Question–Answer Games (extended abstract)

Thomas Ågotnes[1], Johan van Benthem[24], Hans van Ditmarsch[3], and Ştefan Minică[2]

[1] Inform. Science and Media Studies, University of Bergen, Norway
thomas.agotnes@infomedia.uib.no
[2] ILLC, University of Amsterdam, the Netherlands
{s.minica,johan.vanbenthem}@uva.nl
[3] University of Sevilla, D. Logic, Sevilla, Spain
hvd@us.es
[4] Stanford University, USA

1 Introduction

I am in the middle of Manhattan and I am lost. I approach a likely stranger on the road and ask him, pointing to the right: "Is this the way to the railway station?" Discounting for a moment the possibility that I will be ignored and that the stranger will continue on his way (of which the—trivial—game move equivalent would be "I prefer not to answer the question"), I will now get one of three possible answers: "Yes,", "No,", and "I don't know." Each answer is informative. After the third answer I may ask the same question to someone else. Questions and answers can be analyzed in dynamic epistemic logic, and right questions and right answers in game theory. We focus on dynamic epistemic logic and not on game theory, and our main contribution is a novel link between the two. In this introduction we now first address the *question*, then the *answer*, and lastly the *game* wherein such questions and answers figure. The railway example misses that game aspect, but now consider having to decide between asking the conference chair "Are the outcomes of the submissions already known?" or "Is my submission accepted?" What the best question is to ask, also depends on the answers you may expect or fear.

Question Suppose I am a, you are b, and that p is the atomic proposition that the railway station is to the right. The other direction, to the left, is therefore represented by $\neg p$. There are several pragmatic preconditions for the agent asking the question 'p?'. She should not know the truth about p, i.e., she does not know p, and she does not know $\neg p$. We are assuming a multi-agent epistemic logic to model questions, where the expression $K_\mathsf{a} p$ stands for 'a knows p,' and where the epistemic modality K_a is interpreted with an equivalence relation $\sim_\mathsf{a}$. This pragmatic constraint therefore amounts to the precondition $\neg K_\mathsf{a} p \land \neg K_\mathsf{a} \neg p$. The agent asking the question also has expectations about the agent that will answer the question, and that constitutes another pragmatic precondition. She considers it possible that he knows the answer, i.e., $\neg K_\mathsf{a} \neg (K_\mathsf{b} p \lor K_\mathsf{b} \neg p)$. (She also considers it *likely* that he knows the answer, i.e., she believes that tentatively, $B_{\mathsf{a}}(K_{\mathsf{b}}p \vee K_{\mathsf{b}}\neg p)$, where belief and knowledge are combined as in [8]; we may also see that as a combination of preferences and knowledge [13].)

A question φ? splits the domain in the set of states $[\![\varphi]\!]$ where φ is true and the set of states $[\![\neg\varphi]\!]$ where φ is false, i.e., the question induces a dichotomy on the domain of the model. In the approach by van Benthem and Minica [14], a question or *issue* φ? is represented by the issue relation $\approx_\varphi$. Such issue relations go back to [3, 6].

Answer As said, there are three possible answers: 'Yes', 'No', and 'I don't know'. (A fourth possible answer would be: 'I decline to answer the question'.) Such a response is not more informative than required. For example, assume a deal of cards over players. If I were to ask you "Do you have card 0" and you answer me "Yes, I have 0, 1, and 2" then you give me more than I asked for. You could have answered "Yes," i.e., "Yes, I have 0." If your answer to the question φ? is 'yes', you confirm that you know that φ, and *nothing more than that* is required. This means that your answer corresponds to the (unique) largest union of (b) equivalence classes representing your knowledge that **is contained in** the φ-states of the model. This is of course exactly the denotation of the formula $K_{\mathsf{b}}\varphi$. Your answer is therefore the *public announcement* $K_{\mathsf{b}}\varphi$! [11]. (We assume that the questions are answered truthfully.) Similarly, if your answer is "No, I don't," this is an announcement of the formula $K_{\mathsf{b}}\neg\varphi$ and its denotation is the complement of the (unique) smallest union of equivalence classes that **contains** the φ-states. If you answer is "I don't know" you get the remainder, i.e. the union of all $\sim_{\mathsf{a}}$ classes that properly intersect with $\approx_\varphi$.

This in fact shows that answers to question can be seen as *rough sets* [10]. Given the set $[\![\varphi]\!]$ (i.e., the subset of the domain consisting of the φ-states), in rough set terms known as the *target*, take the lower and upper $\sim_{\mathsf{b}}$ approximation of the target, i.e., $\underline{\sim_{\mathsf{b}}}([\![\varphi]\!])$ and $\overline{\sim_{\mathsf{b}}}([\![\varphi]\!])$. If the answer to the question is 'yes', the actual state is in the lower approximation. If the answer is 'no', the actual state is in the complement of the upper approximation. If the answer is 'don't know', the actual state is in the upper approximation minus the lower approximation.

In dynamic epistemic logic, a public announcement is interpreted as a model restriction. Therefore, answering the question can be seen as executing one of three possible such restrictions, a non-deterministic program so to speak. Alternatively (and equivalently!), we can see answering the question as a refinement of the equivalence classes for all agents (and not just of the agent asking the question) with the issue relation $\approx_\varphi$, i.e., for all agents a, $\sim'_{\mathsf{a}} = \sim_{\mathsf{a}} \cap \approx_\varphi$ [14].

Game If we want to play a game with questions and answers, to start with it whould be clear what we are playing for. What are the goals of the players? In the case of me asking for the right way to the railway station, my goal is knowledge about p: $K_{\mathsf{a}}p \vee K_{\mathsf{a}}\neg p$. But in this case it is not so clear who I am playing against. Clearly not against the stranger I am addressing with this question. He has no interest. He answers the question, but he does not play a game.

This is different in case you and I are both spies after a particular secret. The secret is the truth about p. Your goal is to get to know it before me and my goal is to get to know it before you. I.e., each agent has a goal, and $\gamma_{\mathsf{a}} = (K_{\mathsf{b}}p \vee K_{\mathsf{b}}\neg p) \rightarrow (K_{\mathsf{a}}p \vee K_{\mathsf{a}}\neg p)$ whereas $\gamma_q = (K_{\mathsf{a}}p \vee K_{\mathsf{a}}\neg p) \rightarrow (K_{\mathsf{b}}p \vee K_{\mathsf{b}}\neg p)$. In other words, I don't care if you know it, as long as I already know it, and vice versa. The result of questions and answers is an information state wherein we can check for each player whether his goals are fullfilled and that determines a payoff function and thus the outcome of the game.

So in order to play a game with questions and answers the players need a goal, and that goal can be an epistemic formula. Why should a player answer a question if that means giving away information that may make him lose the game? He has no reason whatsoever. However, just like in real life, if you wish the other person to loosen his information strings, you may only expect to obtain that by giving away some information yourself as well. The proceduralized version of this expectation, that we will apply in this contribution, is a game where each player may *choose* between different questions to ask but where the other player addressed by that question is *obliged* to answer. The information content of questions is a standard topic in the analysis of multi-player games, such as Mastermind [2, 7] (and related to information theory and entropy). Such games do not necessarily involve strategic response to questions chosen by other agents. We aim to integrate such analyses in game theory. Relations between game theory and pragmatic phenomena like the information content of questions are also studied in [18].

Simplifications In the remaining discussion, we make some simplifications for concreteness.

- We disregard pragmatic constraints of questions. You may ask a question to which you already know the answer. In other words, the question φ? is not also an informative update / public announcement of $(\neg K_{\mathsf{a}}\varphi \wedge \neg K_{\mathsf{a}}\neg\varphi) \wedge \neg K_{\mathsf{a}}\neg(K_{\mathsf{b}}\varphi \vee K_{\mathsf{b}}\neg\varphi)$. Note that incorporating such constraints is quite doable. It merely restricts the players' strategies.
- Avoiding to answer the question is not modelled as a move in the game. However, it is not problematic, as that response can be modelled as the trivial announcement.
- There are two players only, that ask each other questions. If there are more than two players, one has specify who is addressed by the question. Again, this is doable, and an answer to the question would still be a public announcement.
- The most natural interaction between players involving questions and answers is where they ask each other questions in turn, such that a question is answered before the next question is asked. Extensive game forms for imperfect information games, with sequential equilibria, are harder to analyze than one-shot games in strategic form [4]. For now, we assume that the two players ask each other a single question, and that they ask the question at the same time; say, by writing down the question on a piece of paper, putting

it in an envelope, and then exchanging envelopes. For the answer, they again exchange envelopes.

Consider a game where a strategy does not consist of choosing which question to ask to another player but choosing which announcement to make yourself. The different questions for me to ask are exactly the different announcements for the respondent to make—the analysis must be similar. Such games have been coined *public announcement games* in [1]. Our work merges their approach with the analysis of question dynamics in [14]. Somewhat similarly to public announcement games, the *knowledge games* in [15, 16] treat the more general case where the question is public but the answer may be semi-public: the other players know what the question is, but may only partially observe the answer. E.g., the question may be to show a card *only* to the requesting player but there is common knowledge of some card being shown; the alternatives are the different cards to be shown; [16] contains summary game theoretical results.

Summarizing, the motivation for our research is to model conversation including the dynamics of questions and answers, to provide new links between game theory and dynamic logics of information, and to exploit the dynamic/strategic structure that, we think, lies implicitly inside epistemic models for epistemic languages, and to make that structure an explicit subject of logical study. The novel contributions that we present in this extended abstract are the notion of a two-person question-answer game with information goals, the existence and computation of Bayesian equilibria for these games, and novel connections between logic and game theory such as the existence of equilibria (subject to restrictive conditions) for positive goal formulae. A technical appendix defines the logical and game theoretical terminology used in the continuation.

2 Question-answer games

2.1 Pointed and induced question-answer games

Given two agents a and b, an epistemic model $M = (S, \sim_{\mathsf{a}}, \sim_{\mathsf{b}}, V)$ encodes their uncertainty about facts and about each other; Two formulas γ_{a} and γ_{b} in the logical language express what they wish to achieve by their questions. In order to achieve their goals, agent a asks a question φ? to agent b, to which b is obliged to respond with ‘yes’ (I know that φ), ‘no’ (I know that $\neg\varphi$), or ‘don’t know’ (I don’t know whether φ). “And similarly for b asking a question to a:” we don’t want to keep saying that all the time, so from now on we may refer to the two agents as i and j, where $i \neq j$, and i may be either a or b. We assume both agents ask their question simulaneously, and that subsequently both agents answer the question simultaneously. (Of course, more realistic communicative settings would allow for agents to ask questions and respond to them in any order, but such generalizations are harder to model as strategic games.) The question formulas can be thought of as defining the strategies for the agents.

Executing the strategy φ for agent i can be thought of as follows. Agent i asks φ? to j. If $M, s \models K_j\varphi$, then j answers (announces) "Yes, I know that φ". If $M, s \models K_j\neg\varphi$, then j answers "No, I know that $\neg\varphi$". Otherwise, j answers "I don't know whether φ". The resulting model restriction depends on both answers, e.g., if a asks φ? to which b responds $K_{\mathsf{b}}\varphi!$ and b asks ψ? to which a responds $K_{\mathsf{a}}\neg\psi!$, the result is the restricted model $M|(K_{\mathsf{b}}\varphi \wedge K_{\mathsf{a}}\neg\psi)$. We can capture these alternatives with a construct $\overline{K}_i\varphi$, for 'agent i answers the question φ?', defined as follows. Given an epistemic model M and a state $s \in M$, if $M, s \models K_i\varphi$, then $\overline{K}_i\varphi \equiv K_i\varphi$; if $M, s \models K_i\neg\varphi$, then $\overline{K}_i\varphi \equiv K_i\neg\varphi$; and else $\overline{K}_i\varphi \equiv \neg(K_i\varphi \vee K_i\neg\varphi)$. This is reminiscent of the *resolution on* φ in [14].

Alternatively, we can represent the question by an issue relation $\approx_\varphi$ and the public announcement of answering the question in the link-cutting way of [13, 14]. As we combine a question and the answer to it into a single strategy in this simplified question-answer game, we do not need this more detailed formalization.

We associate two different strategic games with these questions (and their answers) and goals: pointed question-answer games and not-pointed question-answer games. Both are needed: a player may not know what the actual state is, and therefore not know which game he is playing. Except for the more complex payoff function of the pointed game, these definitions are elementary adaptations of similar concepts in [1].

Definition 1 (Pointed question-answer game). *The* state game *or* pointed question-answer game $G((M,s), \gamma_{\mathsf{a}}, \gamma_{\mathsf{b}})$ *associated with state* $s \in M$ *of goals* γ_{a} *and* γ_{b} *for agents* a *and* b *respectively, is the strategic game defined by*

- $N = \{\mathsf{a}, \mathsf{b}\}$*;*
- *for* $i = \mathsf{a}, \mathsf{b}$*,* $A_i = \{\varphi? \mid \varphi \in \mathcal{L}\}$*;*
- *for* $i = \mathsf{a}, \mathsf{b}$*,* $u_i^s(\varphi, \psi) = \begin{cases} 1 \text{ if } M, s \models \langle(\overline{K}_{\mathsf{b}}\varphi \wedge \overline{K}_{\mathsf{a}}\psi)!\rangle\gamma_i \\ 0 \text{ otherwise} \end{cases}$

Note that the set of strategies A_i is the same in all states. As a state independent perspective on question-answer games we propose the following definition. It can be easily shown that this corresponds to a Bayesian game [4], in the sense that it has the same Nash equilibria; we will address that below. In Definition 2, a strategy a_i for player i is *uniform* iff for all $s, t \in S$: $s \sim_i t$ implies $a_i(s) = a_i(t)$.

Definition 2 (Question-answer game). *Given state games* $G((M,s), \gamma_{\mathsf{a}}, \gamma_{\mathsf{b}})$ *for each* $s \in M$*, the* induced game *or* question-answer game $G(M, \gamma_{\mathsf{a}}, \gamma_{\mathsf{b}}) = \langle N, \{A'_i : i \in N\}, \{u_i : i \in N\}\rangle$ *is the strategic game defined by*

- $N = \{\mathsf{a}, \mathsf{b}\}$*;*
- *for* $i = \mathsf{a}, \mathsf{b}$*,* A'_i *is the set of uniform functions from* S *to* A_i*;*
- *for* $i = \mathsf{a}, \mathsf{b}$*,* $a'_{\mathsf{a}}, a'_{\mathsf{b}} \in A'_i$*,* $u_i(a'_{\mathsf{a}}, a'_{\mathsf{b}}) = \frac{\sum_{s \in S} u_i^s(a'_{\mathsf{a}}(s), a'_{\mathsf{b}}(s))}{|S|}$.

As there are many formulas in the language, and possibly many states in a model, this seems like a lot of strategies, but we can make some major simplifications.

Two strategies $\varphi?$ and $\psi?$ are the same for agent i if $M \models K_j\varphi \leftrightarrow K_j\psi$. (If $K_j\varphi \leftrightarrow K_j\psi$ is a model validity, then it is also *common knowledge* to a and b that this is a model validity, i.e., it is common knowledge to a and b whether two strategies are the same.) Subject to that identification, the number of strategies for i is a function of the number of the equivalence classes for agent j in the model M. In practice we will use just any questions (strategies) that have our fancy and that are equivalent to $K_j\varphi$, i.e., if a knows that either $K_{\mathsf{b}}p$ or $K_{\mathsf{b}}\neg p$, then she considers strategy (question) $p?$.

Also, given the requirement of uniform strategies in the induced game, instead of seeing a strategy in the induced game as a function from *states* to formulas, we can also see such a strategy for agent i as a function from *i-equivalence classes* to formulas, and therefore, to make life simpler, as a function from *formulas* characterizing i-equivalence classes to formulas.

With these simplifications there are fewer strategies. (See the example, later.) If the strategies for i have the form $K_j\varphi?$, there are only two and not three answers, namely only 'yes' and 'no'. If strategies are required to have this form for both agents, we call this a *dichotomous question-answer game*. The number of strategies is now even lower. It can be counted as follows. (The count is only meaningful on $\{\mathsf{a},\mathsf{b}\}$-connected models.) Assume that player i has m_i equivalence classes and player j has m_j equivalence classes. Then player i has $2^{m_j m_i - m_i}$ pure strategies. This we can see as follows. There are $2^{m_j - 1}$ different dichotomies for player j (i.e. coursenings of player j's partition), and for each of m_i different equivalence classes for the requesting player i, she may choose one of those questions, therefore the total number of pure strategies is $(2^{m_j-1})^{m_i} = 2^{m_j m_i - m_i}$.

Example 1. Consider a three-state model where a knows the truth about p and b knows the truth about q. (For the state where p is true and q is false we write $p\overline{q}$, etc.)

$$\overline{q}p \;\overset{\mathsf{a}}{\text{————}}\; qp \;\overset{\mathsf{b}}{\text{————}}\; q\overline{p}$$

If p and q are both true, player a is uncertain about q, and b is uncertain about p, if p is true and q is false, b knows that q is false, but a is uncertain about q; etc. Given state qp, player a has two strategies, namely asking for the truth about q, the strategy $q?$, or asking the trivial question, $\top?$. Only, a does not know that this is the actual state! For the induced game, a has four strategies, namely, in words: 'If I know p then I ask $q?$, otherwise I ask $q?$', 'If I know p then I ask $q?$, otherwise I ask $\top?$', 'If I know p then I ask $\top?$, otherwise I ask $q?$', 'If I know p then I ask $\top?$, otherwise I ask $\top?$'. If we see a strategy, that is a function from states to questions, as a set of pairs, the second strategy of a corresponds to the strategy $\{(\overline{q}p, q?), (qp, q?), (q\overline{p}, \top?)\}$, etc. Figure 1 contains a sweeping overview of a question-answer game for that model, for the goal formulas $K_{\mathsf{b}}p \rightarrow K_{\mathsf{a}}q$ for i, and $K_{\mathsf{a}}\neg q$ for b. For goal formulas $\gamma_{\mathsf{a}} = (K_{\mathsf{b}}p \vee K_{\mathsf{b}}\neg p) \rightarrow (K_{\mathsf{a}}p \vee K_{\mathsf{a}}\neg p)$ and $\gamma_{\mathsf{b}} = (K_{\mathsf{a}}p \vee K_{\mathsf{a}}\neg p) \rightarrow (K_{\mathsf{b}}p \vee K_{\mathsf{b}}\neg p)$ as in the introduction, we get (of course) another strategic game form for the induced game. Note that in this simple

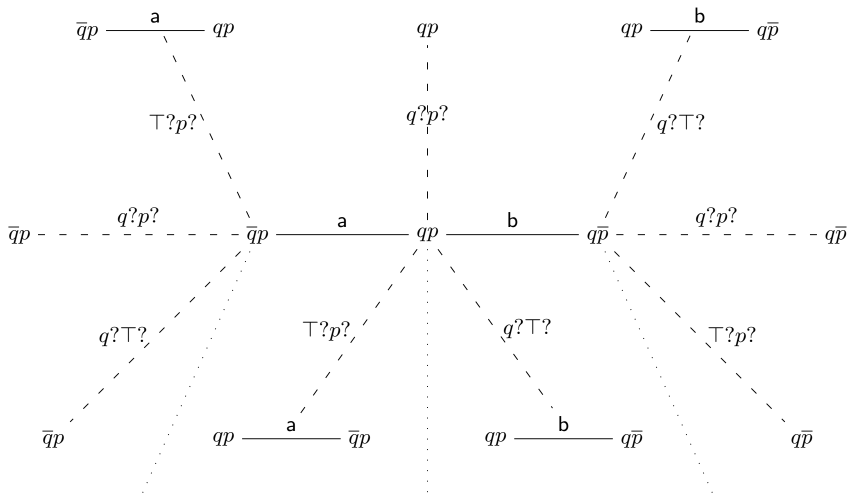

	$\top?$	$p?$
$\top?$	00	00
$q?$	$\underline{01}$	$\underline{01}$

	$\top?$	$p?$
$\top?$	$\underline{10}$	00
$q?$	$\underline{10}$	$\underline{10}$

	$\top?$	$p?$
$\top?$	$\underline{10}$	$\underline{10}$
$q?$	$\underline{10}$	$\underline{10}$

$A_{\mathsf{a}} = \{a_0^{\mathsf{a}}, a_1^{\mathsf{a}}, a_2^{\mathsf{a}}, a_3^{\mathsf{a}}\}$	$A_{\mathsf{b}} = \{a_0^{\mathsf{b}}, a_1^{\mathsf{b}}, a_2^{\mathsf{b}}, a_3^{\mathsf{b}}\}$
$a_0^{\mathsf{a}} = \{(\overline{q}p, \top?), (qp, \top?), (q\overline{p}, \top?)\}$	$a_0^{\mathsf{b}} = \{(\overline{q}p, \top?), (qp, \top?), (q\overline{p}, \top?)\}$
$a_1^{\mathsf{a}} = \{(\overline{q}p, \top?), (qp, \top?), (q\overline{p}, q?)\}$	$a_1^{\mathsf{b}} = \{(\overline{q}p, \top?), (qp, \top?), (q\overline{p}, p?)\}$
$a_2^{\mathsf{a}} = \{(\overline{q}p, q?), (qp, q?), (q\overline{p}, \top?)\}$	$a_2^{\mathsf{b}} = \{(\overline{q}p, p?), (qp, p?), (q\overline{p}, \top?)\}$
$a_3^{\mathsf{a}} = \{(\overline{q}p, q?), (qp, q?), (q\overline{p}, q?)\}$	$a_3^{\mathsf{b}} = \{(\overline{q}p, p?), (qp, p?), (q\overline{p}, p?)\}$

$(a_{\mathsf{a}}, a_{\mathsf{b}})$	$\overline{q}p$	qp	$q\overline{p}$	M
$(a_0^{\mathsf{a}}, a_0^{\mathsf{b}})$	00	10	10	20
$(a_0^{\mathsf{a}}, a_1^{\mathsf{b}})$	00	10	10	20
$(a_0^{\mathsf{a}}, a_2^{\mathsf{b}})$	00	00	10	10
$(a_0^{\mathsf{a}}, a_3^{\mathsf{b}})$	00	00	10	10

$(a_{\mathsf{a}}, a_{\mathsf{b}})$	$\overline{q}p$	qp	$q\overline{p}$	M
$(a_1^{\mathsf{a}}, a_0^{\mathsf{b}})$	00	10	10	20
$(a_1^{\mathsf{a}}, a_1^{\mathsf{b}})$	00	10	10	20
$(a_1^{\mathsf{a}}, a_2^{\mathsf{b}})$	00	00	10	10
$(a_1^{\mathsf{a}}, a_3^{\mathsf{b}})$	00	00	10	10

$(a_{\mathsf{a}}, a_{\mathsf{b}})$	$\overline{q}p$	qp	$q\overline{p}$	M
$(a_2^{\mathsf{a}}, a_0^{\mathsf{b}})$	01	10	10	21
$(a_2^{\mathsf{a}}, a_1^{\mathsf{b}})$	01	10	10	21
$(a_2^{\mathsf{a}}, a_2^{\mathsf{b}})$	01	10	10	21
$(a_2^{\mathsf{a}}, a_3^{\mathsf{b}})$	01	10	10	21

$(a_{\mathsf{a}}, a_{\mathsf{b}})$	$\overline{q}p$	qp	$q\overline{p}$	M
$(a_3^{\mathsf{a}}, a_0^{\mathsf{b}})$	01	10	10	21
$(a_3^{\mathsf{a}}, a_1^{\mathsf{b}})$	01	10	10	21
$(a_3^{\mathsf{a}}, a_2^{\mathsf{b}})$	01	10	10	21
$(a_3^{\mathsf{a}}, a_3^{\mathsf{b}})$	01	10	10	21

	a_0^{b}	a_1^{b}	a_2^{b}	a_3^{b}
a_0^{a}	$\underline{20}$	$\underline{20}$	10	10
a_1^{a}	$\underline{20}$	$\underline{20}$	10	10
a_2^{a}	$\underline{21}$	$\underline{21}$	$\underline{21}$	$\underline{21}$
a_3^{a}	$\underline{21}$	$\underline{21}$	$\underline{21}$	$\underline{21}$

Fig. 1. Agents a and b play a question-answer game. The goal formulas are $K_{\mathsf{b}}p \rightarrow K_{\mathsf{a}}q$ for a, and $K_{\mathsf{a}}\neg q$ for b. The dotted lines connect states to the strategic game forms for the three state games. The dashed lines connect states to the model restriction resulting from the answers to the questions labeling the lines. At the bottom is the strategic game form for the induced game—the payoffs are not divided by 3, e.g., 21 stands for an expected payoff of $(0.66, 0.33)$. The Nash equilibria are underlined. E.g., in the induced game $(a_3^{\mathsf{a}}, a_3^{\mathsf{b}})$ is a Nash equilibrium wherein both players always ask the real and not the trivial question.

example, we cannot get three answers to a question: that requires at least three equivalence classes for the responding agent.

2.2 Bayesian games

The induced game is a Bayesian game. (The proof outline repeats the result in [1].) A signal t_i for player i corresponds to an i-equivalence class. In a Bayesian game, the combination of a player i and a signal t_i defines a virtual player (i, t_i), who has the same strategies $a_i \in A_i$ at his disposition. But this amounts to the same as our definition involving the same players i employing uniform (across equivalence classes!) strategies a'_i that are conditional from states to strategies $a_i \in A_i$, and therefore can be seen as conditional from equivalence classes to strategies a_i.

In our simplified modelling, all states in a given Kripke model get equal a priori probability of $\frac{1}{|S|}$, i.e., uniform over the entire domain. A more general approach would have a given probability distribution as a parameter in the modelling of the game. However, a uniform distribution is a reasonable assumption from the perspective of the observer or modeller of such a multi-agent system: given common knowledge of the structure of the model, as usual in multi-$S5$ conditions, there is no reason to prefer one state over another one. For example, if the Kripke model represents uncertainty about card deals, and the cards are shuffled and drawn blindly from the pack by the players, there is no reason to consider any given card deal (possible world / state) more likely than any other card deal.

After receiving their signal, each agent conditionalizes the probability mass over its equivalence class, i.e., $Pr(s) = \frac{1}{[s]_{\sim_i}}$ for each state inside the class, and 0 outside. For example, continuing our parellel with card games, after the cards have been dealt and a player has picked up his cards, the agent only considers card deals possible wherein she holds that hand of cards, but no longer any of the remaining card deals. Further to this, in the absence of information to the contrary (i.e., assuming 'fair play') each of the possible deals of cards wherein she holds that hand of cards are considered equally likely.

This provides the key to view an induced question-answer game as a Bayesian game. In induced games the payoff for agent i is computed as

$$u_i(a'_{\mathsf{a}}, a'_{\mathsf{b}}) = \frac{\sum_{s \in S} u_i^s(a'_{\mathsf{a}}(s), a'_{\mathsf{b}}(s))}{|S|}$$

whereas for a Bayesian game, one would get a sum (see [1])

$$\sum_{s \in S} Pr(s, t_i) u_i^s(a'_{\mathsf{a}}(s), a'_{\mathsf{b}}(s)) = \sum_{s \in S} \frac{u_i^s(a'_{\mathsf{a}}(s), a'_{\mathsf{b}}(s))}{[s]_{\sim_i}}$$

Although these sums may be different, they induce the same order on payoffs and thus they induce the same Nash equilibria. This justifies the following proposition.

Proposition 1 **([1]).** *An induced question-answer game corresponds to a Bayesian game with the same Nash equilibria.*

2.3 Positive goals

Consider the fragment of the *positive formulae* of $\mathcal{L}$:

$$\varphi ::= p \mid \neg p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid K_i\varphi \mid [\varphi]\varphi$$

where $p \in \Theta$. This notion of positive formulae is found in [17], which is an extension of several such notions going back to [12] who observed that purely epistemic (without announcement operators) positive formulae are *preserved* under submodels. We can now observe that, if both players have goal formulae that are positive, *every pointed question-answer game has a Nash equilibrium.* We can easily grasp why this is true. Let s be the actual state of Kripke model, for which we play the game. If player i asks a question to player j that makes j reveal all he knows, either i's goal is now realized, or there is no way she can realize that goal. This is because if there were a weaker announcement by j also realizing that goal, a further model restriction would preserve the goal, as it is positive. Therefore, asking that question is a weakly dominant strategy for i, for the pointed question-answer game for state s. And as this holds for both players, that must be a Nash equilibrium of the pointed question-answer game for state s. As s was arbitrary, this holds for all pointed games. Unfortunately, the players may not know what that question is... We can therefore call this a *de dicto* Nash equilibrium.

Proposition 2. *If both players' goals are positive formulae, each pointed question-answer game has a Nash equilibrium.*

This result makes for a real difference between public announcement games and question-answer games. From a player's perspective, there is such a thing as a most informative announcement (tell them all you know). But the question that elicits the most informative answer from another player cannot be called the most informative question from the questioning player's point of view. In a different state in the same equivalence class for that player, the question to elicit the most informative answer may be a different question, as the responding player may be in a different equivalence class there. Indeed, in public announcement games the strategy profile consisting of maximal informative announcements is a *de re* Nash equilibrium, but this does not translate into a similar equilibrium in a question-answer game. (The *de dicto*/*de re* distinction is well known in the knowledge and action literature [5].)

3 Further research

We have shown how epistemic models come with natural games that model interesting phenomena, and suggest interesting logical questions. Our games are very simple, but this starting point itself is an advantage, since well-chosen simple games are the way to go in our view. Nevertheless, it is now time to add some urgent further structure. Most natural from the viewpoint of our general aims are a few extensions on which we will present initial results:

- model theory and axioms for appropriate logics describing our games; including issues like bisimulation invariance and fixed-point definability;
- extensive games with longer sequences of moves, and logics for and sequential equilibria of such extensive games;
- a richer account of questions as possible moves of inquiry;
- connections with existing approaches to logics of inquiry and learning;
- non-uniform probability distributions;
- multiple goals per agent (in fact any partial order on the set of all pointed model restrictions of the initial model), instead of a single goal;
- more than two agents, in scenarios where questions are adressed to one in a group and where strategic answers also depend on further questions expected from others in the group;
- games where the trivial answer ('I decline to answer the question' / announcement of $\top$) is always allowed; or games where this is outruled.

Extensive games come into the picture if we consider a game wherein players may ask questions to each other in turn and respond to those questions in turn. We now get an extensive game consisting of several moves. If we assume the pragmatic precondition that the answer to the question should not yet be known, that was mentioned in the introductory section, asking the question is a different move from answering the question. Now, the analysis in [14] becomes essential to properly model these games. Subject to a number of reasonable restrictions, such as that trivial questions are not permitted, a question-answer game consists of a finite number of such moves. The number is indeed finite: the initial model is finite, and therefore there is only a finite number of proper model restrictions; on condition that all announcements are really informative, there are therefore only a finite number of moves in the game.

Acknowledgment

We thank the two anonymous LOFT reviewers for their enthusiasm and their comments, both of which were truly appreciated by the authors.

References

1. T. Ågotnes and H. van Ditmarsch. What will they say? - Public announcement games. Accepted for Knowledge, Rationality & Action. Presented at Logic, Game Theory and Social Choice 6, Tsukuba, August 2009, 2010.
2. M.M. Flood. Mastermind strategy. *Journal of Recreational Mathematics*, 18(3):194–204, 1985–1986.

3. J. Groenendijk and M. Stokhof. Questions. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 1055–1124. Elsevier, 1997.
4. J.C. Harsanyi. Games with Incomplete Information Played by 'Bayesian' Players, Parts I, II, and III. *Management Science*, 14:159–182, 320–334, 486–502, 1967–1968.
5. W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63:185–219, 2004.
6. K. Kelly. *The Logic of Reliable Inquiry.* Oxford University Press, 1996.
7. B.P. Kooi. Yet another mastermind strategy. *ICGA Journal*, 28(1):13–20, 2005.
8. S. Kraus and D. Lehmann. Knowledge, belief and time. *Theoretical Computer Science*, 58:155–174, 1988.
9. M.J. Osborne and A. Rubinstein. *A Course in Game Theory.* MIT Press, 1994.
10. Z. Pawlak. *Rough Sets: Theoretical Aspects of Reasoning about Data.* Kluwer Academic Publishers, 1992.
11. J.A. Plaza. Logics of public communications. In M.L. Emrich, M.S. Pfeifer, M. Hadzikadic, and Z.W. Ras, editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems: Poster Session Program*, pages 201–216. Oak Ridge National Laboratory, 1989.
12. J. van Benthem. 'One is a lonely number': on the logic of communication. In Z. Chatzidakis, P. Koepke, and W. Pohlers, editors, *Logic Colloquium '02.* Association for Symbolic Logic, 2006. Volume 27 of *Lecture Notes in Logic.*
13. J. van Benthem and F. Liu. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logics*, 17(2):157–182, 2007.
14. J. van Benthem and S. Minica. Toward a dynamic logic of questions. In X. He, J.F. Horty, and E. Pacuit, editors, *Logic, Rationality, and Interaction. Proceedings of LORI 2009*, pages 27–41. Springer, 2009. LNCS 5834.
15. H. van Ditmarsch. The description of game actions in Cluedo. In L.A. Petrosian and V.V. Mazalov, editors, *Game Theory and Applications*, volume 8, pages 1–28. Nova Science Publishers, 2002.
16. H. van Ditmarsch. Some game theory of Pit. In C. Zhang, H.W. Guesgen, and W.K. Yeap, editors, *Proceedings of PRICAI 2004 (Eighth Pacific Rim International Conference on Artificial Intelligence)*, pages 946–947. Springer, 2004. LNAI 3157.
17. H. van Ditmarsch and B.P. Kooi. The secret of my success. *Synthese*, 151:201–232, 2006.
18. R. van Rooij. Quality and quantity of information exchange. *Journal of Logic, Language and Information*, 12:423–451, 2003.

Technical appendix

Public Announcement Logic The language $\mathcal{L}$ of public announcement logic [11] over a set of agents $N = \{1, \ldots, n\}$ and a set of primitive propositions Θ is defined as follows, where i is an agent and $p \in \Theta$: $\varphi ::= p \mid K_i\varphi \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid [\varphi_1!]\varphi_2$. We write $\langle\varphi_1!\rangle\varphi_2$ resp. $\hat{K}_i\varphi$ for the duals $\neg[\varphi_1!]\neg\varphi_2$ and $\neg K_i\neg\varphi$.

A *Kripke structure* or *epistemic model* over N and Θ is a tuple $M = (S, \sim_1, \ldots, \sim_n, V)$ where S is a set of states, $\sim_i \subseteq S \times S$ is an epistemic indistinguishability relation that is assumed to be an equivalence relation for each agent i, and $V : \Theta \rightarrow S$ assigns primitive propositions to the states in which they are true. A *pointed Kripke structure* is a pair (M, s) where s is a state in M. In this paper we will assume that Kripke structures are *finite*.

The interpretation of formulae in a pointed Kripke structure is defined as follows (the other clauses are defined in usual truth-functional way): $M, s \models K_i\varphi$ iff for every t such that $s \sim_i t$, $M, t \models \varphi$; and $M, s \models [\varphi!]\psi$ iff $M, s \models \varphi$ implies that $M|\varphi, s \models \psi$, where $M|\varphi = (S', \sim'_1, \ldots, \sim'_n, V')$ such that $S' = \{s' \in S : M, s' \models \varphi\}$, $\sim'_i = \sim_i \cap (S' \times S')$, and $V'(p) = V(p) \cap S'$.

Strategic Game A *strategic game* is a triple $G = \langle N, \{A_i : i \in N\}, \{u_i : i \in N\}\rangle$ where: $N = \{1, \ldots, n\}$ is the finite set of *players*; for each $i \in N$, A_i is the set of *strategies* (or *actions*) available to i. $A = \times_{j \in N} A_j$ is the set of *strategy profiles*; and for each $i \in N$, $u_i : A \to \mathbb{R}$ is the *payoff function* for i, mapping each strategy profile to a number. Notation $(a_1, \ldots, a_n)[a_i/a'_i]$ stands for the profile wherein strategy a_i is replaced by a'_i. A strategy profile is a (pure strategy) *Nash equilibrium* if every strategy is the best response of that agent to the strategies of the other agents, i.e., if the agent can not do any better by choosing a different strategy given that the strategies of the other agents are fixed. Formally, a profile $(a_1, \ldots, a_n)$ is a Nash equilibrium if and only if for all $i \in N$, for all $a'_i \neq a_i$, $u((a_1, \ldots, a_n)[a_i/a'_i]) \leq u(a_1, \ldots, a_n)$. A strategy for an agent is *weakly dominant* if it is at least as good for that agent as any other strategy, no matter which strategies the other agents choose. Formally, a strategy a_i for agent i is weakly dominant if and only if for all agents j, for all a'_j, $u(a'_1, \ldots, a'_n) \leq u((a'_1, \ldots, a'_n)[a'_i/a_i])$.

Bayesian game The most common model of strategic games with imperfect information is the *Bayesian game* [4]. Our presentation of Bayesian games is as in [9]. A *Bayesian strategic game* $BG = \langle N, S, \{A_i : i \in N\}, \{T_i : i \in N\}, \{Pr_i : i \in N\}, \{\tau_i : i \in N\}, \{u_i : i \in N\}\rangle$ has the following components: N is the set of players; S is the finite set of *states* s modelling the players' uncertainty about each other; and for each $i \in N$: A_i is the set of strategies; T_i is the set of *signals* t_i that may be observed by player i, $\tau_i : S \to T_i$ is the *signal function* of player i; Pr_i is a probability measure on S (the *prior belief* of player i) such that $Pr_i(\tau^{-1}(t_i)) > 0$ for all $t_i \in T_i$, i.e., each player's signal is correct with strictly positive probability; and finally u_i is a payoff function on the set of probability measures over $A \times S$ (instead of a payoff function on the set of action profiles A).