

What will they say? – Public Announcement Games

October 15, 2010

Abstract. Dynamic epistemic logic describes the possible information-changing actions available to individual agents, and their knowledge pre- and post conditions. For example, public announcement logic describes actions in the form of public, truthful announcements. However, little research so far has considered describing and analysing rational choice between such actions, i.e., predicting what rational self-interested agents actually will or should do. Since the outcome of information exchange ultimately depends on the actions chosen by all the agents in the system, and assuming that agents have preferences over such outcomes, this is a *game theoretic* scenario. This is, in our opinion, an interesting general research direction, combining logic and game theory in the study of rational information exchange. In this article we take some first steps in this direction: we consider the case where available actions are public announcements, and where each agent has a (typically epistemic) goal formula that she would like to become true. What will each agent announce? The truth of the goal formula also depends on the announcements made by other agents. We analyse such *public announcement games*.

Keywords: modal logic, epistemic logic, imperfect information games, action logic, public announcements

1. Introduction

Epistemic logic has been around since the 1960s (Hintikka, 1962; von Wright, 1951). Initially, in this logic one could model the knowledge of a single agent only. Propositions like Kp express that the agent knows p , i.e., in all indistinguishable states of the Kripke model in which formulas are interpreted, p must be true. Multi-agent epistemic logic became widely studied since the early 1980s, two influential overviews of this period are (Fagin et al., 1995; Meyer and van der Hoek, 1995). Correspondingly we shift from Kp to $K_i p$, for ‘agent i knows p ’. The multi-agent direction made it possible to also introduce group epistemic operators, for general knowledge, common knowledge, and distributed knowledge.

One of the most prominent recent developments of epistemic logic is *dynamic* epistemic logic (DEL); see (van Ditmarsch et al., 2007) for an overview. Information is communicated, so knowledge is by no means static. DEL is an umbrella term for extensions of epistemic logic with dynamic operators for formalising reasoning about information change.

The idea of combining epistemic logic and dynamic modal logic (Harel, 1984; Harel et al., 2000) (developed in the 1980s as a spec-

ification language to reason about the correctness and behaviour of computer programs) to model information change is due to van Benthem (van Benthem, 1989). The perhaps simplest dynamic epistemic logic is *public announcement logic* (PAL) (Plaza, 1989; Gerbrandy and Groeneveld, 1997), which makes epistemic logic dynamic by adding modal operators $\langle \psi \rangle$, where ψ is a formula. The intended meaning of $\langle \psi \rangle \varphi$ is that after ψ is truthfully and publicly announced, φ will be true. More sophisticated and general frameworks model more complex events than public announcements, such as card showing actions, where different agents have different perspectives on the action (one agent can see something is going on while the others know exactly what is happening). Different approaches have been developed by Baltag et al., Gerbrandy, van Linder et al., van Ditmarsch (1998, 1999, 1995, 2000) and many others. Further recent developments include the incorporation of factual change into languages that express epistemic change (van Benthem et al., 2006; van Ditmarsch et al., 2005), preference-based modelling of belief revision with dynamic modal operators (Aucher, 2003; van Ditmarsch, 2005; Baltag and Smets, 2008; van Benthem and Liu, 2007), the integration of dynamic epistemic logics with temporal epistemic logics (van Benthem et al., 2009), and various forms of quantification over propositional variables, as in (Balbiani et al., 2008; Ågotnes and van Ditmarsch, 2008; Ågotnes et al., 2010a).

While dynamic epistemic logic describes the possible information-changing actions available to individual agents, and their knowledge pre- and post conditions, little research so far has considered describing and analysing rational choice between such actions, i.e., predicting what rational self-interested agents actually will do. Since the outcome of information exchange ultimately depends on the actions chosen by all the agents in the system, this is a *game theoretic* scenario.

The purpose of this paper is twofold.

First, we want to point to an interesting general research direction in the theory of information change, combining logic and game theory by analysing DEL from a game theoretic perspective; describing not only what *can* happen, as standard dynamic epistemic logics do, but also what *will* (or *should*) happen. There are many free variables here:

- On the logic side, different types of actions: public announcements, private announcements, general action models, etc.
- On the game theory side, different types of games and related frameworks: strategic games, extensive games, even social choice and mechanism design models.

- The representation of preferences. The question of rational action in a dynamic epistemic logic framework presupposes *preferences* over (multi-agent) epistemic states. Preference *representation* is currently an active research area in logic and computer science. Examples of such representations include logical *goal formulae* (Harrenstein et al., 2001), *weighted formulae* (Jeong and Shoham, 2005; Lang et al., 2006; Uckelman et al., 2009; Elkind et al., 2009), and *CP-nets* (Boutilier et al., 1999; Boutilier et al., 2003; Boutilier et al., 2004).

While various combinations of modal logic and game theory is currently an active research area (van der Hoek and Pauly, 2006), we are not aware of any works analysing action (for example identifying equilibria) in DEL from a game theoretic perspective.

Second, we make some first steps, in analysing the perhaps simplest scenario:

- We assume that actions are *public announcements*,
- we look at *strategic form games*, and
- we represent preferences using logical *goal formulae*.

Thus, we consider situations where each agent has a goal, a typically epistemic formula he or she would like to become true, and where the available actions are public announcements. For example, Ann's goal might be that Bill knows Ann's secret message (or hand of cards) without Cath knowing it. What will each agent announce, assuming common knowledge of the situation? The truth value of the goal formula typically depends on the announcements made by several agents, hence, again, the game theoretic scenario.

In more detail, we want to model and analyse situations where:

- agents may have incomplete information about the world and about each others' knowledge;
- agents have goals in the form of epistemic formulae, and agents' goals are common knowledge among all agents;
- each agent chooses an announcement to make, which she knows to be true;
- all agents choose their announcements simultaneously; and
- all agents act rationally, i.e., they try to obtain their goals.

What can we say about how such agents will, or should, act?

Games with imperfect information have of course been studied in game theory. For the setting described above (strategic form games), the most relevant framework is *Bayesian games* (Harsanyi, 1968). Note that the games described above is a particular type of imperfect information games, where strategies and information available in each state are closely interconnected (strategies, i.e., possible announcements, *are* information).

The remainder of the paper is organised as follows. In Section 2 we review the syntax and semantics of public announcement logic, as well as some concepts from game theory. In Section 3 we introduce our model of *epistemic goal structures* (EGSs), and associate a (pointed) *public announcement game* with each state in an EGS. A particularly interesting special case is when one or more of the goal formulae are *positive* (essentially, only contains negation immediately preceding atomic propositions). In Section 4 we discuss a relevant version of dominance in this setting. In Section 5 we discuss the notion of Nash equilibrium, and, in particular, we introduce a way to view an EGS as a single strategic game. These induced games are, in practice, Bayesian games, as we observe in Section 7. We conclude in Section 8.

2. Background

2.1. PUBLIC ANNOUNCEMENT LOGIC

The language \mathcal{L}_{pal} of public announcement logic (PAL) (Plaza, 1989) over a set of agents $N = \{1, \dots, n\}$ 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 $\neg K_i\neg\varphi$, in addition to the usual defined propositional connectives.

A *Kripke structure* or *epistemic model* over N and Θ is a tuple $M = (S, \sim_1, \dots, \sim_n, V)$ 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 $V : \Theta \rightarrow 2^S$ maps 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*. We also assume that Kripke

structures are *connected*¹ – this assumption is not really needed², but it makes the technical treatment easier.

We will sometimes use the notation $[s]_i$ for the s -equivalence class of agent i , i.e., $[s]_i = \{t : s \sim_i t\}$.

When $M = (S, \sim_1, \dots, \sim_n, V)$ and $M' = (S', \sim'_1, \dots, \sim'_n, V')$ are both epistemic models over N and Θ , $S \subseteq S'$ and each \sim_i and V are the results of restricting each \sim'_i and V' to S , we say that M is a *submodel* of M' . We will sometimes abuse notation and write $M \subseteq M'$ to mean that M is a submodel of M' . Likewise, we sometimes abuse notation and treat a model M as its state space, writing, e.g., $s \in M$ for $s \in S$.

The interpretation of formulae in a pointed Kripke structure is defined as follows.

$$M, s \models p \text{ iff } p \in V(s)$$

$$M, s \models K_i \varphi \text{ iff for every } t \text{ such that } s \sim_i t, M, t \models \varphi$$

$$M, s \models \neg \varphi \text{ iff not } M, s \models \varphi$$

$$M, s \models \varphi \wedge \psi \text{ iff } M, s \models \varphi \text{ and } M, s \models \psi$$

$$M, s \models [\varphi] \psi \text{ iff } M, s \models \varphi \text{ implies that } M|\varphi, s \models \psi$$

where

$$M|\varphi = (S', \sim'_1, \dots, \sim'_n, V')$$

such that

$$\begin{aligned} S' &= \{s' \in S : M, s' \models \varphi\}, \\ \sim'_i &= \sim_i \cap (S' \times S') \text{ and} \\ V'(p) &= V(p) \cap S'. \end{aligned}$$

The update of M by φ , $M|\varphi$, is the submodel of M obtained by removing states where φ is false.

Observe that the interpretation of the dual public announcement operator is:

$$M, s \models \langle \varphi \rangle \psi \text{ iff } M, s \models \varphi \text{ and } M|\varphi, s \models \psi$$

We write $M \models \varphi$ to denote the fact that $M, s \models \varphi$ for every state $s \in M$.

¹ A finite Kripke structure is connected if every state is *reachable* from every other state via the indistinguishability relations. Formally: for every pair $s, t \in S$ there is a sequence of agents i_0, \dots, i_k ($k \geq 0$) and states s_0, \dots, s_{k+1} such that $s_j \sim_{i_j} s_{j+1}$ for any $0 \leq j \leq k$ and $s_0 = s$ and $s_k = t$.

² Every pointed Kripke structure is equivalent to a connected pointed Kripke structure, in the sense that the same formulae are satisfied.

The purely epistemic fragment of the language (i.e., formulae not containing public announcement operators $[\varphi]$) is denoted \mathcal{L}_{el} . It was already shown in Plaza’s original publication on that logic (Plaza, 1989) that the language of PAL is no more expressive than the purely epistemic fragment.

2.2. STRATEGIC GAMES

We very briefly review the notion of a strategic game and some related concepts from game theory; see, e.g., (Osborne and Rubinstein, 1994) for further details.

A *strategic game* is a triple $G = \langle N, \{A_i : i \in N\}, \{u_i : i \in N\} \rangle$ where

- 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*.
- for each $i \in N$, $u_i : A \rightarrow \mathbb{R}$ is the *payoff function* for i , mapping each strategy profile to a number.

When $(a_1, \dots, a_n) \in A$, the notation $(a_1, \dots, a_n)[a_i/a'_i]$ 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, \dots, a_n) is a Nash equilibrium if and only if for all $i \in N$, for all $a'_i \neq a_i$, $u_i((a_1, \dots, a_n)[a_i/a'_i]) \leq u_i(a_1, \dots, a_n)$. In this paper we are mostly interested in pure strategies, and by “Nash equilibrium” we will usually implicitly mean the pure strategy variant.

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 choose³. Formally, a strategy a_i for agent i is weakly dominant if and only if for all agents j , for all a'_j , $u_i(a'_1, \dots, a'_n) \leq u_i((a'_1, \dots, a'_n)[a'_i/a_i])$. Clearly, a strategy profile where all the strategies are weakly dominant is a Nash equilibrium.

Concepts and terminology relating to *imperfect information* games will be introduced in Section 7.

³ The literature differ in the definition of weakly dominant strategies. Another common definition in addition requires that the strategy is strictly better against at least one combination of actions by the other agents.

3. Pointed Public Announcement Games

Given n agents, an epistemic model $M = (S, \sim_1, \dots, \sim_n, V)$ encodes their uncertainty about facts and about each other. We also assume that *goal formulae* $\gamma_1, \dots, \gamma_n$ express what the agents wish to achieve by their announcements, i.e., each agent i wishes goal γ_i to be true.

Definition 1 (Epistemic Goal Structure) *An epistemic goal structure (EGS) AG is a tuple $\langle M, \gamma_1, \dots, \gamma_n \rangle$ where M is an epistemic structure, and where $\gamma_1, \dots, \gamma_n \in \mathcal{L}_{pal}$. Formula γ_i is the goal for agent i . A pointed epistemic goal structure is a tuple (AG, s) where $s \in M$.*

In order to achieve their goals, an agent can choose between different truthful announcements, and we assume that all agents make these announcements simultaneously.

It is now very natural to associate a strategic game with a pointed epistemic goal structure: strategies, or actions, correspond to the individual announcements the agents can choose between, and a goal is satisfied iff it is true after all the agents simultaneously make their chosen announcement.

Definition 2 (Pointed public announcement game, state game)

The pointed public announcement game (pointed PAG) $G(AG, s)$ associated with state $s \in M$ of epistemic goal structure $AG = \langle M, \gamma_1, \dots, \gamma_n \rangle$ is the strategic game defined by

- $N = \{1, \dots, n\}$,
- $A_i = \{\varphi_i \in \mathcal{L}_{pal} : M, s \models K_i \varphi_i\}$,
- $u_i(\varphi_1, \dots, \varphi_n) = \begin{cases} 1 & \text{if } M, s \models \langle K_1 \varphi_1 \wedge \dots \wedge K_n \varphi_n \rangle \gamma_i \\ 0 & \text{otherwise} \end{cases}$

A pointed public announcement game will also be referred to as a state game.

Like in Boolean games (Harrenstein, 2004; Harrenstein et al., 2001), binary utilities are implicit in these public announcement games; an agent's goal is either satisfied or not. In the definition of the payoff function an agent gets the value 1 iff her goal is satisfied after the joint announcement⁴.

Note that every agent i always has at least one strategy available; announcing a formula φ (e.g., a propositional tautology \top) such that $M|K_i \varphi = M$. We call such an announcement a *trivial announcement*.

⁴ Note that, given the definition of a strategy, this condition is equivalent to $M, s \models [K_1 \varphi_1 \wedge \dots \wedge K_n \varphi_n] \gamma_i$.

Sometimes, for convenience, we will identify (AG, s) with its associated pointed PAG $G(AG, s)$, and for example refer to a strategy or a Nash equilibrium of the pointed epistemic goal structure.

For goal satisfaction we do not require that the agent actually *knows* that the goal has been satisfied (that $K_i\gamma_i$ is true) after the announcement; a goal for agent i is *allowed* to take the form $\gamma_i = K_i\psi_i$, but this is not required, and a fortiori, we do not assume achievement of the goals to be common knowledge either. We therefore have to assume an oracle to announce the outcome of the game in the general case.

Agents may have different strategies available in different states of the game. Note, however, that we can also view an arbitrary formula $\varphi \in \mathcal{L}_{pal}$ as a strategy in *any* state, as follows. In this view, the formula φ stands for the strategy of announcing *whether* the agent knows φ . If, indeed the agent knows φ , φ stands for announcing $K_i\varphi$ (i.e., the strategy φ). If she does not know φ , i.e. if $\neg K_i\varphi$ is true, then $K_i\neg K_i\varphi$ is also true (due to the requirement that the accessibility relations are equivalences), and then φ stands for the announcement of $K_i\neg K_i\varphi$ (i.e., the strategy $\neg K_i\varphi$). This view allows us to talk about strategies in arbitrary states. Formally, for any formula φ and state s define $\varphi(s)$ (the strategy denoted by φ in state s) as follows: $\varphi(s) = \varphi$ if $M, s \models K_i\varphi$; $\varphi(s) = \neg K_i\varphi$ otherwise.

Although by the definition there is an infinite number of different strategies (syntactically different formulae true in the state), all (finite) epistemic goal structures have only finitely many strategies that really are *different* for our purposes (pointed public announcement games), since there can be only finitely many announcements with different epistemic content. We define two strategies φ and ψ to be *equivalent* for agent i in model M iff $M \models K_i\varphi \leftrightarrow K_i\psi$. Note that equivalent strategies are always available in the same state games, and in a given state game two equivalent strategies can be substituted for each other in any strategy profile without changing the payoff to any agent. Subject to that equivalence, the number of strategies for i is a function of the number of its equivalence classes in the model M . From now on we will assume that this identification of equivalent strategies is always made. In practice, in examples etc., we will just choose *one* of the many equivalent (syntactic) formulae to represent a strategy.

Proposition 3 (Counting strategies) *If player i has m equivalence classes in M , the number of (non-equivalent) strategies for i is 2^{m-1}*

in every state if M is bisimulation contracted⁵, and at most 2^{m-1} in the general case.

Proof. Given that i has m equivalence classes, and that there are 2^m different subsets of a set of m elements, there are 2^m unions of i -equivalence classes. Observe that the denotation of any announcement made by player i must be such a union, as player i only announces what she *knows* to be true: her announcement has form $K_i\varphi$. Conversely, if the model is bisimulation contracted, any such union is the denotation of some announcement⁶. How many of those unions contain a given actual state? This number now is independent from that state. This is because for each such union of i -equivalence classes, its complement on the domain is also a union of i -equivalence classes; and any state should therefore be either in one or the other. In other words, we are counting the different ways to partition the domain in two parts (a *dichotomy*) such that the partition induced by player i is a refinement of that. Therefore there are $\frac{2^m}{2} = 2^{m-1}$ disjoint pairs of unions of i -equivalence classes.

Example 4 (Counting and choosing announcements) Consider the following formal model of a situation: a two-player pointed EGS $(\langle M, \gamma_{Ann}, \gamma_{Bill} \rangle, s)$, where M is the following structure

$$\bullet_t \xrightarrow{\neg p_B, p_A} \text{Ann} \dots \bullet_s \xrightarrow{p_B, p_A} \text{Bill} \dots \bullet_u \xrightarrow{p_B, \neg p_A}$$

and

$$\begin{aligned} \gamma_{Ann} &= (K_B p_A \vee K_B \neg p_A) \rightarrow (K_A p_B \vee K_A \neg p_B) \\ \gamma_{Bill} &= (K_A p_B \vee K_A \neg p_B) \rightarrow (K_B p_A \vee K_B \neg p_A) \end{aligned}$$

Ann's goal is that Bill does not get to know whether p_A is true unless Ann gets to know whether p_B is true, and similarly the other way around. s is the actual state of the world. In that state, p_A is true and Ann knows this, and the same for p_B and Bill. Ann does not know whether Bill already knows p_A , and similarly Bill does not know whether Ann already knows p_B . Furthermore, Ann does not know

⁵ Roughly speaking; the model does not contain two different logically equivalent states. Assuming that the epistemic models we consider are bisimulation contracted would not be a restriction for our purposes: every pointed epistemic model can be reduced to one which is bisimulation contracted and which has the same logical theory and thus the same state games. See Appendix A for details.

⁶ Technically, the announcement is constructed as follows. For each of the, say k , equivalence classes *not* in the union, let α_i be the *characteristic formula* for one (arbitrarily chosen) of the states in that equivalence class; the needed announcement is $K_i \neg(\alpha_1 \vee \dots \vee \alpha_k)$. See the proof of Lemma 7 for further details about this technique.




Figure 1. A public announcement game consisting of three state games.

can therefore not be a reliable method of identifying rational outcomes. Figure 1 illustrates the state games associated with also the two other states of Example 5. Clearly, if the actual state is s , the state game is not known by any of the players – in fact, they don’t even know all the actions available to the other player. Indeed, while (p_A, p_B) is a Nash equilibrium in the state game in s , it is not in the other state (t) which *Ann* considers possible – she does not even know for certain that p_B is a possible action for *Bill*. Anne’s knowledge about the situation is imperfect, and she has to base a rational strategy on more than just the state game for s . And the same holds for *Bill*.

Thus, the situations we are interested in can be modelled as a particular type of strategic games with imperfect information, where the strategies and information available in each state are closely interconnected (strategies *are* information) and where the same strategies are available in indiscernible states (but not necessarily in others). We now go on to discuss some possible outcomes and solutions of such games. In Section 7 we explain how standard models of strategic games with imperfect information⁷, such as Bayesian games (Harsanyi, 1968), can be applied to this setting.

3.1. POSITIVE GOALS

In this section we look more closely at a particular class of epistemic goal structures, where agents have *positive goals*.

In general, there is a pointed epistemic goal structure corresponding to *any* strategic game (with binary utilities), and state games can have

⁷ Our knowledge games can be seen as both incomplete information games and imperfect information games. Taking the initial state of uncertainty for granted, all moves—namely announcements—are public to all players: perfect information about game actions. From that perspective PAGs and apparently the original Harsanyi perspective (as explained to us by Giacomo Bonanno) are incomplete information games. On the other hand, the uncertain initial state of information can be seen as the result of a partially hidden move of nature: imperfect information games.

zero, one or several Nash equilibria. Restrictions on the form of the goals, however, can induce interesting game theoretic properties of the corresponding state games.

A natural and interesting case is when the goals are *positive formulae*, henceforth referred to as positive goals. The positive formulae are the following fragment of \mathcal{L}_{pal} :

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

where $p \in \Theta$. An example is $K_A q \vee K_A \neg q$ (*Ann's* goal is to know whether q is true). This notion of positive formulae is found in (van Ditmarsch and Kooi, 2006), which is an extension of several such notions going back van Benthem (2002) who observed that purely epistemic (without announcement operators) positive formulae are *preserved* under submodels.

Theorem 6 (van Ditmarsch and Kooi (2006)) *If (M, s) is a pointed epistemic model and M' is a submodel of M containing s , then for any positive formula ψ : $M, s \models \psi \Rightarrow M', s \models \psi$.*

For a general pointed epistemic goal structure, a *maximal informative announcement* (MIA) of an agent i is a strategy φ such that for every strategy ψ , $M|K_i \varphi \subseteq M|K_i \psi$. We state some properties of MIAs.

Lemma 7

1. *A maximal informative announcement always exists (for finite structures).*
2. *The maximal informative announcement is unique (up to equivalence of strategies).*
3. *If α is the maximal informative announcement of agent i in (M, s) and φ, β are arbitrary formulae such that $M, s \models \varphi \wedge K_i \beta$, then $M|(\varphi \wedge K_i \alpha) \subseteq M|(\varphi \wedge K_i \beta)$.*
4. *If $s \sim_i t$, then the maximal informative announcements of i in (M, s) and in (M, t) are equivalent.*

Proof. Appendix B.

If an agent has a positive goal, then that agent has a weakly dominant strategy – namely her maximal informative announcement.

Theorem 8 *If the goal γ_i of agent i in a goal structure AG is positive, then the maximal informative announcement of agent i is weakly dominant in the state game $G(AG, s)$ of any state s .*

Proof. For each agent i , let $\alpha_i \in \mathcal{L}_{pal}$ be the maximal informative announcement for i in M, s . Assume that α_i is not weakly dominant in $G(AG, s)$. Then there exist $\alpha'_1, \dots, \alpha'_n$ such that $u_i(\alpha'_1, \dots, \alpha'_n) > u_i(\alpha'_1, \dots, \alpha'_n)[\alpha'_i/\alpha_i]$. That means that $u_i(\alpha'_1, \dots, \alpha'_n) = 1$ and $u_i(\alpha'_1, \dots, \alpha'_n)[\alpha'_i/\alpha_i] = 0$, in other words that $M \mid \bigwedge K_j \alpha'_j, s \models \gamma_i$ and $M \mid \bigwedge_{j \neq i} K_j \alpha'_j \wedge K_i \alpha_i, s \not\models \gamma_i$. But that contradicts Theorem 6, since $M \mid \bigwedge_{j \neq i} K_j \alpha'_j \wedge K_i \alpha_i \subseteq M \mid \bigwedge K_j \alpha'_j$ by Lemma 7.3.

One consequence is that if all the goals are positive, the state game *always* has at least one Nash equilibrium – the combination of maximal informative announcements.

Corollary 9 *If $AG = \langle M, \gamma_1, \dots, \gamma_n \rangle$ is a goal structure with only positive goals, then for every state $s \in M$, the strategy profile consisting of the maximal informative announcement for each agent is a Nash equilibrium in the state game $G(AG, s)$.*

While the maximal informative announcements form a Nash equilibrium in every state game when all goals are positive, note that different states may have *different* maximal informative announcements. We now turn to the question of the agents' knowledge about strategies and equilibria.

4. Weak Dominance

Let us first consider weakly dominant (wd) strategies.

In order to discuss knowledge of properties such as weak dominance, recall that the interpretation of the formal epistemic language in epistemic models relies on the idea that an agent knows something iff it is true in all the states she considers possible. But this idea is a general one, and does not only apply to propositions expressible in the logical language but also, more informally, to other possible state-properties – such as “there is a weakly dominant strategy for i ”. For example, an agent i knows that there exists a weakly dominant strategy for i iff “there is a weakly dominant strategy for i ” is true in all the states she considers possible.

It should be clear from the discussion in the previous section that there is a crucial distinction, in a pointed PAG, between on the one hand:

- the *existence* of a strategy for an agent which is weakly dominant for that agent,

and on the other hand:

- the existence of a strategy for an agent which that agent *knows* is weakly dominant.

This is because it might be the case that there is a strategy which is weakly dominant in one of the states the agent considers possible, but not in another. For example, in state t of the model in Figure 1, p_A is weakly dominant for *Ann*, but *Ann* does not know this. Contrast this with the fact that \top is also weakly dominant for *Ann*, but this *Ann* knows (it is weakly dominant in all the states she considers possible). Furthermore, \top is also weakly dominant for *Ann* in the remaining state u , so even *Bill* knows that \top is weakly dominant for *Ann*, or rather, \top is dominant throughout the entire model. Knowledge of weakly dominant strategies is a natural solution concept for PAGs. But another important distinction must be made; a distinction well known in the knowledge and action literature (Jamroga and van der Hoek, 2004; Jamroga and Ågotnes, 2006; Jamroga and Ågotnes, 2007).

Definition 10 (De re / de dicto weak dominance) *Given a pointed epistemic goal structure $(\langle M, \gamma_i, \dots, \gamma_n \rangle, s)$ and an agent i .*

- i has a weakly dominant strategy *de dicto* iff for any state t such that $s \sim_i t$, i has a weakly dominant strategy in the state game of t .
- i has a weakly dominant strategy *de re* iff there is some strategy for i which is weakly dominant in the state game of any state t such that $s \sim_i t$.
- i has a global weakly dominant strategy *de re* iff there is some strategy for i which is weakly dominant in the state game of any state t .

Clearly every global weakly dominant strategy *de re* is a weakly dominant strategy *de re*, and every weakly dominant strategy *de re* is a weakly dominant strategy *de dicto*.

If an agent has a wd strategy *de dicto*, she knows that she has a wd strategy, i.e., she has a wd strategy in all states she considers possible, but she does not necessarily know which strategy is dominant; it is not necessarily the same strategy that is dominant in all the possible states. If she has a wd strategy *de re*, on the other hand, she knows which strategy is dominant; the same strategy is dominant in all the states she considers possible. If she has a global wd strategy *de re*, she does not even have to know her state. For example, if this were about uncertainty of card ownership in a card game, say, she can play her global strategy *de re* without looking at her cards!

In state s in the model in Figure 1, *Ann* has a wd strategy *de re*, namely \top . Her other strategy, p_A , is not weakly dominant. The strategy \top is also a global wd strategy *de re*. An example where an agent has a wd strategy *de dicto* but not *de re* will be shown in Example 21, later.

For positive goals, we have the following property.

Corollary 11 *If the goal of an agent is positive, then that agent has a weakly dominant strategy de re in any state.*

Proof. Follows from Theorem 8 and Lemma 7.4.

This shows that the case of positive goals is indeed significant: *with a positive goal, the agent always knows what to do.*

5. Nash Equilibrium

5.1. KNOWING THAT, KNOWING HOW

What about the Nash equilibrium? Clearly, we can have situations similar to the case of weakly dominant strategies: there might be a Nash equilibrium without the agents knowing it; the agents might know that there is a Nash equilibrium but not necessarily know what it is (there might be different equilibria in different accessible states). However, what “knowing” means here is not as clear as in the case of dominant strategies where knowledge of a single agent was needed. In the case of the Nash equilibrium there are *several* agents involved. Group notions of knowledge, such everybody-knows, distributed knowledge and common knowledge, have been studied in the context of the *de dicto/de re* distinction before (Jamroga and van der Hoek, 2004). For our purpose, we argue that the proper type of group knowledge for knowing a Nash equilibrium *de re* is *common knowledge*, since that is the assumption in game theory. Common knowledge of an equilibrium among all agents corresponds to a common equilibrium in all states of the model (since we assume connectedness). Thus, the existence of a Nash equilibrium *de re* is a *model* property, rather than a *pointed* model property as in the case of dominant strategies.

Definition 12 *Given an epistemic goal structure AG, we say that there is a Nash equilibrium de re if there exists a tuple of formulae, one for each agent, which constitutes a Nash equilibrium in the state game of every state in the EGS.*

However, it turns out that this solution concept is very limited, as the following theorem shows.

Theorem 13 *In any Nash equilibrium de re of an epistemic goal structure, every agent chooses the trivial announcement.*

Proof. Assume otherwise; i.e., that $\langle \varphi_1, \dots, \varphi_n \rangle$ is a Nash equilibrium de re and that there is an agent i such that $M \not\models K_i \varphi_i$. That means that there is a state $t \in S$ such that $M, t \not\models K_i \varphi_i$ for some i . But since $\langle \varphi_1, \dots, \varphi_n \rangle$ is a Nash equilibrium de re then φ_i must be a valid strategy for i in every state, including in t , and that means that it must be the case that $M, t \models K_i \varphi_i$. A contradiction.

While the argument for the above theorem is easy, it amounts to an interesting “impossibility result”: *there can never be common knowledge of a Nash equilibrium that does not consist of trivial announcements.*

For example, in the EGS in Figure 1, there is a Nash equilibrium de re, because the strategy profile (\top, \top) is a Nash equilibrium in all the state games. An example where there are Nash equilibria in all the state games but no Nash equilibrium de re will be shown later (Example 21).

The proof of Theorem 13 in fact proves something stronger, namely that the only strategy profile that is executable in every state game is

$\langle \overbrace{\top, \dots, \top}^n \rangle$. That means that the corresponding result, that common knowledge is impossible for other profiles, holds for *any* equilibrium concept and not only Nash equilibrium.

5.2. INDUCED PUBLIC ANNOUNCEMENT GAMES

We have seen how we can associate a strategic game with each state in an epistemic goal structure. The question we address in this section is: can we view an epistemic goal structure as a single strategic game?

We propose the following definition. It can, under certain assumptions, in fact be seen as a type of Bayesian games (Harsanyi, 1968), as we discuss formally in Section 7.

Definition 14 *Given a epistemic goal structure $AG = \langle M, \gamma_1, \dots, \gamma_n \rangle$ with $M = (S, \sim_1, \dots, \sim_n, V)$, the induced public announcement game $G(AG)$ is defined as follows:*

- $N = \{1, \dots, n\}$
- A_i is the set of functions $a_i : S \rightarrow \mathcal{L}_{pal}$ with the following properties:
 - *Truthfulness:* $M, s \models K_i a(s)$ for any s
 - *Uniformity:* $s \sim_i t \Rightarrow a_i(s) = a_i(t)$

Thus, a strategy $a_i \in A_i$ gives a possible announcement for each state, but always the same announcement for indiscernible states (note that the same announcements are always truthful in indiscernible states). Alternatively, a_i can be seen as a function mapping equivalence classes to announcements.

- The payoffs are defined as follows. For any state s in AG , let $G(AG, s) = (N, \{A_i^s : i \in N\}, \{u_i^s : i \in N\})$ be the state game associated with s (Def. 2). Define, for any $(a_1, \dots, a_n) \in A_1 \times \dots \times A_n$:

$$u_i(a_1, \dots, a_n) = \frac{\sum_{s \in S} u_i^s(a_1(s), \dots, a_n(s))}{|S|}$$

For further reference, we call a strategy a_i *global* iff there exists a formula $\varphi \in \mathcal{L}_{pal}$ such that for any $s \in S$, $a_i(s) = \varphi(s)$ ($\varphi(s)$ is defined on p. 8); i.e., a_i is the strategy of announcing “whether φ ” throughout the model. The *trivial global strategy* is the strategy a such that $a(s) = \top$ for any s .

As for pointed EGSs and state games, we shall sometimes implicitly view an EGS AG implicitly as the game $G(AG)$, and talk about its strategies, Nash equilibria, etc. There are various important points to consider in the above definition.

Strategies in the induced game Consider agents *Ann* and *Bill*. Even though the current state is a member of the equivalence class that *Ann* currently considers as possible states, she might consider many possibilities for what *Bill*’s equivalence class might be. Thus, she must take into account what *Bill* is likely to do in all of these circumstances. But *Bill*’s choice depends on the states he considers possible, and he may consider a state possible that is not in *Ann*’s actual equivalence class. And so on. We therefore have to take into account all reachable states, that is (because our model is connected) all states in the model. This means that *Ann* has to know what *Bill* might do in any of the states in the model, independent from her actual state, and similarly for *Bill*. This explains why strategies are formulated as contingencies for every state in the model, i.e., as functions from *every* state to a choice of strategy in that state.

Payoffs in the induced game Second, payoff is computed by taking the average over all states in the model. It is clear that it does not suffice to look only in the current state, as each agent also might consider other states possible. But why not, then, compute an agent’s payoff by taking the average over all the states that that agent considers possible,

i.e., that agent's equivalence class? We cannot, because the strategic game must be common knowledge, in order for solution concepts such as the Nash equilibrium to make sense. In Example 5, if the actual state is t , *Ann* considers it possible that *Bill* considers state u possible. If we only take the average over s and t for *Ann*, u will not be taken into account: that would be irrational for *Ann*, as she knows *Bill* to be uncertain between s and u if s is the case. Averaging over all reachable states corresponds to averaging over all states commonly considered possible (all states accessible according to the accessibility relation for common knowledge). This is also the reason that the induced game is not induced from a *pointed* EGS: the induced game would be the same at all points. This is as it should be, since the game should be common knowledge at any state. The computed payoffs are *expected* payoffs, not expected by a particular agent in the game, but expected payoffs as computed by a *common knower* – an agent whose knowledge is exactly the common knowledge among all agents in the game.

Counting strategies in the induced game We observed that in the state game for (M, s) a player i with m equivalence classes has 2^{m-1} strategies (assuming that the model is bisimulation contracted). In the induced game we can also count the number of strategies.

Proposition 15 (Number of strategies in the induced game) *If player i has m equivalence classes in a bisimulation contracted model M , the number of (non-equivalent) strategies for i in the induced public announcement game is 2^{m^2-m} in every state.*

Proof. In a given state the player now can choose any of her 2^{m-1} state game strategies (Prop. 3). Given uniformity, this means that player i has that freedom for each of her m equivalence classes. Of course, the action chosen in one equivalence class does not have to be the action chosen in another equivalence class: if *Ann* knows p_A , she might want to announce \top , but if *Ann* knows $\neg p_A$, she might not want to announce \top but instead prefer to announce $\neg p_A$, etc. All these choices can be made completely independently. The total number of strategies for the induced game is therefore $2^{m-1} \times \dots \times 2^{m-1}$ (m times) which equals $(2^{m-1})^m = 2^{(m-1) \cdot m} = 2^{m^2-m}$.

This delivers a staggering number of strategy profiles: $\times_{i \in \{1, \dots, n\}} 2^{m_i^2 - m_i}$ (where the number of equivalence classes for player i is m_i). Given that the number of equivalence classes is in the order of the number of states $|M|$ of the model M , this gives us $O(2^{(|M| \cdot |M| \cdot n)})$ strategy profiles.

We will shortly explain the induced game further through an example. First, some theoretical points.

Proposition 16 *If agent i has a weakly dominant strategy $de\ re$ in $G(AG, s)$ for every state s in a EGS AG , then there is a weakly dominant strategy for i in the induced game $G(AG)$.*

Proof. A weakly dominant strategy a in the induced game is defined by, for any state s , taking $a(s)$ to be a wd strategy in the state game in s under the restriction that the same strategy is chosen for all states in the same equivalence class (this is possible because the agent has a strategy *de re*). a is truthful and uniform by definition. Wlog. assume that there are only two agents, and that $i = 1$. Suppose that a is not weakly dominant. Then there is some other strategy a' for 1, and some strategy b for 2 such that

$$\frac{\sum_{s \in M} u_1^s(a'(s), b(s))}{|M|} > \frac{\sum_{s \in M} u_1^s(a(s), b(s))}{|M|}$$

Since payoffs are non-negative, this implies that $u_1^s(a'(s), b(s)) > u_1^s(a(s), b(s))$ for some s . But then $a(s)$ is not weakly dominant in the state game in s after all, which is a contradiction.

This immediately follows for positive goals:

Corollary 17 *If the goal of agent i is positive, then there is a weakly dominant strategy for i in the induced game $G(AG)$.*

Proof. Follows from Proposition 16 and Corollary 11.

Definition 18 *A Nash Announcement Equilibrium (NAE) of a EGS is a Nash equilibrium of the induced public announcement game.*

Example 19 *Let us continue Example 5. We construct the induced game as follows (it is instructive to inspect the state games as illustrated in Fig. 1 on p. 11). Both Ann and Bill have two equivalence classes. Thus, as observed in Proposition 15, they each have 4 strategies. The set A_A of strategies for Ann contains the following:*

- $a_A^1: t, s \mapsto \top; u \mapsto \top$
- $a_A^2: t, s \mapsto \top; u \mapsto \neg p_A$
- $a_A^3: t, s \mapsto p_A; u \mapsto \top$
- $a_A^4: t, s \mapsto p_A; u \mapsto \neg p_A$

A_B (for Bill) is as follows:

- $a_B^1: u, s \mapsto \top; t \mapsto \top$
- $a_B^2: u, s \mapsto \top; t \mapsto \neg p_B$
- $a_B^3: u, s \mapsto p_B; t \mapsto \top$
- $a_B^4: u, s \mapsto p_B; t \mapsto \neg p_B$

In order to compute the payoffs, we need to check the payoffs in the state games for each state and combination of strategies. We have the following:

$\mathbf{a}_A^x, \mathbf{a}_B^y$	\mathbf{t}	\mathbf{s}	\mathbf{u}
1, 1	01	11	10
1, 2	11	11	10
1, 3	01	10	10
1, 4	11	10	10
2, 1	01	11	11
2, 2	11	11	11
2, 3	01	10	11
2, 4	11	10	11
3, 1	01	01	10
3, 2	11	01	10
3, 3	01	11	10
3, 4	11	11	10
4, 1	01	01	11
4, 2	11	01	11
4, 3	01	11	11
4, 4	11	11	11

We get the following payoff matrix. We will henceforth write the payoffs without dividing by the number of states, for ease of presentation (the equilibria do of course not depend on this):

	\mathbf{a}_B^1	\mathbf{a}_B^2	\mathbf{a}_B^3	\mathbf{a}_B^4
\mathbf{a}_A^1	<u>22</u>	<u>32</u>	21	31
\mathbf{a}_A^2	<u>23</u>	<u>33</u>	22	32
\mathbf{a}_A^3	12	22	<u>22</u>	<u>32</u>
\mathbf{a}_A^4	13	23	<u>23</u>	<u>33</u>

The Nash equilibria are underlined. They can be described informally as follows:

- (1,1) Both agents say nothing (informative), no matter what. (This is a global Nash equilibrium.)

- (1,2) *Ann says nothing, but Bill says $\neg p_B$ if the state is t (which Bill can discern from any other state) and nothing otherwise. Let us consider this in the case that the current state is s . Ann knows that the actual state is either s or t , but not which. Thus, in the equilibrium she will play \top under the assumption that Bill will play \top if the actual state is s and $\neg p_A$ if the actual state is t . Actually, Bill will play \top .*
- (2,1) *Similarly to (1,2), with Ann and Bill swapped*
- (2,2) *Ann says $\neg p_A$ if she knows it, i.e., if the state is u ; otherwise she says nothing. Similarly for Bill.*
- (3,3) *Ann says p_A if she knows it, i.e., if the state is in Ann's equivalence class $\{s, t\}$; otherwise she says nothing. Similarly for Bill.*
- (3,4) *Ann says p_A if she knows it, and otherwise nothing; and Bill always plays p_B , i.e., he says p_B if he knows it and $\neg p_B$ if he knows that.*
- (4,3) *Similarly to (3,4), for Ann and Bill swapped.*
- (4,4) *Both agents say everything they know. This is also a global Nash equilibrium.*

We can establish a connection to having a Nash equilibrium *de re*, similarly to Proposition 16 for dominant strategies.

Proposition 20 *If there exists a Nash equilibrium de re in a EGS, then the trivial global strategy is a Nash announcement equilibrium.*

Proof. Assume wlog. that there are only two agents. Assume that there is a Nash equilibrium *de re*; we know from Theorem 13 that (\top, \top) is a Nash equilibrium in every state game. Let (a, b) be the strategy profile for the induced game such that $a(s) = \top$ and $b(s) = \top$ for any s . Clearly, a and b are both uniform and truthful. Suppose that (a, b) is not a Nash equilibrium in the induced game. Then there is a better response a' for one of the agents, again wlog. assume for agent 1. In other words, there is a strategy a' for agent 1 such that $u_1(a', b) > u_1(a, b)$. But this entails that $u_1^s(a'(s), \top) > u_1^s(\top, \top)$ for some state s , and thus that there is a strategy z for agent 1 in the state game in s such that $u_1^s(z, \top) > u_1^s(\top, \top)$ – which contradicts the fact that (\top, \top) is a Nash equilibrium in the state game in s .

Proposition 20 does not hold in the other direction. A counter example is found in Example 21.

Before we consider the relationship to Bayesian games, let us discuss some more examples and observations.

6. Further Illustrations

Example 21 Define a EGS AG as follows. Let the model be as in Example 5, but change the goals as follows:

$$\begin{aligned}\gamma_{Ann} &= (K_B(p_A \wedge p_B) \wedge \neg K_{APB}) \vee (K_B(\neg p_B \wedge p_A) \wedge \hat{K}_A \hat{K}_B \neg p_A) \vee \\ &\quad (K_A(p_B \wedge \neg p_A) \wedge \hat{K}_B \hat{K}_A \neg p_B) \\ \gamma_{Bill} &= (K_A(p_A \wedge p_B) \wedge \neg K_{BP_A}) \vee (K_B(\neg p_B \wedge p_A) \wedge \hat{K}_A \hat{K}_B \neg p_A) \vee \\ &\quad (K_A(p_B \wedge \neg p_A) \wedge \hat{K}_B \hat{K}_A \neg p_B)\end{aligned}$$

Perhaps the reader finds these long formulae hard to read, but it suffices to trust that they give the following state games:

$\bullet_t^{\neg p_B, p_A}$ Ann <table style="border-collapse: collapse; margin: auto;"> <tr><td style="border-right: 1px solid black; padding: 2px 5px;">T</td><td style="padding: 2px 5px;">T</td><td style="padding: 2px 5px;">¬p_B</td></tr> <tr><td style="border-right: 1px solid black; padding: 2px 5px;">T</td><td style="padding: 2px 5px;">11</td><td style="padding: 2px 5px;">00</td></tr> <tr><td style="border-right: 1px solid black; padding: 2px 5px;">p_A</td><td style="padding: 2px 5px;">00</td><td style="padding: 2px 5px;">00</td></tr> </table>	T	T	¬p _B	T	11	00	p _A	00	00	$\bullet_s^{p_B, p_A}$ Bill <table style="border-collapse: collapse; margin: auto;"> <tr><td style="border-right: 1px solid black; padding: 2px 5px;">T</td><td style="padding: 2px 5px;">T</td><td style="padding: 2px 5px;">p_B</td></tr> <tr><td style="border-right: 1px solid black; padding: 2px 5px;">T</td><td style="padding: 2px 5px;">00</td><td style="padding: 2px 5px;">01</td></tr> <tr><td style="border-right: 1px solid black; padding: 2px 5px;">p_A</td><td style="padding: 2px 5px;">10</td><td style="padding: 2px 5px;">00</td></tr> </table>	T	T	p _B	T	00	01	p _A	10	00	$\bullet_u^{p_B, \neg p_A}$ <table style="border-collapse: collapse; margin: auto;"> <tr><td style="border-right: 1px solid black; padding: 2px 5px;">T</td><td style="padding: 2px 5px;">T</td><td style="padding: 2px 5px;">p_B</td></tr> <tr><td style="border-right: 1px solid black; padding: 2px 5px;">T</td><td style="padding: 2px 5px;">11</td><td style="padding: 2px 5px;">00</td></tr> <tr><td style="border-right: 1px solid black; padding: 2px 5px;">¬p_A</td><td style="padding: 2px 5px;">00</td><td style="padding: 2px 5px;">00</td></tr> </table>	T	T	p _B	T	11	00	¬p _A	00	00
T	T	¬p _B																											
T	11	00																											
p _A	00	00																											
T	T	p _B																											
T	00	01																											
p _A	10	00																											
T	T	p _B																											
T	11	00																											
¬p _A	00	00																											

This EGS has some properties not found in the EGS in Example 5 (Figure 1). First, Ann has a weakly dominant strategy de dicto, but not de re, in the pointed EGS (AG, s). The strategy p_A is weakly dominant in s, but not in t. There is, however, another weakly dominant strategy in t, namely T. Second, while every state game has a Nash equilibrium, there does not exist a Nash equilibrium de re in AG.

We get the following induced public announcement game, where the strategies are as in Example 19:

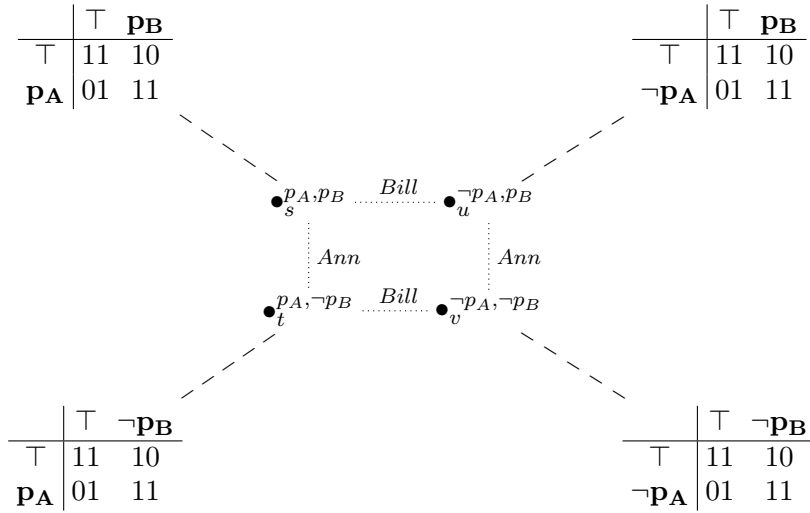
	<u>a_B¹</u>	a _B ²	<u>a_B³</u>	a _B ⁴
<u>a_A¹</u>	22	11	<u>12</u>	01
a _A ²	11	00	<u>12</u>	01
<u>a_A³</u>	<u>21</u>	<u>21</u>	00	00
a _A ⁴	10	10	00	<u>00</u>

The Nash announcement equilibria are underlined. Let us consider the situation in state s. There are several different Nash announcement equilibria, including: all agents announce T in all states (including in s). Note that (T, T) is not a Nash equilibrium in the state game in s. Another equilibrium is that both agents play “down” (i.e., not T) in

any state (also a Nash equilibrium in the state game in s). Note that (a_A^4, a_B^4) is a NAE while for example (a_A^3, a_B^4) is not. If the current state is s (or, from Ann's perspective the current equivalence class is $\{s, t\}$ and from Bill's perspective $\{s, u\}$), Ann will in fact do exactly the same if she uses strategy a_A^3 or a_A^4 . However, since Bill does not know whether Ann's equivalence class is $\{s, t\}$ or $\{u\}$, he must also consider what Ann does in u – which is exactly what differentiates a_A^3 and a_A^4 . Thus, the distinction between these two strategies is significant, even in a state (s) where they give the same action.

Example 22 Let us consider a more regular and symmetric EGS than the ones discussed so far. The situation is similar to the one in Example 5, but now Ann knows that Bill does not know p_A , and similarly for Bill/Ann/ p_B . The situation is modelled by the following goal formulae and Kripke structure. We have also shown the state games.

$$\begin{aligned}\gamma_{Ann} &= (K_B p_A \vee K_B \neg p_A) \rightarrow (K_A p_B \vee K_A \neg p_B) \\ \gamma_{Bill} &= (K_A p_B \vee K_A \neg p_B) \rightarrow (K_B p_A \vee K_B \neg p_A)\end{aligned}$$



Again, the induced game has four distinct strategies for each agent:

$$\begin{array}{l} \mathbf{x} \quad \mathbf{a}_A^x \qquad \qquad \qquad \mathbf{a}_B^x \\ 1 \quad s, t \mapsto \top; \quad u, v \mapsto \top \quad s, u \mapsto \top; \quad t, v \mapsto \top \\ 2 \quad s, t \mapsto \top; \quad u, v \mapsto \neg p_A \quad s, u \mapsto \top; \quad t, v \mapsto \neg p_B \\ 3 \quad s, t \mapsto p_A; \quad u, v \mapsto \top \quad s, u \mapsto p_B; \quad t, v \mapsto \top \\ 4 \quad s, t \mapsto p_A; \quad u, v \mapsto \neg p_A \quad s, u \mapsto p_B; \quad t, v \mapsto \neg p_B \end{array}$$

The induced public announcement game (Nash equilibria underlined):

	a_B^1	a_B^2	a_B^3	a_B^4
a_A^1	<u>44</u>	42	42	40
a_A^2	24	<u>33</u>	<u>33</u>	42
a_A^3	24	33	33	42
a_A^4	04	24	24	<u>44</u>

The game has two Nash equilibria. The first is that both agents say nothing, in all states. The strategies in this equilibrium are both dominant strategies. The second equilibrium (a_A^4, a_B^4) is that both agents tell everything they know, in all states.

In Example 22 the Nash announcement equilibria are all “composed” of Nash equilibria in the state game, in the following sense: for every NAE (a, b) and every state s , $(a(s), b(s))$ is a Nash equilibrium in the state game in s (albeit not *all* such compositions of Nash equilibria in the state games are NAE in the example). Indeed, this is also the case in Example 5. Is this a general property of PAGs? No, and a counter example is found in Example 21: (a_A^1, a_B^1) , because $(a_A^1(s), a_B^1(s))$ is not a Nash equilibrium in the state game in s .

We continue by modelling public announcement games as Bayesian games.

7. Public announcement games as Bayesian games

The most common model of strategic games with imperfect information is the *Bayesian game* (Harsanyi, 1968). Public announcement games can be modelled as Bayesian games. In fact, as we now discuss in more detail, the notion of the induced public announcement game introduced in Section 5.2 can be seen as an alternative formulation of a certain class of Bayesian games, in the following sense: Nash equilibria of the induced game are exactly the Nash equilibria of the Bayesian game (defined from the EGS in a natural way). We now make this correspondence precise. The discussion is somewhat technical, but the main message is: the construction of the induced game in Section 5.2 is simply a variant of the construction used in the analysis of Bayesian games.

In order to view epistemic game structures as Bayesian games, we must make some additional assumptions, because EGSs do not contain all the information needed for a Bayesian game:

- While an EGS describes equivalence classes of states, a Bayesian game in addition gives a probability distribution over the states in

an equivalence class. Here, we will assume that an agent considers the states in an equivalence class with *equal probability*.

- In a Bayesian game, it is assumed that players' have preferences over lotteries on the combination of strategy profiles and states. Here, we will assume that the utilities (0/1) we assign to players in state games, given a state and an action profile, represent the expected value of a *Bernoulli payoff function* (see, e.g., (Osborne, 2004)).

There is also one assumption usually made about Bayesian games, which does not hold for our games: it is usually assumed that a player has the same actions available in every state. However, this is not an essential assumption⁸.

The following definition of a Bayesian games is from (Osborne, 2004) (only slightly generalised to allow different actions in states receiving different signals).

A *Bayesian game* has the following components:

- N is the set of players;
- Ω is the finite set of *states* modelling the players' uncertainty about each other;

and for each $i \in N$:

- T_i is the set of *signals* that may be observed by player i , and $\tau_i : \Omega \rightarrow T_i$ is the *signal function* of player i ;
- For each $t_i \in T_i$, $A_i(t_i)$ is a set of strategies;
- for each signal $t_i \in T_i$, there is a probability distribution over the states $\tau^{-1}(t_i)$ consistent with the signal, modelling i 's *beliefs*. $Pr(\omega : t_i)$ gives the probability of state ω given signal t_i ;
- a *Bernoulli payoff function* u_i over pairs of states and action profiles available in the state, represents the player's preferences.

Definition 23 *Given an EGS $\langle M, \gamma_1, \dots, \gamma_n \rangle$, where $M = (S, \sim_1, \dots, \sim_n, V)$, the associated Bayesian game is defined as follows:*

- $N = \{1, \dots, n\}$;
- $\Omega = S$;

⁸ The essential assumption is that the same actions are available in states receiving the same signal.

- $T_i = \{[s]_i : s \in S\}$, where $[s]_i = \{t : s \sim_i t\}$, i.e., T_i is the set of equivalence classes for i ;
- $\tau_i(s) = [s]_i$;
- $A_i([s]_i) = \{\varphi : M, s \models K_i\varphi\}$;
- $Pr(s : [s]_i) = \frac{1}{|[s]_i|}$ (the uniform distribution assumption made above);
- $u_i(a_1, \dots, a_n, s) = u_i^s(a_1, \dots, a_n)$, where u_i^s is the payoff function of the state game in s .

To compute Nash equilibria of a Bayesian game, an intermediary set of players (i, t_i) , for every player $i \in N$ and signal $t_i \in T_i$ for that player, is employed. The equilibria are then computed for the game with those virtual players. The payoffs for the strategic game form derived from a Bayesian game are then computed using the probability distribution (beliefs) and preferences. Formally, the Nash equilibrium of a Bayesian game is defined as a Nash equilibrium of the following strategic game:

- $N' = (i, t_i)$, where $i \in N$ and $t_i \in T_i$;
- $A_{i,t_i} = A_i(t_i)$;
- $u_{i,t_i}(\{a_{j,t'_j} : j \in N, t'_j \in T_j\}) = \sum_{\omega \in \Omega} Pr(\omega : t_i) u_i(a_1, \dots, a_n, \omega)$, where $a_i = a_{i,t_i}$.

For the purpose of computing equilibria of the induced public announcement game we defined strategies somewhat differently, namely as total functions, required to be uniform, from states to announcements by i , i.e. from states to ‘pointed’ strategies for players i . The uniformity required was that different states in the same equivalence class for i are assigned the same strategy a_i , so in fact we can see a pure strategy in the induced game as a function from, for each agent, equivalence classes for that agent to her strategies in that class. In other words, as a function from, for each agent, signals to strategies. In this function to every signal can be assigned a different strategy. Therefore, these strategy choices are completely independent from one another and it therefore corresponds to playing the game for a set of players (i, t_i) instead, in the Bayesian version.

Theorem 24 *Given an epistemic goal structure, $\langle a_1, \dots, a_n \rangle$ is a Nash equilibrium in the induced public announcement game iff $\{a_{i,t_i} : i \in N, t_i \in T_i\}$, with $a_{i,t_i} = a_i(s)$ for some arbitrary $s \in \tau^{-1}(t_i)$, is a Nash equilibrium in the associated Bayesian game.*

Proof. Appendix C.

In the above theorem, a_{i,t_i} is well defined since a_i is uniform.

We leave it to the reader to construct the Bayesian game versions of the EGS examples discussed earlier.

8. Discussion

Epistemic goal structures naturally give rise to state games. We discussed properties of state games such as their number of strategies. In this setting, there are several different variants of notions such as weakly dominant strategies, corresponding to different types of knowledge. It is our view that there is not merely one single “correct” variant, but that all of them can be interesting in certain contexts. For example, even if an agent does not have a weakly dominant strategy de re, it might be useful to know that she has one de dicto. When it comes to the relationship between game theoretic properties and properties of the goals, an interesting case is positive goals: if an agent’s goal is positive, then not only is she guaranteed to have a weakly dominant strategy, but she is also guaranteed to know that she does and to know which one it is. For the Nash equilibrium, it turns out that for general epistemic goal structures (not necessarily with positive goals) there can never be common knowledge of a Nash equilibrium that does not consist of trivial announcements. When it comes to more general Nash equilibria, the notion of an induced game we introduced can be seen as a variant of Bayesian games, but our definition is less baroque and is perhaps more natural from the perspective of logical semantics.

8.1. RELATED AND FUTURE WORK

Boolean games The intimate connection between knowledge and strategies in public announcement games distinguishes them from many other types of games. In *Boolean games* (Harrenstein, 2004; Harrenstein et al., 2001), each agent has a goal formula like in PAGs, and each agent controls a set of primitive propositions which affects the truth value of the goal formulas. In contrast, in PAGs an agent “controls” common knowledge of any formula he or she knows. We have seen that we cannot simply view a pointed PAG as Boolean type game, because the agents do not necessarily have common knowledge about the game that is being played.

Question-answer games Consider a different game. Instead of players choosing what to announce, players choose which questions to pose to

another player, where the other player is obliged to answer the question truthfully. Of course there are just as many questions to be posed to player j as (truly) different announcements for player j to make. And one variable in such a game could be that a player i whose turn it is may *choose* another player j to ask a question to. We can already observe that, given such a choice, the total number of strategies of the induced question-answer game is also countable given some model M involving at least players i and j as initial parameters. Instead of 2^{m^2-m} pure strategies for player i with m equivalence classes, we now have $2^{m_j m_i - m_i}$ pure strategies for player i asking player j a question, where player i has m_i equivalence classes and player j has m_j equivalence classes. This we can see as follows. There are 2^{m_j-1} different dichotomies for player j (i.e. coarsening 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 $(2^{m_j-1})^{m_i} = 2^{m_j m_i - m_i}$.

This sort of question-answer game is defined as a knowledge game in (van Ditmarsch, 2000) (see also (van Ditmarsch, 2002; van Ditmarsch, 2004)), for 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. From our current perspective this is a subset of all different questions to be asked, namely only those with singleton equivalence classes. An initial investigation into question-answer games can be found in (Ågotnes et al., 2010b), also based on (van Benthem and Minica, 2009).

Further generalizations In future work we will also look at more fine grained goal models which do not necessarily give binary payoffs, for example weighted formulae or CP-nets. Also, we want to describe mixed equilibria for public announcement games. Finally, we will model situations with *sequential* announcements by using extensive form games. It is of course very common that agents make informative moves one after the other, as in a dialogue. So such an analysis would be very relevant. Unfortunately, it is considered much harder to compute equilibria for the extensive form.

Acknowledgements

This work started life as a presentation at a Workshop on Reasoning about Knowledge and Rational Action in honour of Professor Wiebe van der Hoek on his 50th birthday, Liverpool, UK. A later version was presented at the 6th conference on Logic, Game Theory and Social Choice (LGS 6), Tsukuba, Japan. We thank the anonymous reviewers, as well as attendants at the mentioned meetings, for comments that have helped us improve the paper. We also thank Giacomo Bonanno, Pål Grønås Drange and Stefan Minica for helpful remarks.

References

- Ågotnes, T., P. Balbiani, H. van Ditmarsch, and P. Seban: 2010a, ‘Group Announcement Logic’. *Journal of Applied Logic* **8**(1), 62–81.
- Ågotnes, T., J. van Benthem, H. van Ditmarsch, and S. Minica: 2010b, ‘Question-Answer Games’. 9th Conference on Logic and the Foundations of Game and Decision Theory (LOFT 2010) (informal proceedings).
- Ågotnes, T. and H. van Ditmarsch: 2008, ‘Coalitions and Announcements’. In: Padgham, Parkes, Muller, and Parsons (eds.): *Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*. Estoril, Portugal, pp. 673–680.
- Aucher, G.: 2003, ‘A combined system for update logic and belief revision’. Master’s thesis, ILLC, University of Amsterdam, Amsterdam, the Netherlands. ILLC report MoL-2003-03.
- Balbiani, P., A. Baltag, H. van Ditmarsch, A. Herzig, T. Hoshi, and T. D. Lima: 2008, “Knowable” as ‘known after an announcement’. *Review of Symbolic Logic* **1**(3), 305–334.
- Baltag, A., L. Moss, and S. Solecki: 1998, ‘The Logic of Public Announcements, Common Knowledge, and Private Suspicions’. In: I. Gilboa (ed.): *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 98)*. pp. 43–56.
- Baltag, A. and S. Smets: 2008, ‘Dynamic Belief Revision over Multi-Agent Plausibility Models’. In: G. Bonanno, W. van der Hoek, and M. Wooldridge (eds.): *After-conference Proceedings of LOFT 2006*. Amsterdam University Press.
- Boutilier, C., R. I. Brafman, C. Domschlag, H. H. Hoos, and D. Poole: 2003, ‘CP-Nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements’. *Journal of Artificial Intelligence Research* **21**, 135–191.
- Boutilier, C., R. I. Brafman, H. H. Hoos, and D. Poole: 1999, ‘Reasoning with conditional ceteris paribus preference statements’. In: *Proceedings of UAI99*. pp. 71–80.
- Boutilier, C., R. I. Brafman, H. H. Hoos, and D. Poole: 2004, ‘Preference-Based Constrained Optimization with CP-Nets’. *Computational Intelligence* **20**(2), 137–157.
- Elkind, E., L. A. Goldberg, P. Goldberg, and M. Wooldridge: 2009, ‘A Tractable and Expressive Class of Marginal Contribution Nets and Its Applications’. *Mathematical Logic Quarterly* **55**(4), 362–376.

- Fagin, R., J. Halpern, Y. Moses, and M. Vardi: 1995, *Reasoning about Knowledge*. Cambridge MA: MIT Press.
- Gerbrandy, J.: 1999, 'Bisimulations on Planet Kripke'. Ph.D. thesis, University of Amsterdam. ILLC Dissertation Series DS-1999-01.
- Gerbrandy, J. and W. Groeneveld: 1997, 'Reasoning about Information Change'. *Journal of Logic, Language, and Information* **6**, 147–169.
- Goranko, V. and M. Otto: 2007, 'Model Theory of Modal Logic'. In: P. Blackburn, J. van Benthem, and F. Wolter (eds.): *Handbook of Modal Logic*. Kluwer, pp. 249–329.
- Harel, D.: 1984, 'Dynamic Logic'. In: D. Gabbay and F. Guenther (eds.): *Handbook of Philosophical Logic*, Vol. II. Dordrecht, pp. 497–604, Kluwer Academic Publishers.
- Harel, D., D. Kozen, and J. Tiuryn: 2000, *Dynamic Logic*. Cambridge MA: MIT Press. Foundations of Computing Series.
- Harrenstein, P.: 2004, 'Logic in Conflict'. Ph.D. thesis, Utrecht University.
- Harrenstein, P., W. van der Hoek, J.-J. Meyer, and C. Witteveen: 2001, 'Boolean Games'. In: J. van Benthem (ed.): *Proceeding of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge (TARK VIII)*. Siena, Italy, pp. 287–298.
- Harsanyi, J. C.: 1967–1968, 'Games with Incomplete Information Played by 'Bayesian' Players, Parts I, II, and III'. *Management Science* **14**, 159–182, 320–334, 486–502.
- Hintikka, J.: 1962, *Knowledge and Belief*. Ithaca, NY: Cornell University Press.
- Ieong, S. and Y. Shoham: 2005, 'Marginal Contribution Nets: A Compact Representation Scheme for Coalitional Games'. In: *Proceedings of the Sixth ACM Conference on Electronic Commerce (EC'05)*. Vancouver, Canada.
- Jamroga, W. and T. Ågotnes: 2006, 'What Agents Can Achieve under Incomplete Information'. In: P. Stone and G. Weiss (eds.): *Proc. of the Fifth Intern. Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS)*. pp. 232–234, ACM Press.
- Jamroga, W. and T. Ågotnes: 2007, 'Constructive knowledge: what agents can achieve under imperfect information'. *Journal of Applied Non-Classical Logics* **17**(4), 423–475.
- Jamroga, W. and W. van der Hoek: 2004, 'Agents that Know How to Play'. *Fundamenta Informaticae* **63**, 185–219.
- Lang, J., U. Endriss, and Y. Chevaleyre: 2006, 'Expressive Power of Weighted Propositional Formulas for Cardinal Preference Modelling'. In: *Proceedings of Knowledge Representation and Reasoning (KR 2006)*. Lake District, England.
- Meyer, J.-J. and W. van der Hoek: 1995, *Epistemic Logic for AI and Computer Science*. Cambridge Tracts in Theoretical Computer Science 41. Cambridge: Cambridge University Press.
- Osborne, M.: 2004, *An Introduction to Game Theory*. New York: Oxford University Press.
- Osborne, M. and A. Rubinstein: 1994, *A Course in Game Theory*. Cambridge MA: MIT Press.
- Plaza, J.: 1989, 'Logics of Public Communications'. In: M. Emrich, M. Pfeifer, M. Hadzikadic, and Z. Ras (eds.): *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems: Poster Session Program*. pp. 201–216, Oak Ridge National Laboratory. ORNL/DSRD-24.
- Uckelman, J., Y. Chevaleyre, U. Endriss, and J. Lang: 2009, 'Representing Utility Functions via Weighted Goals'. *Mathematical Logic Quarterly* **55**(4), 341–361.

- van Benthem, J.: 1989, ‘Semantic Parallels in Natural Language and Computation’. In: *Logic Colloquium '87*. Amsterdam, North-Holland.
- van Benthem, J.: 2002, ‘One is a lonely number: on the logic of communication’. Technical report, University of Amsterdam. ILLC Research Report PP-2002-27 (material presented at the Logic Colloquium 2002).
- van Benthem, J., J. Gerbrandy, T. Hoshi, and E. Pacuit: 2009, ‘Merging frameworks for interaction’. *Journal of Philosophical Logic* **38**, 491–526.
- van Benthem, J. and F. Liu: 2007, ‘Dynamic logic of preference upgrade’. *Journal of Applied Non-Classical Logics*. To appear.
- van Benthem, J. and S. Minica: 2009, ‘Toward a Dynamic Logic of Questions’. In: X. He, J. Horty, and E. Pacuit (eds.): *Logic, Rationality, and Interaction. Proceedings of LORI 2009*. pp. 27–41, Springer. LNCS 5834.
- van Benthem, J., J. van Eijck, and B. Kooi: 2006, ‘Logics of Communication and Change’. *Information and Computation* **204(11)**, 1620–1662.
- van der Hoek, W. and M. Pauly: 2006, ‘Modal Logic for Games and Information’. In: J. van Benthem, P. Blackburn, and F. Wolter (eds.): *The Handbook of Modal Logic*. Amsterdam, The Netherlands: Elsevier, pp. 1180–1152.
- van Ditmarsch, H.: 2000, ‘Knowledge games’. Ph.D. thesis, University of Groningen. ILLC Dissertation Series DS-2000-06.
- van Ditmarsch, H.: 2002, ‘The description of game actions in Cluedo’. In: L. Petrosian and V. Mazalov (eds.): *Game Theory and Applications*, Vol. 8. Commack, NY, USA, pp. 1–28, Nova Science Publishers.
- van Ditmarsch, H.: 2004, ‘Some game theory of Pit’. In: C. Zhang, H. Guesgen, and W. Yeap (eds.): *Proceedings of PRICAI 2004 (Eighth Pacific Rim International Conference on Artificial Intelligence)*. pp. 946–947, Springer. LNAI 3157.
- van Ditmarsch, H.: 2005, ‘Prolegomena to dynamic logic for belief revision’. *Synthese (Knowledge, Rationality & Action)* **147**, 229–275.
- van Ditmarsch, H. and B. Kooi: 2006, ‘The Secret of My Success’. *Synthese* **151**, 201–232.
- van Ditmarsch, H., W. van der Hoek, and B. Kooi: 2005, ‘Dynamic Epistemic Logic with Assignment’. In: *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 05)*. New York, pp. 141–148, ACM Inc.
- van Ditmarsch, H., W. van der Hoek, and B. Kooi: 2007, *Dynamic Epistemic Logic*, Vol. 337 of *Synthese Library*. Springer.
- van Linder, B., W. van der Hoek, and J.-J. Meyer: 1995, ‘Actions that make you change your mind’. In: A. Laux and H. Wansing (eds.): *Knowledge and Belief in Philosophy and Artificial Intelligence*. Berlin, pp. 103–146, Akademie Verlag.
- von Wright, G.: 1951, *An Essay in Modal Logic*. Amsterdam: North Holland.

Appendix

A. Bisimulations and Contractions

The notion of bisimulation for structural similarity is well known in modal logic. Here we state the most basic definitions and some well known properties that we use in this article; we refer to, e.g., (Goranko and Otto, 2007) for further details.

Let two epistemic models $M = (S, \sim_1, \dots, \sim_n, V)$ and $M' = (S', \sim'_1, \dots, \sim'_n, V')$ be given. A non-empty relation $\mathfrak{R} \subseteq S \times S'$ is a *bisimulation* between M and M' iff for all $s \in S$ and $s' \in S'$ with $(s, s') \in \mathfrak{R}$:

atoms for all $p \in \Theta$: $s \in V(p)$ iff $s' \in V'(p)$;

forth for all $i \in N$ and all $t \in S$: if $s \sim_i t$, then there is a $t' \in S'$ such that $s' \sim'_i t'$ and $(t, t') \in \mathfrak{R}$;

back for all $i \in N$ and all $t' \in S'$: if $s' \sim'_i t'$, then there is a $t \in S$ such that $s \sim_i t$ and $(t, t') \in \mathfrak{R}$.

We write $(M, s) \Leftrightarrow (M', s')$ to mean that there is a bisimulation between M and M' linking s and s' , and we then call the pointed epistemic models (M, s) and (M', s') *bisimilar*.

Let Z_M be the greatest bisimulation relation on M . Note that Z_M is an equivalence relation on S . For all $s \in S$, let $\|s\|$ be the equivalence class of s modulo Z_M . The *bisimulation contraction* of M is the structure $\|M\| = (S', \sim'_1, \dots, \sim'_n, V')$ such that:

- $S' = S|_{Z_M}$, i.e. the quotient of S modulo Z_M
- $\|s\| \sim'_i \|t\|$ iff there exist $v, w \in S$ such that $sZ_M v$, $tZ_M w$ and $v \sim_i w$
- $V'(p) = V(p)|_{Z_M}$

The following are well known and/or obvious:

- If $(M, s) \Leftrightarrow (M', s')$, then for any formula φ , $M, s \models \varphi$ iff $M', s' \models \varphi$
- $(M, s) \Leftrightarrow (\|M\|, \|s\|)$
- If $(\|M\|, \|s\|) \Leftrightarrow (\|M\|, \|t\|)$ then $\|s\| = \|t\|$ (the contraction contains no distinct bisimilar states)

B. Proof of Lemma 7

Proof. Let M, s be a pointed epistemic model and i an agent.

1. First, we show that the statement holds when M is a bisimulation contraction (see Appendix A). In that case, we construct the maximal informative announcement (MIA) α as follows. Let $\alpha_{M,s,i}$

be the *characteristic formula* for the i -modality in (M, s) . The characteristic formula has the property that (for finite structures M) it is true also in (M, s') iff (M, s) and (M, s') are bisimilar; see (Goranko and Otto, 2007) for details. We now use characteristic formulae to construct a formula that is true exactly throughout the s -equivalence class for i as follows. Let $X = \{s_0, \dots, s_k\} \subseteq S$ ($k \geq 0$) be such that $s_0 = s$ and $[s_0]_i, \dots, [s_k]_i$ partitions S ; i.e., each s_j is a representative of one of i 's equivalence classes. Let

$$\alpha = \neg \bigvee_{x \in X \setminus \{s_0\}} \alpha_{M,x,i}.$$

It is easy to see that

$$M|K_i\alpha = [s]_i. \quad (1)$$

Formally: first, assume that $t \in [s]_i$, and that $M, t \not\models K_i\alpha$. That means that there is a $t' \in [t]_i$ such that $M, t' \not\models \alpha$, i.e., that $M, t' \models \alpha_{M,s_j,i}$ for some $j > 0$. By the nature of the characteristic formula, that means that M, t' and M, s_j are bisimilar. But since $t' \in [t]_i = [s_0]_i$ that means that $t' \neq s_j$, which contradicts the assumption that M is a bisimulation contraction. Conversely, assume that $t \in S$ and $t \notin [s]_i$. It must be the case that $t \in [s_j]_i$ for some $j > 0$. Thus, $M, t \models \alpha_{M,s_j,i}$, and it follows that $M, t \not\models \alpha$ and thus by reflexivity that $M, t \not\models K_i\alpha$. Thus, (1) holds.

From (1) it is clear that α is a MIA: first, $M, s \models K_i\alpha$. Second, assume that $M, s \models K_i\psi$. If $M|K_i\alpha \not\subseteq M|K_i\psi$, then there is a $t \in M|K_i\alpha = [s]_i$ such that $t \notin M|K_i\psi$, which means that there is a t such that $s \sim_i t$ and $M, t \not\models K_i\psi$, which contradicts the assumption that $M, s \models K_i\psi$.

Now, consider the case that M is not necessarily a bisimulation contraction. Let $\|M\|$ be the bisimulation contraction of M (see Appendix A), and let α be the MIA for i in $(\|M\|, \|s\|)$. We argue that α is also the MIA for i in (M, s) . First, $M, s \models K_i\alpha$ since $\|M\|, \|s\| \models K_i\alpha$ and $(\|M\|, \|s\|)$ and (M, s) are bisimilar. Second, let ψ be a formula such that $M, s \models K_i\psi$. We must show that $M|K_i\alpha \subseteq M|K_i\psi$. Again by bisimilarity, we have that $\|M\|, \|s\| \models K_i\psi$ and thus that $\|M\||K_i\alpha \subseteq \|M\||K_i\psi$ by the fact that α is a MIA in $(\|M\|, \|s\|)$, i.e., that $\{\|t\| : \|M\|, \|t\| \models K_i\alpha\} \subseteq \{\|t\| : \|M\|, \|t\| \models K_i\psi\}$ which, by bisimilarity, means that $\{\|t\| : M, t \models K_i\alpha\} \subseteq \{\|t\| : M, t \models K_i\psi\}$, i.e., that $\{t : M, t \models K_i\alpha\} \subseteq \{t : M, t \models K_i\psi\}$. Thus, $M|K_i\alpha \subseteq M|K_i\psi$.

2. Let α and β both be MIAs, and let t be some state. If $M, t \models K_i\alpha$, then also $M, t \models K_i\beta$ by the fact that $M|K_i\alpha \subseteq M|K_i\beta$. Similarly with α and β switched. Thus, $M \models K_i\alpha \leftrightarrow K_i\beta$.

3. Assume otherwise; that $M, t \models \varphi \wedge K_i \alpha$ and $M, t \not\models \varphi \wedge K_i \beta$. But that means that $M, t \models K_i \alpha$ and $M, t \not\models K_i \beta$, which contradicts the fact that $K_i \alpha$ is a MIA.
4. Let α be the MIA in s . $M, s \models K_i \alpha$ implies that $M, t \models K_i \alpha$. If $M, t \models K_i \psi$ for some ψ , then also $M, s \models K_i \psi$, so $M|K_i \alpha \subseteq M|K_i \psi$ by the fact that $K_i \alpha$ is a MIA in M, s . Thus, $K_i \alpha$ is also a MIA in M, t .

C. Proof of Theorem 24

Proof. Let a_1, \dots, a_n be given. We show that a_1, \dots, a_n is *not* a NE in the induced game iff $\{a_{i,t_i} : i \in N, t_i \in T_i\}$ ($a_{i,t_i} = a_i(s)$, $s \in \tau^{-1}(t_i)$) is *not* a NE in the associated Bayesian game.

$\{a_{i,t_i} : i \in N, t_i \in T_i\}$ is not a NE in the associated Bayesian game iff there are $j \in N$, $t_j \in T_j$, and $a'_{j,t_j} \in A_{j,t_j}$ such that $u_{j,t_j}(\{a_{i,t_i} : i \in N, t_i \in T_i\}[a_{j,t_j}/a'_{j,t_j}]) > u_{j,t_j}(\{a_{i,t_i} : i \in N, t_i \in T_i\})$ iff $\sum_{s \in S} Pr(s : t_j) u_j^s((a_{1,t_1}, \dots, a_{n,t_n})[a_{j,t_j}/a'_{j,t_j}]) > \sum_{s \in S} Pr(s : t_j) u_j^s(a_{1,t_1}, \dots, a_{n,t_n})$ iff, since $Pr(s : t_j) = 0$ when $s \notin t_j$, $\sum_{s \in t_j} Pr(s : t_j) u_j^s((a_{1,t_1}, \dots, a_{n,t_n})[a_{j,t_j}/a'_{j,t_j}]) > \sum_{s \in t_j} Pr(s : t_j) u_j^s(a_{1,t_1}, \dots, a_{n,t_n})$ iff $\sum_{s \in t_j} \frac{1}{|t_j|} u_j^s((a_{1,t_1}, \dots, a_{n,t_n})[a_{j,t_j}/a'_{j,t_j}]) > \sum_{s \in t_j} \frac{1}{|t_j|} u_j^s(a_{1,t_1}, \dots, a_{n,t_n})$ iff $\frac{1}{|t_j|} \sum_{s \in t_j} u_j^s((a_{1,t_1}, \dots, a_{n,t_n})[a_{j,t_j}/a'_{j,t_j}]) > \frac{1}{|t_j|} \sum_{s \in t_j} u_j^s(a_{1,t_1}, \dots, a_{n,t_n})$ iff

$$\sum_{s \in t_j} u_j^s((a_{1,t_1}, \dots, a_{n,t_n})[a_{j,t_j}/a'_{j,t_j}]) > \sum_{s \in t_j} u_j^s(a_{1,t_1}, \dots, a_{n,t_n}) \quad (2)$$

On the other hand, a_1, \dots, a_n is not a NE in the induced game iff there are $j \in N$ and a'_j such that $\sum_{s \in S} \frac{u_j^s((a_1(s), \dots, a_n(s))[a_j(s)/a'_j(s)])}{|S|} > \frac{\sum_{s \in S} u_j^s(a_1(s), \dots, a_n(s))}{|S|}$ iff

$$\sum_{s \in S} u_j^s((a_1(s), \dots, a_n(s))[a_j(s)/a'_j(s)]) > \sum_{s \in S} u_j^s(a_1(s), \dots, a_n(s)) \quad (3)$$

First, let j and a'_j be such that (3) holds. There must be at least one $s' \in S$ such that $\sum_{s \in [s']_j} u_j^s((a_1(s), \dots, a_n(s))[a_j(s)/a'_j(s)]) > \sum_{s \in [s']_j} u_j^s(a_1(s), \dots, a_n(s))$ (otherwise is a contradiction of (3)). Now (2) holds by taking $t_j = [s']_j$, and $a'_{j,t_j} = a'_j(s')$. Thus, (3) implies (2).

Second, let j, t_j and a'_{j,t_j} be such that (2) holds. Define a'_j as follows:

$$a'_j(s) = \begin{cases} a'_{j,t_j} & \text{if } s \in t_j \\ a_j(s) & \text{otherwise} \end{cases}$$

We have that $\sum_{s \in S} u_j^s((a_1(s), \dots, a_n(s))[a_j(s)/a'_j(s)]) = \sum_{s \in t_j} u_j^s((a_1(s), \dots, a_n(s))[a_j(s)/a'_j(s)]) + \sum_{s \in S \setminus t_j} u_j^s((a_1(s), \dots, a_n(s))[a_j(s)/a'_j(s)]) = \sum_{s \in t_j} u_j^s((a_{1,t_1}, \dots, a_{n,t_n})[a_{j,t_j}/a'_{j,t_j}]) + \sum_{s \in S \setminus t_j} u_j^s(a_{1,t_1}, \dots, a_{n,t_n})$ which, by (2), is greater than $\sum_{s \in t_j} u_j^s(a_{1,t_1}, \dots, a_{n,t_n}) + \sum_{s \in S \setminus t_j} u_j^s(a_{1,t_1}, \dots, a_{n,t_n}) = \sum_{s \in S} u_j^s(a_{1,t_1}, \dots, a_{n,t_n}) = \sum_{s \in S} u_j^s(a_1(s), \dots, a_n(s))$, and thus (3) holds. Thus, (3) implies (2).

