Communication Strategies in Games

Jelle Gerbrandy

Dipartimento di Informatica
Università di Torino
jelle@gerbrandy.com

ABSTRACT. *We formulate a formal framework in which we combine the theory of dynamic epistemic logic and the theory of games. In particular, we show how we can use tools of dynamic epistemic logic to reason about information change – and in particular, the effect of communication acts – in such a game of imperfect information. We show how this framework allows for the formulation of specific assumptions in pragmatics of communication, as well as the formulation of general results about the value of information.*

KEYWORDS: *Dynamic Epistemic Logic, Game Theory, Pragmatics, Value of Information*

1. Introduction

In recent years, dynamic epistemic logic (DEL) has become a mature field of research. It provides a detailed analysis of *epistemic actions*: different types of information change that may occur in a multi-agent setting ([BAL 98], [GER 98], [DIT 00], [BAL 02], [BAL 04]). In this paper, we will sketch some ways in which these tools can be put to work in what is perhaps the most commonly used model of multi-agent interaction: the theory of games.

The idea is this. Players of a game are often not completely informed about all aspects of the strategic situation they are in. They may be under-informed about the exact parameters of the game, such as their payoffs or the payoffs of the others, or they may not be sure what the other players know about the game. This lack of information is usually modeled by combining a game-theoretical structure with a representation of the information of the players. This information structure may change as a result of, for example, observations by or communication between players. Dynamic epistemic logic provides a precise way of modeling various types of information change of this kind, and is therefore a natural framework for studying its effects in a game-theoretic context.

Combining dynamic epistemic logic with game theory is a knife that cuts both ways. On the one hand, it *provides the theory of epistemic actions with a pragmatics*:

it allows us to talk about, for example, an epistemic action being useful to an agent because it improves her expected payoff; of a communication act being 'credible' because the speaker has no conceivable reason for lying about what he says; and it provides criteria for judging when to choose one communication act rather than another. 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 — literature on signaling games, for example, or about the value of information — but it seems to miss the versatile approach to multi-agent information change that is characteristic of the theories of dynamic semantics.

Some of the issues that arise can be illustrated with the following example.

EXAMPLE 1 (THREE ENVELOPES). — Ann has to choose one of three closed envelopes. One envelope is empty, one envelope contains three euro for Ann and three euro for Bill, and a third envelope contains 6 euro for Bill (Ann gets nothing). Ann has no idea which envelope contains which amount of money, but Bill knows the exact distribution. All of these facts are commonly known.

Since Ann does not know which of the three envelopes has 3 euro for her, she can do no better than to choose an envelope at random. She may then expect a payoff of 1 euro, 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 euro 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 euro. 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. If Ann is gullible, Bill's best option is to lie, and tell Ann that there is 3 euro for her in the third envelope. If Ann believes him, she will choose this envelope, giving Bill 6 (and Ann nothing).

So Ann is suspicious of what Bill says. She does not believe Bill when he says that the third envelope contains 3 euro 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. □

2. Semantics

In this section, we will give an introduction to dynamic epistemic logic, and in particular its probabilistic version, before moving on to game theory proper.

2.1. ***Probabilistic Dynamic Epistemic Semantics***

Typically, in the game theoretical literature, information is taken to be probabilistic in nature. For the purposes of this article, it will suffice to just add a probability measure to the classical notion of a Kripke model. For more versatile combinations of probabilistic and epistemic reasoning, see e.g. [FAG 94] and [HAL 03].

DEFINITION 2 (PROBABILISTIC KRIPKE MODEL). —
Given a set of players N *and a set of propositional variables* Prop*, a* probabilistic Kripke model *is a tuple* $(S, (\longrightarrow_i)_{i\in N}, V, P)$*, where* S *is a finite set of states,* $\longrightarrow_i$ *is a relation on* $S \times S$ *that is Euclidean and transitive,* V *is a valuation function that assigns truth values to variables from* Prop *at each state, and* P *is a prior probability distribution over* S*.*

If s *is a state in a Kripke model* M*, we will refer to the pair* (M, s) *as a* pointed model*, and we will write* Ms *instead of* (M, s)*. Also, we will write* $p_i(s)$ *for the probability distribution conditional on the information set of* i *in* s*, defined as* $p_i(s)(s') = P(s' \mid \{t \mid s \longrightarrow_i t\})$*.*[1] □

A probabilistic Kripke model is simply a model of $K45$ epistemic logic, together with a probability distribution over its domain of states. The accessibility relations $\longrightarrow_i$ represent the information of the players relative to a state s: the set $\{s' \mid s \longrightarrow_i s'\}$ of i-accessible states is the *information set of* i *in* s: in state s, i believes that the actual state must be among those in her information set.

In this paper we will be concerned with model theory rather than with logic, and we will leave the exact nature of the object language in which to express properties of the models implicit. It suffices to assume that this language has a number of propositional variables to express properties of states, and standard epistemic operators K_i with their usual semantics:

$$M, s \models K_i\phi \text{ iff } M, s' \models \phi \text{ for each } s' \text{ such that } s \longrightarrow_i s'$$

Dynamic epistemic semantics is about different ways in which information in Kripke models may change. It defines the effect of 'epistemic actions' as transformations on Kripke models. A typical example is the definition of a 'public announcement' from [PLA 89], which is generalized to the probabilistic case by [KOO 03]. A public announcement reflects an idealized situation in which all players get the information that a certain proposition is true from a source that is completely trusted (and commonly known to be so). Formally, we can define the effect of such a public announcement of ϕ in a multi-agent Kripke model M by simply removing all states where ϕ is false. If $[\phi]_M$ is the set of states in M where ϕ is true, we adapt the probability measure by Bayesian conditioning on ϕ and setting the new probability measure $P^{\mathsf{new}}(s) = P^{\mathsf{old}}(s \mid [\phi]_M)$

---

1. Here, the notation $p(s \mid X)$ stands for the *conditional probability* of the state s being the actual one, given that the actual state is among those in X. It is defined in the usual way as $p(s \mid X) = p(s)/\Sigma_{s\in X}\, p(s)$.

But information change can be much more complex, and cannot always be modeled by simply removing states from the model. Players may get private information, or information that is false, players may suspect that other players are being falsely informed about the lack of knowledge of a third player, etc. In a series of papers ([BAL 98], [BAL 02], [BAL 04]) an elegant framework for representing different types of 'epistemic actions' in a uniform way has been developed. The crucial idea is that partial information about actions can be modeled in the same way that information about states is modeled in a Kripke model: by way of relations between possible actions. For example, the fact than an agent j is not sure whether i has learned that ϕ is true can be modeled as an action in which j cannot distinguish the action where i learns ϕ from the trivial action in which nothing happens.

DEFINITION 3 (ACTION MODELS). —
An epistemic action model *is a tuple* $(A, (\longrightarrow_i)_{i \in N}, \mathsf{pre})$*, where* A *is a set of actions,* $\longrightarrow_i$ *is a relation over* A *for each* i*, and* pre *is a function from actions in* A *to* preconditions *of a given language* $\mathcal{L}$*.*

A pointed epistemic action model *is a tuple* $((A, (\longrightarrow_i)_{i \in N}, \mathsf{pre}), \Sigma)$ *consisting of an action model and a subset* Σ *of* A*.* □

In an action model $(A, (\longrightarrow_i)_{i \in N}, \mathsf{pre})$, the relations $\longrightarrow_i$ represent the informational structure of the action: if $a \longrightarrow_i a'$, this means that in the case that a is the 'real' action, as far as i knows, it might as well be a' that is happening. In addition, each action comes with a *precondition* that limits the set of states in which the action can be executed. In a pointed model (A, Σ), the set Σ contains those actions that might *actually* be happening.

Most of the work in dynamic epistemic semantics is developed in a non-probabilistic setting. Generalizing these definitions to apply to probabilistic models is not trivial. The problem is that epistemic actions can be non-deterministic – an agent may not be sure which action is executed in a given state, and has therefore no way of calculating her posterior probabilities. There are two possible solutions to this. One approach is to add explicit information about the probabilities with which an action may happen to the action models. This leads to a theory of *probabilistic action models*, proposals for which can be found in [BEN 03], [AUC 05] and in [BEN 06]. The approach we take in this paper is more simplistic. We will consider only action models for which the problem does not arise: those action models in which there can, for the agents, never be any ambiguity about which action takes place in a certain state.

DEFINITION 4 (EPISTEMICALLY DETERMINISTIC ACTION MODELS AND PROBABILISTIC UPDATES). — *An action model* (A, Σ) *is* epistemically deterministic *iff for any two actions* a, b *such that* $a \neq b$*, if* a *and* b *are both in* Σ*, or if they are in the same information set, then* $\mathsf{pre}(a) \wedge \mathsf{pre}(b)$ *is a contradiction.*

Let $M = (S, (\longrightarrow_i)_{i \in N}, V, (P_i)_{i \in N})$ *be a Kripke model and let* A *be an action model* $(A, (\longrightarrow_i)_{i \in N}, \mathsf{pre})$*. The* product update $M \times A = (S', (\longrightarrow'_i)_{i \in N}, V', P')$ *is defined by setting:*

- $S' = \{(s, a) \mid (s, a) \in S \times A$ *and* $M, s \models \mathsf{pre}(a)\}$

- $(s, a) \longrightarrow_i' (s', a')$ *iff* $s \longrightarrow_i s'$ *and* $a \longrightarrow_i a'$
- $V'((s, a)) = V(s)$.
- $P'((s, a)) = \dfrac{P(s)}{\Sigma_{(s',a') \in S'} P(s')}$ □

If M, s is a pointed Kripke model and (A, Σ) a deterministic pointed action model and $s \models \mathsf{pre}(a)$ for some $a \in \Sigma$, then we will write $M, s \times (A, \Sigma)$ for the (unique) model $M \times A, (s, a)$. □

Restricting the framework to action models that are deterministic limits the repertoire of epistemic actions somewhat. We cannot model, for example, the fact that one agents suspects, but is not sure, that something has happened. The remaining fragment remains quite rich, however, and is more than sufficient for the scope of this article.

– A *public announcement* of ϕ can be represented as an action model that has a single action a with precondition $\mathsf{pre}(a) = \phi$, and with $a \longrightarrow_i a$ for each player i. Its product update gives us exactly the effect as above: it eliminates the non-ϕ-states from the model, and the new probability measure is obtained by conditioning on the prior one. The *trivial action* is a public announcement of the tautology $\top$: the model that results after a product update is isomorphic to the original one.

– A *private announcement* models the case when a player i gets the information that ϕ, while the other players are not aware of anything – as far as they know, nothing happened. A model for a private announcement of ϕ to a player i has two actions, a_ϕ and $a_\top$: intuitively, in a_ϕ, player i privately learns that ϕ, while $a_\top$ is the trivial action. The preconditions are given by $\mathsf{pre}(a_\phi) = \phi$ and $\mathsf{pre}(a_\top) = \top$. Player i knows what is going on: $\longrightarrow_i$ is the identity relation. Other players, however, mistakenly believe that nothing happens, so $a \longrightarrow_j b$ only if $b = a_\top$ for those players.

In the kind of situations that we consider below, it is often interesting to look not at the case in which an agents obtains a specific piece of information, but rather as receiving a certain *type* of information – the type of information that is known as a *signal* in game theory. Within the present context, it is perhaps more natural to think of this type of update as getting the answer to a question that may get one of a possible set of mutually exclusive answers, which is why we will call these type of action models 'answer sets.'

– A *public answer set* is a non-deterministic epistemic action (A, Σ) which consists of a set of public announcements whose set of preconditions $\{\mathsf{pre}(a) \mid a \in \Sigma\}$ forms a partition of the logical space. Intuitively, these are the 'answers' in the answer set.

– In a *legal private answer set*, an of agent i gets an answer to a question, and it is common knowledge among all players that she does. In other words, the other agents know that i gets the answer – they just don't know what the answer is. Formally, this corresponds to an epistemic action (A, Σ), where Σ is the domain of A, the preconditions of Σ form a partition, $\longrightarrow_i$ is the identity relation, and $\longrightarrow_j$ is the universal relation for each j different from i.

– In an *illegal private answer set* an agent i gets an answer to a question without the other agents being aware of anything. This can be modeled with an epistemic action that consists of a set of private announcements whose preconditions form a partition.

2.2. *Dynamic Epistemic Logic for Games*

One way of modeling partial information about a game is by associating possible games with each state in a Kripke model. Following standard practice in the game-theoretical literature, we will call such games *Bayesian games*.

DEFINITION 5 (BAYESIAN GAMES). —
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 payoff function for* i *that assigns to each element of* $\times_i S_i$ *a real number.*

A Bayesian game *is a tuple* $(S, G, (\longrightarrow_i)_{i\in N}, V, P, s)$*, where* $(S, (\longrightarrow_i)_{i\in N}, V, P, s)$ *is a probabilistic pointed Kripke model, and* G *assigns a game in strategic form with players* N *to each state in* S *in such a way that each player has the same strategies available in all games.*

A (mixed) strategy *for an agent* i *in a game in strategic form is a probability distribution over* S_i*. A* strategy for a Bayesian game *for player* i *is a function* σ_i *that assigns to each state* s *a (possibly mixed) strategy, in such a way if* $s \longrightarrow_i s'$*, then* $\sigma_i(s) = \sigma_i(s')$*. A* strategy profile $(\sigma_i)_{i\in N}$ *contains a strategy for each player in* N*.*
□

We will abuse notation somewhat, and identify the states s in a game with the game $G(s)$. For example, we might write Gg for the probabilistic model Ms with $G(s) = g$. Intuitively, in a game Gg, g represents the payoff matrix of the game that is actually being played; and $g \longrightarrow_i g'$ means that i believes that it is possible that the actual game being played has the payoff defined by g'.

The present definition of a Bayesian game is a generalization the notion of the commonly used definition of a game of incomplete information (see, for example, [OSB 98]). The informational structures used in game theory are almost always assumed to be 'Harsanyi-consistent': the accessibility relations are equivalence relations, and the probability measures are based on a common prior. We generalize this first assumption: it is possible for the players to have *false* beliefs as well. This generalization is crucial for studying communication in a general way. It is obvious that we need to allow for false beliefs to model the effect of a successful lie, but also more innocent-seeming communication acts might lead to models in which information is false. For example, the result of a private announcement is, in general, a model in which the information of the other player becomes false (because she still believes, false, that the information set of the receiver has not changed).

At this point, it is important to remark that, at least from a technical point of view, a great part of the game-theoretical analysis of games of incomplete information

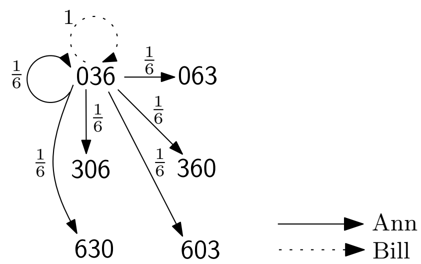

Figure 1. *The initial situation of example 1 (we only drew the arrows leaving* 036*)*

does not depend on Harsanyi-consistency. In particular, many of the standard solution concepts (such as Nash equilibrium, rationalizability, for example) continue to work.

EXAMPLE 6 (EXAMPLE 1 CONTINUED). — 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:

Bill\ Ann	choose 1	choose 2	choose 3
game 036:	$0 \setminus 0$	$3 \setminus 3$	$6 \setminus 0$
game 630:	$6 \setminus 0$	$3 \setminus 3$	$0 \setminus 0$
game 360:	$3 \setminus 3$	$6 \setminus 3$	$0 \setminus 0$

A probabilistic Kripke model representing the information structure in this game has six states $\{036, \ldots, 630\}$, one for each of the six games. Ann's beliefs are represented by the universal relation over these states, while Bill's beliefs are represented by the identity relation. Ann assigns equal probability to each of the games: $P(g) = \frac{1}{6}$.

The actual distribution of the money is that of game 036, and we will denote the whole probabilistic model with G036. Figure 1 is a graphical representation of a part of the model, and figure 2 shows an example of the effect of an epistemic action in this model.

□

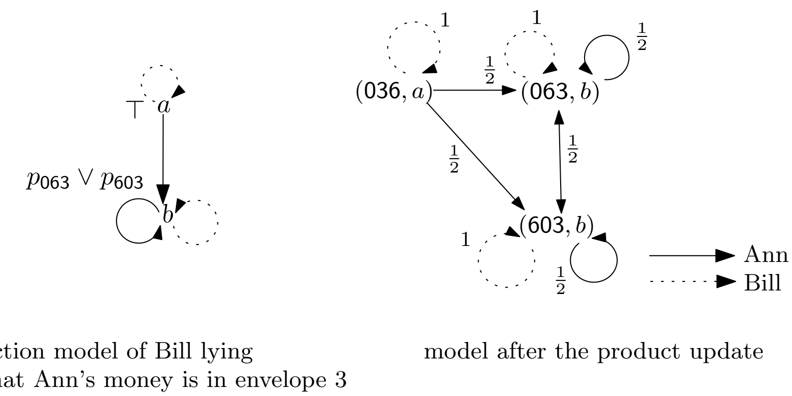

Figure 2. *Bill lies that there is money for Ann in the third envelope*

3. Pragmatics

Dynamic epistemic logic defines a range of operations that correspond to specific changes in the informational structure of a game. Some of these actions may be 'good' for a certain player in the sense that they change the informational structure of a game in such a way that he can expect a higher payoff in the new game. Sometimes it is possible to quantify this quality: for example, if a player i believes that ϕ is true, we can say that the value for i of a public announcement of ϕ in a game Gg equals the difference between the payoff that he expects in Gg and that in the game after the update. In other words, a player can 'measure' the value of an epistemic action by comparing his expected value before and after the update with that action. Of course, this makes the value of an epistemic action dependent on some notion of expected value.

3.1. *Expected Value*

Defining the expected value of a given strategy profile for an agent is straightforward: the agent simply multiplies the payoff she gets when this strategy is played in each of the states in her information set with the probability that this state is the actual one.

DEFINITION 7 (EXPECTED VALUE OF A STRATEGY PROFILE). — *The* expected value *for i in Gg of a strategy profile σ, $\mathsf{V}_i(Gg,\sigma)$, is the weighed average of her payoff for that profile in each of the games in her information set:*

$$\mathsf{V}_i(Gg,\sigma) := \Sigma_{g'\in S}(p_i(g)(g') \cdot u_i^{g'}(\sigma))$$

This definition is only useful if we somehow know what strategies the players will choose. Given the wide range of games and the fact that we know little about the players of these games, it is impossible to predict in a general way what an arbitrary player will, or should, do. Are the players risk averse or risk seeking? Can they play mixed strategies? Are these properties common knowledge between the players? Even if we know all these things, it is often far from clear what 'the' outcome of a game will be.

However, with certain specific assumptions about the players or about the type of games, it is sometimes possible to associate a specific expected value to a game. For example, a risk averse player that has no confidence in the rationality of his opponents may take his security level – which is the highest expected payoff he can obtain given a worst-case scenario for the strategy choices of the other players – as the value of a game. Similarly, if the rationality of the players is common knowledge, a risk-loving player may identify the value of a game with, say, the highest value among his payoffs in, say, each rationalizable strategy profile. Alternatively, we restrict our attention to classes of games that can reasonably be said to have a certain fixed value. Zero-sum games with two players, for example, have a unique Nash Equilibrium payoff, and [LEH 06] take this payoff to be the value of such games. In a similar move [BAS 03] consider games which have a unique Pareto dominant strategy profile: we might take the associated payoff of that profile as the value.

EXAMPLE 8 (EXAMPLE 1 CONTINUED). — In our example, predicting the outcome of the game is relatively easy. Ann determines the outcome all by herself (at most, Bill can influence her decision indirectly by communicating), and she is therefore not faced with a proper 'game-theoretical' decision but with the relatively easy problem of making a decision under uncertainty instead. In short, we can fairly safely assume that if Ann is rational, she chooses that envelope that maximizes her expected payoff. We also assume that if there are several envelopes that maximize this value, she chooses randomly between them. Given a game Gg, let $\sigma(Gg)$ be Ann's optimal strategy in this sense.

If Bill knows this about Ann, and given the fact that he is completely informed about the game, Bill's expected payoff is simply the actual payoff he gets when Ann plays her strategy:

$$\mathsf{V}_{\mathsf{bill}}(Gg, \sigma) = u^g_{\mathsf{bill}}(\sigma(Gg))$$

In the initial situation of our example, Ann has no information about the distribution of money over the envelopes at all – all strategies have the same expected payoff. That means that she chooses randomly between envelopes, and Bill's expected payoff $\mathsf{V}_{\mathsf{bill}}(G\mathsf{036}) = 3$. □

3.2. *Value of Epistemic Actions*

Given a way of associating a value to specific games, we can calculate the expected value of an epistemic action (A, Σ) by calculating, for each action $a \in \Sigma$, the increase in expected value, and multiplying that value with the probability of a actually occurring.

DEFINITION 9 (EXPECTED VALUE OF AN EPISTEMIC ACTIONS). — *Let* V_i *be a function that assigns payoff values to pointed games:* $\mathsf{V}_i(Gg)$ *represents the value of the game* Gg *to* i*. Let* (A, Σ) *be a deterministic action model.*

The expected value for player i of action (A, Σ) relative to V_i in Gg *is the increase expected value for each of the actions in* Σ*, weighed by the probability that it will occur:*

$$\Sigma_{a\in\Sigma}\Sigma_{g'\models\mathsf{pre}(a)}(p_i(g)(g')\cdot\mathsf{V}_i(Gg'\times(A,\{a\}))-\mathsf{V}_i(Gg)$$

□

When we consider the value of public announcement of ϕ to an agent that already believes ϕ (and hence is in the position of announcing it), the definition of expected value can be simplified considerably. If $(A, \{a\})$ is a public announcement of ϕ, and i believes ϕ in Gg, then the expected value becomes simply the difference in expected value before and after the update:

$$\mathsf{V}_i(Gg\times(A,\{a\}))-\mathsf{V}_i(Gg)$$

For measuring the expected value of an answer set, definition 9 gives us the sum of weighed values of each of the answers. Formally, let (A, Σ) be the answer set with preconditions $\{\phi_1 \ldots \phi_n\}$. Let A_ϕ be action model that represents a public announcement of ϕ. If we write $p_i(g)(\phi)$ for the probability of ϕ given the information set of i in G, then the expected value of the answer set becomes the sum of the probability of each of the answers multiplied by the value of that answer:

$$\Sigma_{1\leq m\leq n}p_i(g)(\phi_m)\cdot[(\mathsf{V}_i(Gg\times A_{\phi_m})-\mathsf{V}_i(Gg))]$$

3.3. *Value and Credibility*

With precise assumptions about expected value, we can put precise values on epistemic actions using definition 9 and compare them in terms of expected value. It seems obvious that in choosing a communication act, a speaker would choose an act that maximizes her expected value – in whichever way her notion of expected value is defined.

EXAMPLE 10 (EXAMPLE 1 CONTINUED). — In our informal gloss of the three envelope game, we noted that an action that manages to convince Ann that p_{063} – a sentence that expresses that the game being played is 063 – is valuable for Bill. If we model the effect of a 'successful lie' in the way represented in figure 2, we can apply definition 9 to calculate its value for Bill. After this action, Ann believes (falsely) that choosing the third envelope will give her a payoff of 3, and will choose this envelope: Bill gets 6 instead of the initial expected payoff of 3, giving the action a value of 3 for Bill.

If Bill can, for whatever reason, only make announcements with sentences that he knows to be true, he could opt for saying that $p_{036} \vee p_{063}$. In the model updated with the public announcement of this sentence, Bill expects a payoff of $4\frac{1}{2}$, which gives this communication act a value of $1\frac{1}{2}$.

However, Bill, can do better also when speaking the truth. He can also *mislead Ann by telling the truth*: he can tell 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 if she chooses the third envelope, the other two have 3 euro for her in the first and the second envelope, respectively (the latter is the actual situation). Choosing the third envelope is the strategy that gives her the highest subjective expected value, even if in fact she will get nothing in this case, and Bill gets 6. (Note that in this situation, it is actually harmful for Ann to get information that is true.) □

This analysis of our example raises an obvious question: when should a hearer believe what the speaker says? Again, an answer to this question depends on assumptions that are exogenous to the model. For example, if players can, for whatever reason, only say the truth, and know this, the question is moot. But even if players are at complete liberty to choose their acts, lack of trust does not imply that communication is impossible – this is well-known from the literature on signaling games. There are communication acts that are credible because of the fact that the speaker has no conceivable reason to lie. [FAR 96], who write about communication of intended actions rather than of propositions that describe the game situation, 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 adapt this definition to apply to communication about facts, using our definitions of public announcement and the value of an epistemic action:

DEFINITION 11 (CREDIBILITY OF AN ANNOUNCEMENT). — *A public announcement of ϕ by i is credible to j in a game Gg iff it holds for all games g' that j considers possible in Gg that the public announcement of ϕ has a positive value for i in Gg' iff ϕ is true in Gg'* □

EXAMPLE 12 (EXAMPLE 1 CONTINUED). — When should Ann believe what Bill says? If we use the notion of credibility of definition 11, we get the following results.

Announcing $p_{036} \vee p_{063}$ is credible because in all games Gg' that are consistent with Ann's information, a public announcement that the first envelope is empty has positive value for 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. □

We refer the reader to [ROO 03] for similar and further investigations.

4. Actions with a positive value

General results about the positive value of information are hard to come by. For almost any kind of increase in information of the players it is possible to think of a scenario in which the players are actually worse of with the new information ([BAS 97]) — the case in our example of the true sentence $\neg p_{360} \wedge \neg p_{630}$ that can be used to have Ann make a 'wrong' choice is an example. However, there are some positive general results about the value of information as well. We have taken three results from the literature and rephrased them as statements about particular classes of epistemic actions that we introduced in section 2.1. The first result, based on a theorem from [NEY 91], says that strictly private information – such as captured in a illegal private answer set — always has a positive value. The second result says that in a zero-sum game, getting private information has a positive value even if when the other players are aware of the first player getting the information – this is the kind of update that occurs with a 'legal answer set'. The third is a result says that says that that getting true information always has a positive value if the resulting game has a unique Pareto dominant strategy profile ([BAS 03]).

As we discussed above, the value of an epistemic action is only defined relative to a function that captures the expected value of a game. For the purposes of this paper, we will concentrate on two of such functions: the function that identifies the value of a game with its security level (for very cautious players), and the function that identifies the value of the game with the smallest payoff to i in any of the Nash Equilibria of the game. For reasons of space, we refer to the original papers for more detailed proofs.

THEOREM 13 (PRIVATE INFORMATION HAS POSITIVE VALUE). — *Let (A, Σ) be an illegal private answer set for i in the sense above – i gets purely private information without the other players suspecting anything. Then (A, Σ) has a non-negative value for i with respect to the security level, or with respect to the minimum equilibrium payoff.*

proof sketch: The idea behind the proof is straightforward. Player i can get the same payoff in the new game as in the initial game by just ignoring his new information. Since the other players suspect nothing, they have no reason to change their strategies in response. So i's expected value – in terms of his security level or Nash equilibria, can only improve. [NEY 91] gives a detailed proof in terms of Nash equilibria that can be reformulated to apply here. □

One class of games that have a natural 'value' are two-person zero-sum games, which have a unique equilibrium payoff. Let us say that a Bayesian game is a zero-sum game when each of the games of perfect information associated with its states are zero-sum. If such a game is Harsanyi-consistent (in particular, if the accessibility relations are equivalence relations) the entire Bayesian game can be viewed as a zero-sum game as well, and therefore has a unique equilibrium value. Variants of the following theorem, that says that getting a question answered is always positive even when your opponent knows you are getting it answered (but does not know *which* answer you get) have been proven in [KAM 90], [LEH 06], [GOS 01].

THEOREM 14. — *(positive value of legal answer sets in zero-sum games) Let Gg be an Harsanyi-consistent zero-sum game. Let (A, Σ) be a legal private answer set for i. Then (A, Σ) has a positive value for i in Gg with respect to the value function V_i that assigns to each game its unique Nash Equilibrium value.* □

A third kind of general theorem is inspired by [BAS 03]. Again, we consider a class of games in which the expected outcome is fairly clear: those that have a unique Pareto dominant strategy profile. The result states applies to a large class of epistemic actions, and says that all players are better off with any kind of truthful epistemic action, as long as the resulting game has a unique dominant equilibrium.

Having a unique Pareto dominant strategy profile seems a very strong assumption. But for a player who wants to influence the outcome of a game by changing its informational structure, it makes sense to try to turn the game into one with a very predictable outcome (such as those with a unique Pareto dominant strategy profile).[2]

THEOREM 15. —
Let (A, Σ) be an legal answer set – be it private or public. Assume that $Gg \times (A, \Sigma)$ has unique dominant equilibrium. Let V_i be a value function that assigns some Nash equilibrium value to a game (and so it will give us the value for i of the Pareto dominant strategy profile, if a unique one exists). Then A has positive value with respect to V_i in Gg.

proof sketch: Again, the proof hinges on the observation that players can ignore the new information and duplicate their strategies in the new model $Gg \times (A, \Sigma)$. If a strategy profile σ is a Nash equilibrium in G and τ its duplicate $G \times A$, then the expected value in Gg of playing τ in $Gg \times (A, \Sigma)$ is exactly the expected value of σ in Gg. Since the unique dominant equilibrium in $Gg \times (A, \Sigma)$ dominates τ in particular, it will give i a payoff that is at least as high. □

5. Conclusions

We have defined a formal framework in which dynamic epistemic logic can be put to use to reason about information change in Bayesian games. We illustrated how the definitions can be used to formulate and answer two types of problems. One type of question concerns the pragmatics of communication – and in particular the value of different types of communication in a game-theoretic situation. We discussed the

2. Similar considerations lead [KAM 90] to define the property of 'being able to induce a certain value' with a communication act only if the resulting game has a unique Nash equilibrium.

rudiments of such a theory in section 3. Another type of question revolves around finding classes of games (and solution concepts) in which certain classes of epistemic updates behave in a particular way – we gave examples of such results in section 4. With this paper, we hoped to have defined the relevant notions in enough detail to provide the basis for future work in which more of these questions can be formulated and, hopefully, answered.

The author would like to thank the Fondazione CRT for financial support, and the participants of the workshop "Belief Revision and Dynamic Logic" at the ESSLLI'05, as well as two anonymous referees, for their useful comments.

6. References

[AUC 05] AUCHER G., "How Our Beliefs Contribute To Interpret Actions", PECHOUCEK M., PETTA P., VARGA L. Z., Eds., *CEEMAS 2005, LNAI 3690*, p. 276–286, 2005.

[BAL 98] BALTAG A., MOSS L., SOLECKI S., "The Logic of Public Announcements, Common Knowledge and Private Suspicions", *Proceedings TARK*, 1998, p. 43–56.

[BAL 02] BALTAG A., "A logic for suspicious players: epistemic actions and belief-updates in games", *Bulletin of Economic Research*, vol. 54, num. 1, 2002, p. 0307–3378.

[BAL 04] BALTAG A., MOSS L. S., "Logics for Epistemic Programs", *Synthese*, vol. 139, num. 2, 2004, p. 165-224.

[BAS 97] BASSAN B., SCARSINI M., ZAMIR S., ""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, 1997.

[BAS 03] BASSAN B., GOSSNER O., SCARSINI M., ZAMIR S., "Positive value of information in games", *International Journal of Game Theory*, vol. 32, 2003, p. 17–31.

[BEN 03] VAN BENTHEM J., "Conditional Probability Meets Update Logic", *Journal of Logic, Language and Information*, vol. 12, num. 4, 2003, p. 409 – 421.

[BEN 06] VAN BENTHEM J., GERBRANDY J., KOOI B., "Dynamic Update with Probabilities", ILLC Prepublication PP-2006-21, 2006.

[DIT 00] VAN DITMARSCH H., "Knowledge games", PhD thesis, Groningen University, 2000, ILLC Dissertation Series 2000-06.

[FAG 94] FAGIN R., HALPERN J. Y., "Reasoning about knowledge and probability", *Journal of the ACM*, vol. 41, num. 2, 1994, p. 340–367.

[FAR 96] FARRELL J., RABIN M., "Cheap Talk", *The Journal of Economic Perspectives*, vol. 10, num. 3, 1996, p. 103–118.

[GER 98] GERBRANDY J., "Bisimulations on Planet Kripke", PhD thesis, Universiteit van Amsterdam, 1998, ILLC Dissertation Series DS-1999-01.

[GOS 01] GOSSNER O., "The value of information in zero-sum games", unpublished, 2001.

[HAL 03] HALPERN J. Y., *Reasoning About Uncertainty*, MIT Press, 2003.

[KAM 90] KAMIEN M. I., TAUMAN Y., ZAMIR S., “On the value of information in strategic conflict”, *Games and Economic Behavior*, vol. 2, 1990, p. 129–153.

[KOO 03] KOOI B. P., “Probabilistic Dynamic Epistemic Logic”, *Journal of Logic, Language and Information*, vol. 12, num. 4, 2003, p. 381–408.

[LEH 06] LEHRER E., ROSENBERG D., “What restrictions do Bayesian games impose on the value of information?”, *Journal of Mathematical Economics*, vol. 42, 2006, p. 343–357.

[NEY 91] NEYMAN A., “The positive value of information”, *Games and Economic Behavior*, vol. 3, 1991, p. 350–355.

[OSB 98] OSBORNE M., RUBINSTEIN A., *A Course in Game Theory*, The MIT Press, 1998.

[PLA 89] PLAZA J., “Logics of Public Communications”, EMRICH M., PFEIFER M., HADZIKADIC M., RAS Z., Eds., *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, p. 201–216, Academic Press, New York, 1989.

[ROO 03] VAN ROOY R., “Quality and Quantity of Information Exchange”, *Journal of Logic, Language and Information*, vol. 12, 2003, p. 423–451.