central to the topic of the paper.

There are occasional subcategorization violations, as in

> You tend to have more conversation about what's going on, *as opposed to if you go off to lunch.*

Out of the 118 utterances there was only one truly ungrammatical sentence:

> Okay so now now is the question of who who goes first.
> *Which makes sense to go first.*

To summarize, after repair, the utterances are almost all grammatical clauses, and those that are not are usually noun phrases, verb phrases, prepositional phrases, gerunds, or other "first-class" constituents, possibly preceded by a connective. There are two principal exceptions to this general characterization. First, phrases may be broken off in the middle, possibly to be picked up by another participant, who may back up to a functor word. Second, there may be overlapping phrases in a kind of "shadowing", as in

⌈B: which leaves ten minutes to talk.
⌊A: no time

Here "ten minutes" and "no time" are spoken simultaneously.

The fact that seemingly disfluent discourse can be automatically corrected into grammatical or near grammatical strings, including questions, makes compositional semantics possible, and therefore makes possible deeper processing on the basis of the content of the utterances.

4.3 Co-construction of Content Segments

As is well-known, conversations do not proceed by A making a complete contribution, then B making a complete contribution, and then A making another one. Rather, the participants co-construct meaningful units across multiple turns. This is illustrated in the following fragment:

B: So, if we could, break it up in the, in the following kind of way. I mean, i', if all, if everybody's gonna be at lunch.
C: Right, then then the sounding out about current interest and available funds we could, that's sort of a joint thing.
⌈B: Right.
⌊C: So it could be the three of us together.
B: That's right.
C: And so I don't know, a half hour or something on that during lunch.
B: Right.

The content or message that B and C have co-constructed across six or seven turns could be stated very neatly as "**For a half hour during lunch the three of us together could sound him out about current interests and available funds.**" This final formulation never actually gets said all at once. Nevertheless, it is reasonable to assume that something like this content is what gets entered into the conversational record as a result of this fragment of the dialogue.

This final message consists of four "subpropositions"—"**For a half hour**", "**during lunch**", "**the three of us together**", and "**could sound him out about current interests and available funds**." Subpropositions are worked out incrementally. Often they are first expressed poorly, then expressed well, and then referred to anaphorically while the participants go on to work out the best formulation of another subproposition. The subproposition "**For a half hour**" is stated just once as "I don't know, a half hour or something" by C. The subproposition "**during lunch**" is expressed first by B when he says "everybody's gonna be at lunch", and is then stated by C as "during lunch". The subproposition "**the three of us together**" is stated first by B as "all", then as "everybody", then by C as "a joint thing", and finally by C as "the three of us together". The subproposition "**could sound him out about current interests and available funds**" is stated first by B as "break it up in the, in the following kind of way", then by C as "the sounding out about current interest and available funds", then by C as "that" and "it", and finally again by C as "on that".

All of the other material in this passage is performing one of three functions:

1. Expressing relations among subpropositions, i.e., B's "So, if", and "I mean, i', if all, if", and C's "then then", "So", and "And so".

2. Expressing requests for confirmation, including mitigations, i.e., B's "we could", C's "we could", "sort of", "could be", "I don't know" and "or something".

3. Expressing confirmations, i.e., C's "Right", and B's "Right", "That's right", and "Right".

The working out of subpropositions, expressing the relations among them, and requesting and giving confirmations account for all of the material in this passage, and in most of the data.

Examining this fragment will lead the reader to appreciate the difficulty we face in Section 5 in identifying and classifying questions. As they are co-constructing bits of the shared plan, the participants contribute suggestions like "So it could be the three of us together." These could equivalently be stated as questions: "So should it be the three of us together?" When we look at pragmatic questions, we will see that the boundary between suggestions and yes-no questions is very fuzzy indeed.

Moreover, requests for confirmation can be in the form of questions or can be in the form of mitigations, like "or something". Should the mitigations be viewed functionally as questions? Boundaries are very unclear.

4.4 Global Coherence

The participants are solving a problem in these dialogues, and the structure of the problem imposes a structure on the discourse. Hobbs and Agar (1982) called this "global coherence". Global coherence is the structure that emerges because of the way the discourse coheres with what else is going on in the world, in particular the plans the participants are executing that the dialogue is a part of.

The structure of the problem can be expressed in terms of prerequisite and subevent relations among actions or tasks, as illustrated in Figure 1. This is precisely the material out of which the shared plan has to be built. The ultimate goal is to allocate time slots to speakers and lunch. In order to do this, the participants must decide on the temporal order of the presentations, decide how long each of the presentations will be, and decide on the length of time for lunch. In order to determine a temporal order they must decide on the priority of each project, and to do this they must each describe the monetary value of their projects and describe the extent to which it is at risk. In order to determine the length of time to be given to each person, they need to describe their desired length and, again, to decide on the priority of the projects. To decide on the length of time for lunch they need to decide where lunch will be and decide on an agenda for lunch.

*** FIGURE 1 GOES ABOUT HERE ***

In our data the order of the discourse was almost always consistent with the partial order imposed by the structure of the problem. Discussions of tasks were almost always preceded by discussion of their prerequisites. In one dialogue of 69 utterances, for example, the discussion of topics was as follows:

1-6:	Describe Monetary Value
7-18:	Describe Desired Length
19-22:	Decide Where Lunch
23-26:	Decide Agenda for Lunch
27:	Decide Length of Time for Lunch
28-31:	Decide Project Priority
31-39:	Describe Risk (Violation)
39-44::	Decide Project Priority
45-53:	Decide Temporal Order
54-58:	Describe Desired Length
59-65:	Decide Length of Time for Projects
66-68:	Decide Temporal Order
68-69:	Allocate Time Slots

There is one violation here of the global coherence constraint. It was as follows:

(28) B: Right. And now you've got a million dollars a year.
(29) A: Mhmm.
(30) C: Mhmm.
(31) B: Okay. Then, giving you from essentially one to two, or one to whenever we leave is probably the most uh, uh, advantageous thing. Um, are we all under the same constraints, that, uh, uh in this particular case, my project is uh, uh, in jeopardy currently.
(32) A: Oh, no, no, let's tell, okay, let's let's share you know, what the projects are, okay?

B mentions the risk factor in his project, and that lets them know that all the prerequisites have not been satisfied for the task of deciding on the project priority. The discussion is redirected to another part of the ideal plan, upon the realization that all the relevant information was not yet part of the Conversational Record. Note that A uses a question to effect this redirection.

In general, we found in our data that the discussion of issues was consistent with the partial ordering imposed by the ideal plan, with only two classes of exceptions. The first is illustrated by the above example; a mistake or an omission of important information has occurred in an earlier part of the plan and the plan must be repaired. The second occurs when earlier decisions and earlier information are reviewed to justify a choice in the part of the plan currently being worked out.

4.5 Local Coherence

The order of discussion of issues is underdetermined by the structure of the ideal plan. The structure of the dialogue is underdetermined by global coherence. Considerable freedom remains in the flow of the discussion from one topic to another.

This flow can be characterized by coherence relations between adjacent segments of text (Hobbs, 1985). This level of structure was called "local coherence" by Agar and Hobbs (1982). Local coherence is the tree-like structure that emerges as a result of the relations successive segments bear to each other. Coherence relations are based on relations of similarity, figure-ground, change of state or "occasion", causality, contrast, and elaboration or alternate description as seen in questions and their answers. The relations obtain between the eventualities conveyed in the primary "claims" of the segments of discourse. When a coherence relation occurs between successive segments of discourse, the two segments are composed into a larger segment, whose primary claim is determined from the two component segments and the relation between them. In this way, larger and larger segments are built, yielding a tree structure. Aberrations occur, in which coherence relations obtain between nonadjacent segments, yielding a tree-*like* structure, but these structures can be characerized within the same theory (Hobbs and Agar, 1985).

One common relation is a deep parallelism or similarity, e.g., between projects, as when A describes Project A and then B describes Project B, or between parallel aspects of a

project, as when someone describes his project size, describes the risk, and then describes the desired time.

Another relation is that of contrasting alternatives, as when someone proposes that A go first and someone else proposes that B go first.

A common pattern is problem-solution, which is a kind of causal relation. A problem is a barrier to reaching a goal; the solution is something that causes it not to be a barrier. An example is when one person notes that they have only 3 hours, and someone else suggests that they bring in lunch rather than go out for lunch. A justification is similarly based on causality. An example is when a participant proposes a time slot and then cites the justifying risk factor, as in

> ... would give you one fifteen to two.
> Because that's a neat idea for him to take away.

An important class of local coherence relations is elaborations. In these, the same situation is described in each segment, perhaps from different perspectives. More precisely, there is an elaboration relation if the same proposition can be inferred from the principal claims of each segment being related. A summary is an example of the elaboration relation. A single eventuality is described in two different ways—one long and detailed and the second succinct. Discovering this pattern is important in identifying decision points.

A question and its answer can also be viewed as being in an elaboration relation. In a wh-question the content of the question includes an abstract reference to some sort of identifying information (the wh-word), and in the answer a specialization of that information is provided. One can think of the content of a yes-no question as saying that some situation either does or does not hold, and the answer is a specialization of that, saying which of those possibilities is true. The inferences that need to be drawn in recognizing the elaboration relation are exactly the same as those that need to be drawn to recognize a question-answer pair, especially when the first element is only functionally and not formally a question.

4.6 Effects of Status

Each group had an A, a B, and a C participant, determined by the size of their projects, and correlated with their seniority and management experience in real life. Some of the groups had worked together before and some had not.

In groups that had worked together before, the relative status of the participants was established immediately. In groups that hadn't, relative status was established as soon as the size of contracts were known.

Here are two examples of the start of the session for groups that had worked together before.

(1) ⌈C: Okay, so. Let's see
⌊A: Well, I I have a million dollar project.
(2) A: And yours?

(1) A: It seems like the first issue is for us to figure out what each other's projects are, right?
(2) B: Yeah. Why don't you go first

It is interesting that in both cases, it is a question that A uses to establish his position in the hierarchy.

An even starker example of this occurs in one group that had not worked together before. A and B were of roughly equal seniority and management experience. B starts off trying to lead the discussion, and A jumps in and disagrees with him right away. In (4) she interrupts him to change the topic and issues rapid-fire questions. Then in (12) she announces she has a million dollar project, and from that point on, she is in charge.

(1) B: Well, John's gonna be here for three hours and we ...
(2) A: Well we really only have two hours ...
(3) B: Yeah, but we can talk to him about certain things over
(4) ⌈B: lunch too.
⌊A: Now what I don't know much about either of your projects ...
(5) B: Well, ours is one that ...
(6) A: At what kind of level?
(7) B: Three hundred K a year ...
(8) A: Uh huh, so how long do the demos take?
(9) B: Oh, fifteen to twenty minutes.
(10) A: And how long do you think you need otherwise, to ...
(11) B: Oh maybe about ten minutes ...
(12) A: See, I've got this big million dollar a year project, and ...

A's questions function as a way of establishing a power relationship over B; she exhibits her right to get information from B before beginning to reciprocate. It is sometimes said that questions put the questioner in a lower position. (This is often given as the reason men don't ask for directions.) Here, by contrast, we see questions used as in instrument for the establishment of power.

Status had a definite effect on who talked the most. We counted the number of words each said and overall the A participant said 48% of the words, B said 34% , and C said 18%. Similar results were obtained for the number of utterances. In every case, A and B spoke more than C, and in four out of five sessions, A spoke more than B.

When we look more closely, we see that their roles differ. A is in charge of achieving global consensus and is deferred to in that role. A asks more questions. A is more likely to redirect or refocus the discussion on immediate goals. And A's bad ideas are taken more seriously; in one dialogue, A makes a stupid suggestion and immediately retracts it, but B and C carry on for the next ten turns saying why it is not a good idea, despite A's efforts to change the topic.

B is not responsible for global consensus. Thus, B is freer to pursue his own personal obsessions, and the B participants frequently do.

C's contributions are either rejected or appropriated by A or B. C never defends his or her suggestions. A and B defend their own ideas; C lets others defend them.

These roles are abstract characterizations of the kind of actions each of the participants are involved in in the developing shared plan for conducting the conversation.

4.7 Decision Points

We identified the utterances at which a decision was made about some aspect of the schedule. The first example is the decision about the slot for C's presentation:

A: W', so let's see. So, so let's see. If if one fifteen to two was the Charlie time

The next is the decision about the slot for A's presentation. Note that A refers to himself in the third person.

A: Okay. So let's, let's, let's make it the Adam is eleven to twelve.

The next is the time for B's demos.

A: Then then you start, then twelve o'clock you start giving your demos.
⌈B: So from twelve to twelve thirty is demo and presentation.
|C: Mhmm.
⌊A: Yeah.

The next is the time for lunch.

A: So let's make it lunch starts at twelve thirty.

Identifying the decision points is crucial if we are to determine the group action that emerges from the individual contributions, and thus to know what the shared plan is that emerges. These examples illustrate some of what we found.

Syntax is of no help. This is shown in the example

> If if one fifteen to two was the Charlie time ...

This sentence was not even completed, and yet this slot is thereby settled. After A breaks off, there is no further discussion of it, and this is what appears in the final schedule.

Lexical signals are of no help. One might think the word "so" signalled a decision. In fact in the first session, five out of six decision points are introduced by "so". But there are 54 occurrences of "so" in the first session, and only five of them signal a decision point.

Status matters. In every case, the decision points are utterances by A, although sometimes they are echoed by others.

Local coherence matters. The decision point is always a summary of the preceding discussion. But of course this is not easy to determine computationally.

4.8 Summary of the Emergence of Consensus

The structure of a computational procedure for discovering the decision points in decision-making dialogues is illustrated in Figure 2. The syntactic structure of disfluencies has to be discovered and repairs have to be made, as described in Section 4.1. The repaired utterances then have to be parsed, with a few concessions to robustness, as described in Section 4.2, and compositional semantics has to be applied; this is the step that enables the recognition of questions as questions. The co-construction of content segments has to be analyzed into their subpropositions and recomposed into the content of the segment, as described in Section 4.3. These three together yield formal representations of the information in the content segments. This is input for the next two levels of analysis. The content segments have to be matched up with the tasks and subtasks in the ideal plan for solving the problem the group is presented with. This yields the global coherence, as described in Section 4.4. The flow from one segment to another within the structure imposed by the ideal plan needs to be characterized in terms of local coherence relations, including the relation between questions and their answers. These two processes yield the large-scale structure of the discourse. This information together with knowledge of the status of the participants, and thus of their roles, can then be used to determine the decision points, and hence the shared plan that emerges from the dialogue.

*** FIGURE 2 GOES ABOUT HERE ***

5 Questions

5.1 Syntactic Questions

Section 4 presented the overall model of how participants go from disfluent and fragmentary utterances to an agreed-upon shared plan, pointing out some of the roles of questions along the way. We now turn to questions specifically and systematically, and address their various roles in this overall model for the construction of a shared plan. First we classify questions by their syntactic structure. Then we classify them by their pragmatic function. We observe that there is quite a bit of divergence between these classifications. That is, things are not always what they seem. In fact, they rarely are.

There are 99 questions in the data we collected, if we judge something to be a question by syntactic or intonational criteria. I will look first at how these are distributed across syntactic question types. Seven types of question were identified:

Standard wh-questions: 25 instances.

How much of that do you want to do your pleading?

In-place wh-questions: 2 instances.

That gets us up to what?

Yes-no questions: 27 instances.

So um, does that seem ... long enough?

Elliptical yes-no questions: 28 instances.

Right?

Elliptical alternative questions: 4 instances.

While the demo's going on or prior to the demo?

Syntactically declarative sentences: 8 instances.

We're allowing forty five minutes for lunch?

Elliptical declaratives: 3 instances.

Well, I I have a million dollar project. And yours?

There are cases which are ambiguous among the types. Consider the following exchange.

B: How critical is it that you talk to him?
A: At all, you mean?

We could call A's response an elliptical wh-question, where the full sentence is "How critical is it that I talk to him at all?" A would be repeating and elaborating on B's question. Or it could be an elliptical yes-no question, where the full question is "Do you mean 'at all'?" Or it could be an elliptical form of a syntactic declarative, where the full statement is "You mean 'at all'."

A second ambiguous case is

For one year or does it say?

Is this an elliptical alternative question or a yes-no question?

A common story on the intonation of questions is that wh-questions typically end in a low boundary tone (L%), whereas yes-no questions typically end in a high boundary tone (H%). Although this account is controversial, we found it borne out in our data. We found that 62% of wh-questions ended in L%, 79% of yes-no questions ended in H%, and 80% of elliptical yes-no questions ended in H%. (See Pierrehumbert and Hirschberg, 1990; also see Sadock, this volume, for a more detailed analysis of intonation in questions.)

A common story on alternative questions is that the first disjunct will have an H% boundary tone and the last will have L%. This was true in all the examples in our data, but there were too few cases to draw any conclusions.

It would be expected that a declarative or an elliptical declarative functioning as a question would have an H% boundary tone. It is unsurprising that elliptical yes-no questions would be so frequently signalled by intonation, since syntax does not signal their function as a question. This was almost always the case in our data, although again our data was sparse. There was one case of a declarative with an L% boundary tone functioning as a question, namely, B's response in the following:

A: Right?
B: You mean like if he gets here late [L%]?

This can be identified as a question because it is a statement about the hearer's mental state that the hearer would have privileged access to.

We see that questions can be signalled syntactically (standard wh-questions and yes-no questions), lexically (in-place wh-questions), intonationally (in declaratives and elliptical yes-no questions), and by their content as statements of information that the hearer has privileged access to.

5.2 Pragmatic Questions

The form and function of utterances are notoriously at odds with each other in normal discourse (cf., e.g., Levinson, 1983). The syntactic form of a question is not always a good indicator of its pragmatic function. The 99 questions in our data were categorized into six pragmatic functions. In this section I first define each of the categories. I then present the statistics on the occurrences of these functions. Finally, I discuss each of the categories, with examples, noting common ambiguities of function.

The six pragmatic functions of questions are as follows:

Wh-question: A request for someone who knows to add a property or relation to the Conversational Record.

Yes-no question: A request for someone who knows to add a binary judgment (P or not P, for some P) to the Conversational Record.

Alternative question: A request for someone who knows to add either P or Q to the Conversational Record, for some P and Q.

Suggestion: A request to add to the Conversational Record an action to the shared plan for the task or some step in the decision-making process; this will require agreement from others before it is accomplished.

Requests for confirmation: A request that others agree to one's own contribution, thereby placing it in the Conversational Record.

Check: A request to delay agreement on someone else's contribution until further discussion.

All of these involve manipulations of the Conversational Record, and often of shared plans.

The 99 questions in our data were categorized according to their pragmatic function, or where ambiguous their two possible functions. The results were as follows:

Wh-questions:	29
Wh-questions or yes-no questions:	5
Yes-no questions:	6
Alternative questions:	4
Alternative or yes-no questions:	2
Suggestions:	6
Suggestions or yes-no questions:	10
Requests for confirmation:	25
Requests for confirmation or yes-no questions:	7
Checks:	2
Checks or yes-no questions:	2

Note that 27 out of 99 questions (the indented items) have ambiguous functions. These cases are discussed below.

Syntactically, there were four different types of questions that invited a "yes" or "no" answer—yes-no questions, elliptical yes-no questions, declaratives, and elliptical declaratives. Altogether there were 66 utterances in these categories—two thirds of all the questions in the data. Yet when we look at their pragmatic functions, only six end up unambiguously as yes-no questions. The conclusion we can draw from this is that yes-no questions usually aren't.

Let us now consider each category of pragmatic function in detail.

Wh-Questions: Wh-questions are requests to add some property or relation to the Conversational Record. 68% of the time this was a request for private information that only the hearer would know, as in

> And how much time did you need for your stuff?

and

> Well, I I have a million dollar project. And yours?

as well as

<knock on door>
Yes?

In this last example, "yes?" means something like "Who's there?" or "What do you want?"

There were also requests for clarification or elaboration, as in

A: And, instead of having the half hour that we'd like to talk with him, we could take a bit less on the assumption that we kind of handle lunch in a certain way.
B: Now what way is that?

There were verifications of something that was already in the Conversational Record:

Then for Project C, how much time are we allowing for that?
That gets us to to what?

The first question is asking a repetition of something they had already agreed on. The second question is asking for a simple mathematical inference from what is already mutually known.

There were also requests for suggestions, as in

Okay, that leaves lunch, right? Okay, where are we going?

Thus we see wh-questions used as requests to add information to the set of mutual beliefs or to add a task to a shared plan, as requests for more precise specification of what someone else is bidding to add, or as requests for double-checks on something that has already been added.

Yes-no questions: A yes-no question is a request to add a binary judgment to the Conversational Record. This category included requests for private information, as in

Anybody know what he likes? Chinese?

This example illustrates two very common phenomena. The first sentence is syntactically a yes-no question, but pragmatically one can claim it functions as the wh-question "What does he like?" (But see below.) Questions about someone's knowledge are usually a request for that knowledge. The second sentence—"Chinese?"—is a genuine yes-no question. Yes-no questions are frequently preceded by wh-questions, as in this example. First the speaker requests a bit of specific knowledge from the hearer. Then he or she provides a possibility of what that knowledge might be, in the form of a yes-no question.

Yes-no questions can also be questions about what is already in the Conversational Record, as in

Who'll be doing their presentation during that time?
Did we discuss that?

The second question asks whether a particular topic has already occurred, hence a reference not to the shared plan instantiated in the schedule, but to the shared plan for constructing the schedule.

One yes-no question was a joke:

> Do we get a bonus if we don't take fifteen minutes to do this?

In general, I found in doing the classification that the more I looked for yes-no questions, the fewer I found. By far most of the utterances that look on the surface to be yes-no questions are on inspection seen to be playing a much more specific function, asking for specific properties or relations rather than mere binary judgments, or suggesting additions to a shared plan.

Wh-question/yes-no question ambiguities: Five utterances were ambiguous between wh-questions and yes-no questions. Four were as follows:

> Anybody know what he likes?
>
> Is there any conventional wisdom about the order of?
>
> Has he indicated what he wants to see while he's here?
>
> D'you know anything about, uh, his interest of your project?

All of these have the syntactic structure of yes-no questions, but they are all questions about knowledge, mutual knowledge, or received communication that the hearer would have special access to. They are partially legitimate yes-no questions. The addressee can cooperatively answer "no". But one can't cooperatively answer just "yes". More information has to be given. In a sense, the questions unpack into two instructions, the second one conditional:

> Tell me if you know P.
> If so, tell me P.

The fifth ambiguous case was

> What was it [L-L%] three hundred [H-H%]

The first part of this utterance ends with a low phrasal tone and a low boundary tone. The second part ends with a high phrasal tone and a high boundary tone. The ambiguity here is resolved if we consider this two utterances rather than one. The first part is a wh-question and the second a yes-no question, exemplifying the pattern described above where a piece of knowledge is requested and then a possible answer suggested.

Alternative questions: An alternative question is a request to add either P or Q to the Conversational Record for some P and Q. This was used to request private information, as in

> Do you know that he's really looking forward to it, or is this completely outa the blue?

This was a reference to the project description that only the addressee was given.

Alternative questions were also used to check on something already in the Conversational Record:

> For you Chris, you had one demo or two demos?
>
> While the demo's going or or prior to the demo?

Again the functions are to add information or tasks to the Conversational Record or to check on information or tasks that are already there.

Alternative/yes-no question ambiguities: There were two utterances that were ambiguous between alternative questions and yes-no questions. The first is

> For one year, or does it say?

In a way, this is the reverse of the pattern described above of wh-question via epistemic state followed by a possible answer. One could paraphrase this utterance as "Does it say [how long your project goes on]? For one year?"

The other ambiguous example is spread across several turns.

B: We're allowing forty-five minutes for lunch?
⌈A: Oh, I see. Okay, so you
⌊B: Or were you starting lunch at twelve?

B's first utterance is a legitimate pragmatic yes-no question. B's second utterance turns it into the first part of an alternative question. A "yes" answer to the one implies a "no" answer to the other and conversely.

Suggestions: A suggestion is a request to add some step in the task or some step in the decision-making process to the Conversational Record, requiring agreement from others. Several examples of questions were bids to add a task or task constraint to the the schedule being developed in the Conversational Record.

> Why don't we have the client do his presentation right away?
>
> Can we have a working lunch?
>
> Do you wanna do this sequentially? Uh, that is, all all talk and demo of one project at a time?

The last of these examples is interesting, because what is being suggested is the opposite of the content of the question, since he is describing exactly the order they had been assuming. That is, the speaker is suggesting, by asking the question, that they not present the projects sequentially, but rather all the talks first and then all the demos.

Other examples of suggestions were bids to enter an action in the decision-making process into the Conversational Record.

Why don't you go first?

Should we go get them?

The first is a suggestion that the addressee describe his project first. The second is a suggestion that the task is completed and they should notify the people running the experiment.

Thus, suggestions can be about the shared plan they are developing for future joint action, or about the process they are engaged in right now for coming up with this shared plan.

Suggestion/yes-no question ambiguities: It is often very difficult to decide whether something should be classified as a suggestion or as a legitimate pragmatic yes-no question. The following examples are listed in order, from those more like suggestions to those more like yes-no questions.

Why don't you go first?

Well, we could do that presumably in each uh individual thing?

Do you want to do that, Charlie?

Does the lunch have to be a full hour [H-L%]?

Do you wanna try and do a different kind of schedule?

Can we perhaps schedule an earlier lunch?

Does lunch have to be a full hour [H-H%]?

The fourth and seventh examples differ only in whether the intonation indicates that the issue is closed or still open.

I find it impossible to draw a line in this spectrum between the two categories. In a sense, a yes-no question is a deniable way of making a suggestion that runs counter to what has been agreed on already.

Confirmations and requests for confirmation: A request for confirmation is a request that the other participants agree that one's own contribution should be added to the Conversational Record. A request and confirmation are a solution to the problem of how mutual belief is determined. They indicate that the proposition has been heard, understood and agreed with, thus making it subject to the co-presence heuristic.

Among the examples in our data were cases where the contribution was a suggested task constraint, as in

I think that we bring in lunch, right?

Is that okay with you, Chad?

There were cases where the contribution was a constraint on the scheduling process.

Let's let's share y'know what the projects are, okay?

There was a case where the contribution was a justification of a task constraint.

I think that we bring in lunch, right?
I mean, we certainly can't afford to go to the I building for lunch, right?

There was one example of a request for confirmation of an arithmetic calculation.

Eleven twenty to twelve, that's uh forty minutes, right?

There were two cases of requests for confirmation that their task was completed.

So, so I think we've got it, right?

So are we set?

Request for confirmation/yes-no question ambiguities: It was sometimes ambiguous whether an utterance was a request for confirmation or a yes-no question. In the following two examples the contribution was a suggested task constraint.

B: You should write down on the thing, just project A on the top one.
⌈B: You're, are you Project A? Project A, yeah.
⌊A: Project A, yeah.

B is not entirely certain that A's project is "Project A", so his question can be viewed as a request for confirmation of that, or as a genuine question about whether it is true. Now consider

B: <suggests ordering of presentations>
... ... [10 turns]
B: Is is this, acceptable to you? [to C, after A agrees]

B's question to C could be viewed as a request to confirm that the schedule B proposed can go into the Conversational Record. Or it could be viewed as a legitimate yes-no question about C's stance toward something that is already in the Conversational Record.

There was one example of a request for confirmation that some arithmetic was correct.

Eleven twenty to twelve, that's us forty minutes, right?
Is that right?
Twenty and forty. Yeah.

The second time he asks if it is right is almost an admission that he doesn't know, and therefore the question can be viewed as a legitimate yes-no question.

Finally, at least one of the confirmation-of-completion requests could be viewed as a genuine request for information.

So, we're done?

Checks: A check was defined as a request to delay agreement on someone else's contribution until further discussion has taken place. There were four examples of this in the data. The first is

C: And then from one to one-fifteen, maybe would be a more detailed presentation of, of uh ...
⌈B: No, I I I don't necessarily want to schedule it, um
⌊C: No?

Here C is making a suggestion, B rejects it, and C questions the rejection. The check is C's second utterance—"No?" This exchange is a kind of negotiation about what should be placed in the Conversational Record.

Next consider

B: I'd rather actually do that first, I'd think.
C: Really? I would've I would've thought that you'd wanna show him the stuff he likes best first.

C's "Really?" objects to adding B's contribution to the Conversational Record.

Check/yes-no question ambiguities: The final two examples are ambiguous between checks and legitimate pragmatic yes-no questions.

A: Yeah, during lunch, yeah.
B: Okay, so starting lunch at twelve thirty. We're allowing forty five minutes for lunch?
A: Oh, I see. Okay, so you
B: Or were you starting lunch at twelve and giving demos through lunch?

B's last utterance could be viewed as a legitimate yes-no question about what A had thought the global consensus was. Or it could be viewed as a request for further discussion on what it was that was being added to the Conversational Record. In

B: How critical is it that you uh talk to him?
A: At all, you mean?

it could be that A is objecting to B's implied suggestion that A not talk to the visiting contracts monitor, and hence a check on what goes in the shared plan, or it could be a legitimate yes-no question about what B was meaning to say.

Pragmatic declaratives, in a sense, deposit information in the Conversational Record, and suggestions do so conditioned on their acceptance. Pragmatic wh-questions, yes-no questions, and alternative questions request that information be deposited in the Conversational Record. When we get to requests for confirmation, confirmations, and especially

checks, the Conversational Record is not just the target of information exchange but the matter under discussion, as the participants negotiate about what should go there, i.e., whether they were right in what they have already agreed on in the construction of a shared set of beliefs and shared plans for joint action.

6 Summary

When people engage in a conversation, they build on a base of mutual beliefs, and they extend that base in the course of the conversation. Frequently the function of the conversation is to construct a shared plan that will enable joint action. Even when the conversation is not task-oriented in this way, the process of creating and nurturing the conversation itself can be viewed as a matter of constructing a shared plan. Questions play a vital role in this construction.

The construction of a shared plan can be viewed as layers of processes, from the level of individual words produced by speakers and understood by hearers, to the level of the group as a whole adopting a joint plan of action. We identified seven such levels, from the management of self-repairs in single utterances, to the "decision point" statements that solidify the shared plan. This emergent joint plan of action is an important part of the Conversational Record that the participants construct while engaging in the dialogue. Questions can be examined as phenomena at many of these levels, e.g., in terms of their syntactic form after disfluencies have been repaired, as a local coherence elaboration relation between questions and answers, or as instruments for asserting power as the roles of the participants are being sorted out.

We have several means for indicating that an utterance is a question; we can do so syntactically, lexically, intonationally, and by means of content, i.e., by making a statement whose content the hearer has special access to. We can also identify categories of pragmatic functions for questions, in terms of whether they request information for mutual belief, request a task for the shared plan, suggest a task for the shared plan, request confirmation that information or a task should be added to the Conversational Record, or request that something in mutual belief or the shared plan be reconsidered. But there is no neat mapping from the syntactic form of questions to their pragmatic function, nor is there a neat categorization of questions into pragmatic function. Ambiguities abound. The speaker's intentions and the hearer's interpretations are not always clear-cut. The picture that emerges is of people's flexibility in accomplishing a variety of conversational functions with a variety of linguistic means. Computationally, this means that to determine what count as questions and what pragmatic functions they have, we must exploit all available information, including their syntactic form, intonation, their content in particular as it relates to the participants' knowledge states, the flow of the conversation from one utterance to the next, and crucially the relation between the utterance and the shared plans that are being built in the course of the conversation.

Acknowledgments

I am indebted to Ann Podlozny for helping me collect this data, for transcribing it, and for helping me analyze it. I am indebted to Rebecca Brown for annotating it for intonation, and for discussions about the function of the intonation.

References

[1] Agar, Michael, and Jerry R. Hobbs, 1982. "Interpreting Discourse: Coherence and the Analysis of Ethnographic Interviews", *Discourse Processes*, Vol. 5, No. 1, pp. 1-32.

[2] Clark, Herbert H., and Catherine R. Marshall, 1981. "Definite Reference and Mutual Knowledge", in A. Joshi, B. Webber, and I. Sag, eds., *Elements of Discourse Understanding*, pp. 10-63, Cambridge University Press, Cambridge, England.

[3] Hobbs, Jerry R., 1985. "On the Coherence and Structure of Discourse", Report No. CSLI-85-37, Center for the Study of Language and Information, Stanford University.

[4] Hobbs, Jerry R., and Michael Agar, 1985. "The Coherence of Incoherent Discourse", *Journal of Language and Social Psychology*, vol. 4, nos. 3-4, pp. 213-232.

[5] Levelt, Willem, 1983. "Monitoring and Self-Repair in Speech", *Cognition*, Vol. 14, No. 1, pp. 41-104.

[6] Levinson, Stephen C., 1983. *Pragmatics*, Cambridge Textbooks in Linguistics, Cambridge University Press, Cambridge, England.

[7] Lewis, David, 1979. "Scorekeeping in a Language Game," *Journal of Philosophical Logic*, Vol. 6, pp. 339-59.

[8] Pierrehumbert, Janet, and Julia Hirschberg, 1990. "The Meaning of Intonational Contours in the Interpretation of Discourse", in *Intentions in Communication*, P. Cohen, J. Morgan, and M. Pollack, editors, Bradford Books (MIT Press), Cambridge, Massachusetts, pp. 271-312.