

Communication Strategies in Games

Jelle Gerbrandy*

1 Introduction

In recent years, dynamic epistemic semantics has become a mature field of research that provides a sophisticated analysis of different types of epistemic actions in a multi-agent setting (Baltag et al. (1998), Baltag and Moss (2004), Gerbrandy (1998), Kooi (2003), van Ditmarsch (2000)). It seems time that these new tools are put to work in what is *the* theory of multi-agent interaction: game theory. The purpose of this paper is to sketch some ways in which this can be done.

The idea is this. Players are always never completely informed about all aspects of the game they are playing – they may be under-informed about the exact parameters of the game, or they may not be sure what the other players will do, or have done. This lack of information is usually modeled by combining a game-theoretical structure with a representation of the information of the players. Epistemic actions that arise from observations or from communication between players can then be seen as changing the information that players have; these changes are reflected in the epistemic structure of the game; the result of an epistemic action is a new game.

In this paper, we formulate a number representations of games with imperfect information that make them look like a Kripke models in which the states are game-theoretical objects. This type of representation allows us to use the tools of dynamic epistemic semantics to model the effect of epistemic actions in a game in a straightforward way. We discuss a number of issues that arise, and we provide the reader with a worked-out example, but for the rest, the paper consists of definitions and open questions only.

Combining dynamic semantics with game theory works in two ways. Firstly, it can *provide the theory of epistemic actions with a pragmatics*: it allows us to talk about, for example, an act of communication being ‘good’ for an agent because it improves her expected payoff, or of a communication act being ‘credible’ because the speaker has no conceivable reason for lying about what he says. At the same time, the combination *provides game theory with a general theory of epistemic actions*. To be sure, there exists an extended and sophisticated discourse on communication in the game-theoretical literature, but it seems to miss the versatile approach to multi-agent communication and observation that is characteristic of the theories of dynamic semantics.

*Dipartimento di Informatica, Università di Torino. Comments to jelle@gerbrandy.com are welcome. Research supported by Lagrange Project of the Fondazione CRT. With thanks to the anonymous referees for their useful comments.

To illustrate some of the questions that are at stake, consider the following simple example.

Example 1.1 (three envelopes) Ann has to choose one of three closed envelopes. One envelope is empty, one envelope contains three yuan for Ann and three yuan for Bill, the third contains 6 yuan for Bill (Ann gets nothing). These facts are commonly known. In addition, Bill knows which envelope contains the money. Ann has no idea.

In this game, Ann can do no better than to choose an envelope at random. She may then expect a payoff of 1 yuan, and Bill may expect 3 (the average of the money in the three envelopes).

Suppose Bill is allowed to communicate with Ann, but only by saying things that are true (say he communicates via an independent referee). What is smart for Bill to say? He could tell Ann that the second envelope contains 3 yuan for each of them; Ann, believing him, would then choose this envelope. This particular communication act secures, but does not improve, Bill's expected payoff of 3 yuan. Bill can do better by providing Ann with *less* information, and tell her, for example, that it is either the second or the third envelope that contains money for both of them, without specifying which. Ann can now improve her chances by choosing one of these two envelopes, and Bill can expect a payoff of $4\frac{1}{2}$.

Suppose now that there is no referee: there is no way for Ann to check the truth of what Bill says (apart from opening an envelope, of course). If Ann is gullible, Bill's best option is to lie, and tell her that there is 3 yuan for her in the third envelope.

So Ann is suspicious of what Bill says. She does not believe Bill when he says that the third envelope contains 3 yuan for her, and, by a symmetric argument, she will not believe him either if he, this time truthfully, claims that it is the second. But she probably *should* believe him if he says that it is either the second or the third: Bill could not possibly gain anything by lying about this. \square

2 Games in strategic form

2.1 Semantics

In games in strategic form, all players choose a strategy simultaneously. In a game of incomplete information in strategic form, the players may be partially informed about certain crucial aspects of the game, such as the payoffs or the information of other players. We can define of strategic game with imperfect information as a special kind of probabilistic Kripke model in the spirit of Fagin and Halpern (1994), with games doubling as states:

Definition 2.1 (game of incomplete information in strategic form)

A *game in strategic form* is a tuple $(N, (S_i)_{i \in N}, (u_i)_{i \in N})$, with N a set of players, S_i a set of strategies for player i , and u_i a utility function for i .

A *game in strategic form with incomplete information* (henceforth, 'a game') is a probabilistic Kripke model $(G, (\rightarrow_i)_{i \in N}, (p_i)_{i \in N}, g)$, where G is a set of games, \rightarrow_i is a relation over $G \times G$, and p_i assigns to each world a probability function over G , and g is a game from G .

We will abuse notation somewhat, and write Gg for a game such as the above. A game needs to satisfy the following conditions:

- $p_i(g')(g'') > 0$ iff $g' \rightarrow_i g''$.
- If $g' \rightarrow_i g''$, then $p_i(g') = p_i(g'')$
- Each player has the same strategies available in all games in G □

Intuitively, in a game Gg , g represents the payoff matrix of the game that is actually being played, and $g \rightarrow_i g'$ means that i believes that it is possible that the actual game being played is the one represented by Gg' . If $p_i(g)(g') = \alpha$, this means that, in g , i believes with probability α that she is playing game Gg' .

The conditions are there to make sure that this interpretation makes sense. The first condition states that the accessibility relation, that represents the beliefs of the player, is completely determined by the probability distribution in each world.¹ This forces the relations to be transitive, Euclidean and serial – a model of consistent belief.² The second condition states that each players knows what probability she assigns to each game; the third condition states that she knows what strategies are available to her. These last two conditions are not really essential, but they conveniently force our models to be similar to a Bayesian game, which is the standard notion of a strategic game of incomplete information in game theory (see, for example, Osborne and Rubinstein (1998)).

The main difference with a Bayesian game is that we allow for players to have *false* beliefs. This assumption is important for studying epistemic actions: without it, for example, it becomes very difficult to model the effect of a player getting information without the other players being aware of it. Of course, we also need this assumption to model lying or misleading, or simply being wrong.

At this point, it is important to remark that a great part of the game-theoretical analysis of games of incomplete information does not depend on the assumption that the information of the players be true. In particular, many of the standard solution concepts (Nash equilibrium, iterated elimination of dominated strategies, for example) that are defined for Bayesian games can also be applied to games in which players may have false information.

If we add a valuation function to a game of incomplete information, we obtain a proper probabilistic Kripke model, which allows us add the machinery of dynamic epistemic logic to reason about knowledge, belief, and information change to the game.

For example, if ϕ is a sentence expressing a property of games, we can define the effect of a public announcement of a sentence ϕ by extending the usual definition from dynamic semantics in the following way:

Definition 2.2 (public announcement, Kooi (2003)) Publicly announcing ϕ in a game $(G, \rightarrow_i, p_i, g)$ results in a model $(G', \rightarrow'_i, p'_i, g)$, where:

¹In effect, we could have left out the reference to \rightarrow_i out of the model altogether. We left it in for making the relation with standard Kripke models clear.

²In the context of dynamic logic, it may also be convenient to allow for ‘inconsistent’ information states – i.e. dropping the assumption of seriality, and allow for the possibility that $p_i(g)$ assigns a value of 0 to each world. This does not mean that a rational agent should *accept* inconsistent information; instead, it is a way of representing the potential effect of getting information that is inconsistent with prior beliefs.

- $G' = \{g \in G \mid g \models \phi\}$
- \rightarrow'_i is \rightarrow_i restricted to G'
- $p'_i(g)(g') = p_i(g)(g' \mid [\phi])$, with $[\phi]$ the set of games in G where ϕ is true.³ \square

In other words, the effect of a public announcement of ϕ is to remove all worlds where ϕ is false, and to do a Bayesian update of the probability functions $p_i(g)$ with the proposition that ϕ is true.

We would like to define the effect of other kinds of updates – players getting private information, players getting false information, players suspecting that other players are being falsely informed about the lack of knowledge of a third player – along the lines of the work cited in the introduction as well. Even if most of the work in dynamic epistemic logic is about updates on non-probabilistic models, for many simple actions it is fairly straightforward to adapt the definitions.

However, there is not yet a *general* theory of epistemic actions in probabilistic models. There is some work: the thesis of Kooi (2003) contains a fully worked out logic for public announcements along the lines of the definition above; and the paper by van Benthem (2003a) sketches a more general approach by adding probability functions to the epistemic actions of Baltag et al. (1998). This latter approach, though flexible and elegant, is still sketchy. We believe this remains an important open question:

Question 2.3 Formulate a logic of epistemic actions for probabilistic models.

In the present context, we would be particularly interested in the update effects of communication acts (including private communication and lying), but also in the update effects of signaling by playing certain strategies in a multi-stage game.

Combining epistemic semantics with game theory raises other new questions for dynamic logic as well. Game-theoretical models not only introduce probabilities, but also payoffs. In some settings it is reasonable to assume that communication comes with a cost: communication will have an effect on the payoffs of the players involved. How to model this in the spirit of dynamic epistemic semantics is a completely open question:

Question 2.4 Formulate a logic of epistemic actions with costs.

2.2 Pragmatics

Given that we know the effect of an epistemic action in a game, what can we say about the expediency of, for example, sharing a particular piece of information with other players in a game? As in our example involving the three envelopes, we are interested in the question when it is advisable for a player to share some of her information with another player, and if yes, which information that should be.

As a concept, the following definition seems simple enough:

³Here, the notation $p(g \mid S)$ stands for the *conditional probability* of the game g being the actual game, given that the actual game is among those in S . It is defined in the usual way by $p(g \mid S) = p(g) / \sum_{g \in S} p(g)$.

Definition 2.5 (value of an epistemic actions) The value for an agent i of an epistemic action α game G is the difference in expected utility between G and G updated with α . \square

For example, we can say that the value for i to make a public announcement of ϕ in a game G equals the difference between his expected utility and G and that in G updated with ϕ . If the difference is positive, he might want to make the announcement, especially if he cannot find one with a higher value.

However, for the schema to work, we need to agree on a definition of ‘expected utility.’ Game theory does not provide us with a universal concept: given the wide range of games and of players that can play these games, it is impossible to predict what an arbitrary player will, or should, do in an arbitrary game. Are the players risk averse or risk seeking? Can they play mixed strategies? Are the players rational? Are these properties common knowledge between the players? Even if we know all these things, it is far from clear what ‘the’ outcome of a game will be.

Lacking a universal solution concept, we do not have a notion of expected utility that works in all cases, and therefore, no universal way of comparing epistemic actions in terms of their effect on the expected utility. Instead, we have to make do with results that apply to specific cases – results that apply to a particular type of player playing a particular type of game.

Here are some examples of questions that can be formulated in a mathematically precise way using the notions similar to those we have defined. Actually, each of these items stands for a family of questions, as the answers depend on the way the players decide on a strategy.

- Does getting information secretly (that is, without other players being aware of it) always lead to a higher expected utility? (Neyman (1991))
- Is making a true public announcement in a game always good for the social welfare of the group? Or, if not, is there always a player that values such an update? (Bassan et al. (1997))
- Is communication (of information known to be true) always undesirable in zero-sum games with two players?
- Is it always good to share the information that you have with other players in coordination games? (Crawford and Sobel (1982))
- If we want to maximize the social welfare of the agents (i.e. if we are in the position of a player whose outcome function is increasing with the sum of the utilities of the other players), should we always share the information that we have? (Kamien et al. (1990)).
- Consider the opposite case, of a player (an auctioneer, say) with an outcome function that is strictly decreasing in the utilities of the players. Is it *never* good to share information in his position?
- Which games are ‘safe for communication’ in the sense that communication does not influence the expected payoff? (A classical example is a second-price auction, in which, in contrast to a first price auction, communication between the bidders of their preferences does not influence the outcome.)
- In which games is lying discouraged? (For example, in a second price auction, no player has an incentive to misrepresent her preferences).

Many of these questions have been discussed before in the game-theoretical literature. However, we feel that it may be worthwhile to reexamine some of these questions, because we believe that from the present, more logical, stance, we can expect a different type of analysis, one that is of a qualitative rather than a quantitative character, and is focused on the syntactic structure of an epistemic action instead of on the mathematical properties of a type of signal.

Many other questions relating to the ‘pragmatics’ of epistemic actions can be asked. We hint at two more topics.

As in our three-envelope example above, lack of trust does not imply that communication is impossible: there are communication acts that are credible because the speaker has no conceivable (for the hearer) reason to lie. Talking about communication of intended actions rather than about aspects of the game situation, Farrell and Rabin (1996) propose a concept somewhat like the following: an announcement ϕ by i is *self-signaling* if it holds that i prefers ϕ to be believed iff ϕ is in fact true. We can make this a more precise by using our definitions of public announcement and the value of an epistemic action:

Definition 2.6 (credibility of an announcement) A public announcement of ϕ by i is credible to j in a game Kg iff for all games g' that j considers possible in Kg , the public announcement of ϕ is of positive value (in the sense of definition 2.5) to i in Kg' iff ϕ is true in Kg' . \square

Another related topic concerns pragmatics in the linguist’s sense of the word. From the fact that a player in a game chooses one epistemic action instead of another, the hearer can deduce extra information that may not be explicitly reflected in the semantics of the action: from the ‘conventional meaning’ of the action, together with what he knows about the game situation, the hearer may deduce the ‘speaker’s meaning’ of that action. This topic is pursued in van Rooij (2003), who works out observations of this kind into a theory of natural language pragmatics.

3 Other types of games

3.1 Strategic Information

There is a second way in which players in a game are partially informed: in general, they will not know what the other player will *do*.

This type of partial information can be modeled by Kripke model as well, one in which we associate strategy profiles instead of games to each state in the model.

Definition 3.1 (incomplete information about strategies)

A Kripke model for the strategic information of players in a strategic game $G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$ is a Kripke model $(S, (\rightarrow_i)_{i \in N}, s)$, where S is the set of strategy profiles in G .⁴ \square

The idea is that $s \rightarrow_i s'$ just in case it holds that if s is represents the actual strategies being played, as far as i knows, s' could have been played as well. So, a Kripke model for strategic information represents the information – or lack thereof – that players have about what they and the other players will do.

⁴Of course we can, as before, add probability functions to the model as well.

In general, there will be a number of constraints that could be put on such a model. For example, if i believes that j is rational, then i knows that j will not play a strictly dominated strategy, which would be reflected in the model by the fact that there is no s' with $s \rightarrow_i s'$ in which s'_j is strictly dominated by another strategy.

Such restrictions imply that we cannot simply use the existing definitions of epistemic actions to model the corresponding update – we should also take the strategic conclusions that rational players will draw on the basis of the new information into account.⁵ For example, if player i announces that he will play a certain strategy a , a successful public announcement of this intention will not only eliminate the strategy profiles where i does not play a ; under the assumption of common knowledge of rationality, the strategy profiles in which other players play a strategy that is not rational *given the information that i plays a* should be eliminated as well.

What exactly this additional effect consists of depends what is understood by ‘rational play.’ For example, if ‘common knowledge of rationality’ is equated with ‘iterated elimination of strictly dominated strategies,’ the update of a model of strategic information with a public announcement of ϕ is obtained by first applying the standard update by removing all strategy profiles that do not satisfy ϕ , and then iteratively eliminating all dominated strategies in the new game.

Question 3.2 Formulate a theory of communication acts regarding strategic aspects of a game, taking common knowledge of rationality into account.

Given such a theory, we can, as before, talk about which communication actions are expedient for certain players and which are not, and about the credibility of their announcements of their intentions.

3.2 Games in extensive form

In a game in extensive form, the sequential structure of a play is explicitly represented. Here is a definition of such a game in the format of this paper, as a Kripke model with game-theoretical objects as states:

Definition 3.3 (game in extensive form)

An game in extensive form is a tuple $(N, H, P, (u_i)_{i \in N})$ where N is a set of players, H is a set of histories (a set of sequences of actions closed under subsequences), P is a player function assigning to each history the player who is to move at that point, and u_i is utility function for each player.

An extensive game of incomplete information is a Kripke model (S, \rightarrow_i, π) , where π assigns a pair (g, h) consisting of a game in extensive form together with a history in that game to each player.

We demand that the following conditions hold:

- The accessibility relations are transitive and euclidean.
- The beliefs of the agents are modeled for each history: if $\pi(s) = (g, h)$, and h' is a history in g , then there is a s' such that $\pi(s') = (g, h')$.

⁵We should mention the work of van Benthem (2003b) in this context, who uses dynamic semantics to analyze notions of rationality in a very similar setting.

- An agent knows which moves he has available: if $\pi(g, h) \rightarrow_i \pi(g', h')$, then $\{a \mid ha \text{ in } g\} = \{a \mid h'a \text{ in } g'\}$.
- Agents have perfect recall: if $\pi(g, h) \rightarrow_i \pi(g', h')$, and $h'' \leq h$ and $h''' \leq h'$, then $\pi(g, h'') \rightarrow_i \pi(g', h''')$ □

It has become a standard move when importing epistemic semantics into games in extensive form to represent the information of each player at each history of the game, instead of only the information of the player that is to move (Bonanno (2004), van Benthem (2000)), and we have followed this example. Another deviation, made for technical rather than conceptual reasons, is that we have defined the information structure for a set of games instead of just a single one.

Just as with games of strategic information, we cannot naively apply the operations of dynamic epistemic semantics to model epistemic actions in a game. This time, we have to make sure that the assumption of perfect recall is preserved: if an agent gets new information at some point in a game, we want her to remember that she has gotten this information in all later states of the game as well.

Question 3.4 Formulate a theory of epistemic actions on games in extensive form that preserve perfect recall.

One approach to such a theory would be to define an update in an extensive game by applying the corresponding update from dynamic semantics at that history and at all histories that follow. For this to work, however, we need to formulate the game, or the update action, in such a way that an update at a certain history does not interfere with the information structure at previous or unrelated histories.

4 An example

Up to now, we have very roughly sketched a framework in which dynamic epistemic logic could be combined with game theory. To give an idea of how all of this would work for a specific case, we will give a more formal analysis of our motivating example 1.1.

Recall that the game to be played is one where Ann has a choice of three strategies (choose envelope 1, choose 2, or choose 3), but she has no idea of the way the money is distributed over the envelopes. Bill is powerless; what he gets depends only on what Ann does.

As far as Ann knows, she may be playing any of six games, corresponding to the ways that the money can be distributed over the three envelopes. Three of these games are given here in their standard matrix form, with the abbreviations we will use to refer to them:

game 036:	Bill \ Ann	choose 1	choose 2	choose 3
		0 \ 0	3 \ 3	6 \ 0

game 630:	Bill \ Ann	choose 1	choose 2	choose 3
		6 \ 0	3 \ 3	0 \ 0

game 360:	Bill \ Ann	choose 1	choose 2	choose 3
		3 \ 3	6 \ 3	0 \ 0

A probabilistic Kripke model representing the information structure in this game has six states $\{036, \dots, 630\}$, one for each of the six games. Ann's beliefs in a game g are represented by a probability functions $p_{\text{ann}}(g)$ for each game g , that assign a probability of $\frac{1}{6}$ to each of the six games. The information of Bill, who is completely informed about the game that is being played (whichever it is), can be represented by probability functions $p_{\text{bill}}(g)$ such that $p_{\text{bill}}(g)(g) = 1$, and $p_{\text{bill}}(g)(g') = 0$ for each $g' \neq g$. The actual distribution of the money is that of game 036, and we will denote the whole probabilistic model with $G036$.

We are interested in possible communication acts of Bill. Let us assume that Bill communicates using a simple propositional language in which he has all Boolean combinations of six propositional variables p_{036}, \dots, p_{630} at his disposal, with the obvious denotation. We assume that the effect of his communication acts, if successful, is given by the definition of public announcement from definition 2.2.

Two of the questions we can ask about this specific game are: what should Bill say, and should Ann believe him?

Before we can use definition 2.5 to determine the value of a specific utterance, we need a precise definition of the payoff that Bill can expect, and, as we have seen in the discussion following that definition, Bill needs to make certain assumptions about Ann's behavior.

Fortunately, predicting Ann's behavior in this game is relatively simple. Ann determines the outcome all by herself (at most, Bill can influence her decision indirectly by communicating), and she is therefore not really faced with a proper 'game-theoretical' decision, but with a standard decision problem. In particular, there is always a pure strategy for Ann that gives her a maximal expected payoff, given the information that she has, and being rational, it is such a strategy that she will choose. If there is more than one undominated strategy, we just stipulate that Ann picks one of them at random.⁶ This procedure gives us for any Kripke model Gg a unique mixed strategy $\text{strategy}(Gg)$ (i.e. a probability distribution over strategies).

Given that Ann follows the above procedure to choose her strategy, the expected payoff for Bill in a game Gg is given by:

$$\sum_{s \text{ is a strategy for Ann}} \text{strategy}(Gg)(s) \cdot \text{the payoff for Bill in } g \text{ when Ann plays } s^7$$

The expected utility of Bill in our initial game $G036$ is equal to 3.

In our informal gloss, we noted that a public announcement with the highest possible value for Bill is p_{063} , and this also follows from our formal definition. Of course, this would be lie. If Bill is limited to making true announcements only, he could opt for saying that $p_{036} \vee p_{063}$. In the model updated with the

⁶This decision is rather arbitrary; we could also allow for Ann playing mixed strategies to spread her risks also over strategies that have a small chance of being right.

⁷In our example, Bill is completely informed about the game being played, and specifically about what Ann believes, so Bill knows what Ann will do. We can generalize this to the case where Bill is not sure which game is being played or about what Ann knows exactly. There is then a distinction between the 'objective' expected utility for Bill (depending on what Ann does in the actual game) and his 'subjective' expected utility (depends on what Bill knows about the game). This 'subjective expected utility' is captured by:

$$\sum_{g' \text{ in } G} p_{\text{bill}}(g)(g') \cdot \text{the objective expected utility of Bill in } Gg'$$

Where the 'objective expected utility' is the one given in the main text.

public announcement of this sentence, Bill can, indeed, expect a payoff of $4\frac{1}{2}$, given this communication act a value of $1\frac{1}{2}$. Ann, by the way, also gains from this information; instead of getting 1, she may expect a payoff of 2.

However, Bill, if devious, can also *mislead* Ann by telling the truth. Suppose he tells her that $\neg p_{360} \wedge \neg p_{630}$. In the model that results after a public announcement of this sentence, Ann has four possibilities left, each with a probability of $\frac{1}{4}$. Two of these possibilities have her win 3 yuan if she chooses the third envelope, the other two have 3 yuan for her in the first and the second envelope, respectively. (The latter is the actual situation.) So, choosing the third envelope is the strategy that gives her the highest expected value. In fact, she will get nothing in this case – and Bill gets 6. So, for Bill, this announcement has a value of 3 yuan. This is a case where Ann is better off without having this information, even if it is true.⁸

The notion of public announcement of definition 2.2 presumes that all information that is announced is believed to be true without reservation. The case where Ann is suspicious of Bill’s announcement (after all, they are playing a non-cooperative game) is interesting as well: when should Ann believe what Bill says? In our informal gloss of the example we came to the conclusion that announcing p_{036} would not be credible to Ann, but saying that $p_{036} \vee p_{630}$ is.

This is corroborated by the notion of credibility of definition 2.6. Announcing $p_{036} \vee p_{630}$ is credible because in all games Gg' that are consistent with Ann’s information, a successful announcement that the first envelope is empty is valuable to Bill only when it is true. An announcement of p_{036} is, although true, not credible, because as far as Ann knows, the actual game being played could be $G063$, in which case Bill could double his expected payoff by lying that p_{036} . Finally, Bill’s clever announcement of $\neg p_{360} \wedge \neg p_{630}$ is not credible for Ann either, but for a different reason: the announcement would have been true in the game $G063$ as well, but has no positive value for Bill in this case.

There are many other aspects of this example that we have not discussed. In particular, we have not taken account the role that common knowledge plays in the reasoning of the players, and in particular, we have not discussed the fact that Ann may also use the information she has about Bill to judge the meaning of his utterance. Unfortunately, reasons of space force us to cut off the discussion here.

5 Conclusions

We have sketched a formal framework in which dynamic epistemic semantics can be used to reason about epistemic actions in games, by rearranging the definitions of, respectively, strategic game with incomplete information about the structure of the game, strategic games with incomplete information about the actions of the players, and of a game in extensive form, and making them look like Kripke models.

This being a paper for workshop, we have mostly focused on open questions, and there are many.

⁸This partially answers the first of the list of questions that we asked in the pragmatics section: is getting information in private always good? The answer seems to be negative, but note that this conclusion is built on very specific assumptions about what it means to get information and how you should change your strategy on the basis of this new information.

First of all, we identified a number of types of dynamic semantics that have hardly or not at all been studied: the theory of epistemic actions on probabilistic Kripke models, the theory of epistemic actions that come with a cost, theories of ‘strategic updating’ that take different forms of rationality into account, and the definition of epistemic updates in a game in extensive form. Another type of question that arises is more game-theoretical in nature; questions about the value of certain types of information in certain classes of games, and about the information strategies that players of that game may pursue to obtain a preferred outcome. Many of these questions are interesting, and few have been answered.

References

- Baltag, A., L. Moss, and S. Solecki: 1998, ‘The Logic of Public Announcements, Common Knowledge and Private Suspicions’. In: *Proceedings TARK*. pp. 43–56.
- Baltag, A. and L. S. Moss: 2004, ‘Logics for Epistemic Programs’. *Synthese* **139**(2), 165–224.
- Bassan, B., M. Scarsini, and S. Zamir: 1997, “‘I don’t want to know’: can it be rational?”. Center for Rationality and Interactive Decision Theory, The Hebrew University of Jerusalem, Discussion paper 158.
- Bonanno, G.: 2004, ‘Memory implies von Neumann-Morgenstern games’. *Knowledge Rationality and Action*.
- Crawford, V. and J. Sobel: 1982, ‘Strategic Information Transmission’. *Econometrica* **50**, 1431–1451.
- Fagin, R. and J. Y. Halpern: 1994, ‘Reasoning about knowledge and probability’. *Journal of the ACM* **41**(2), 340–367.
- Farrell, J. and M. Rabin: 1996, ‘Cheap Talk’. *The Journal of Economic Perspectives* **10**(3), 103–118.
- Gerbrandy, J.: 1998, ‘Bisimulations on Planet Kripke’. Ph.D. thesis, Universiteit van Amsterdam. ILLC Dissertation Series DS-1999-01.
- Kamien, M. I., Y. Tauman, and S. Zamir: 1990, ‘On the value of information in strategic conflict’. *Games and Economic Behavior* **2**, 129–153.
- Kooi, B. P.: 2003, ‘Probabilistic Dynamic Epistemic Logic’. *Journal of Logic, Language and Information* **12**(4), 381–408.
- Neyman, A.: 1991, ‘The positive value of information’. *Games and Economic Behavior* **3**, 350–355.
- Osborne, M. and A. Rubinstein: 1998, *A Course in Game Theory*. The MIT Press.
- van Benthem, J.: 2000, ‘Games in Dynamic-Epistemic Logic’. *Bulletin of Economic Research* **53**(4), 219–248. Proceedings LOFT-4, Torino.
- van Benthem, J.: 2003a, ‘Conditional Probability Meets Update Logic’. *Journal of Logic, Language and Information* **12**(4), 409 – 421.
- van Benthem, J.: 2003b, ‘Rational Dynamics and Epistemic Logic in Games’. ILLC Research Report P-2003-07.
- van Ditmarsch, H.: 2000, ‘Knowledge games’. Ph.D. thesis, Groningen University. ILLC Dissertation Series 2000-06.
- van Rooij, R.: 2003, ‘Quality and Quantity of Information Exchange’. *Journal of Logic, Language and Information* **12**, 423–451.