Dynamic Aleatoric Reasoning in Games of Bluffing and Chance

Extended Abstract

Tim French
The University of Western Australia
Crawley, Western Australia
tim.french@uwa.edu.au

Andrew Gozzard
The University of Western Australia
Crawley, Western Australia
andrew.gozzard@uwa.edu.au

Mark Reynolds
The University of Western Australia
Crawley, Western Australia
mark.reynolds@uwa.edu.au

ABSTRACT

Games of chance and bluffing, such as bridge, The Resistance, and poker allow epistemic reasoning. Players know their own cards while being uncertain of opponents'. Success generally involves reducing your uncertainty without reducing that of your opponents. Reasoning in such games requires a mix of logical (deducing what is possible) and probabilistic (what is likely). We present a *dynamic aleatoric logic* for epistemic reasoning in such games.

KEYWORDS

Probabilistic reasoning; Dynamic epistemic logic

ACM Reference Format:
Tim French, Andrew Gozzard, and Mark Reynolds. 2019. Dynamic Aleatoric Reasoning in Games of Bluffing and Chance. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

This paper proposes a probabilistic generalisation of dynamic epistemic logic for reasoning about games of bluffing and chance. Typically these games have a hidden epistemic state, so that the knowledge of all agents is not equal. There is also an element of chance, either coming through an initial deal of cards, or some random element such as dice or a coin. Finally there should be a strategic advantage to having knowledge, an incentive to discover what the opponent knows, and to hide their own knowledge. Such games include traditional games like Poker and Bridge, and the more recent The Resistance and Hanabi. *Aleatoric*, Latin for "depending on the throw of a dice", describes both explicit elements of such games (card deals, dice rolls) as well as the policies and strategies of players in the game (so a player may bluff 10% of the time).

There has been substantial work in this direction: Hailperin [6] and Nilsson [10] generalised propositional logic so the the semantics of true and false are replaced by probability measures, and in [12] Williamson provided an inference system based on Bayesian epistemology. Feldman and Harel [4] and Kozen [9] gave a probabilistic variation of propositional dynamic logic. Epistemic logics also have probabilistic variants [3, 7] based on Dempster-Shafer models of belief, and Kooi and van Benthem [8, 11] extended dynamic epistemic logic with explicit probabilities. Baltag and Smets [1] provided similar extensions in the context of belief revision.

Here we provide a lightweight logic, the *aleatoric dynamic epistemic logic*, building on recent development of the *modal aleatoric calculus* [5] for formalising reasoning processes in games of chance. To demonstrate the logic, we use *The Resistance*, a card game where players are required to sabotage one another without revealing their true purpose.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

2 SYNTAX AND SEMANTICS

We take a many-valued approach. Rather than presenting a logic that describes what is *true* about a probabilistic scenario, we present the *aleatoric dynamic epistemic logic* (ADEL) for determining what is likely. The difference is subtle: In probabilistic dynamic epistemic logic [8] it is possible to express that the statement "Alice thinks X has probability 0.5" is true; whereas the calculus here simply has a term "Alice's expectation of X" which may have the value 0.5.

The syntax is given for a set of random variables X, and a set of agents N. We also a constants $\top$ and $\bot$. The syntax of aleatoric dynamic epistemic logic, $\mathcal{ADEL}$, is as follows:

$$\alpha ::= \; x \mid \top \mid \bot \mid (\alpha?\alpha:\alpha) \mid (\alpha \,|\, \alpha)_i \mid [\alpha]\alpha$$

where $x \in X$ is a random variable and $i \in N$ is an agent. As usual, we let $v(\alpha)$ refer to the set of variables that appear in α. We refer to $\top$ as *always* and $\bot$ as *never*. The *if-then-else* operator $(\alpha?\beta:\gamma)$ is read *if* α *then* β *else* γ and uses the ternary conditional syntax of programming languages such as C. The *marginal expectation* operator $(\alpha \,|\, \beta)_i$ is *agent* i 's expectation of α given β. The *global observation* operator $[\alpha]\beta$ is *the expectation of* β *once* α *is observed by all agents*. This corresponds to Bayesian conditioning on a public announcement of α.

Some abbreviations we can define in $\mathcal{ADEL}$ are as follows: $\alpha \wedge \beta = (\alpha?\beta:\bot)$, $\alpha \vee \beta = (\alpha?\top:\beta)$, $\alpha \rightarrow \beta = (\alpha?\beta:\top)$ and $\neg\alpha = (\alpha?\bot:\top)$.

The aleatoric dynamic epistemic logic is interpretted over *probability models* similar to the probability structures of [7],

Definition 2.1. Given a set S, $PD(S)$ is the set of *probability distributions over* S, where $\mu \in PD(S)$ implies: $\mu : S \longrightarrow [0,1]$; and $\Sigma_{s\in S}\mu(s) = 1$. Given a set of variables X and a set of agents N, a *probability model* is specified by the tuple $P = (W, \pi, f)$, where:

- W is a set of possible worlds.
- $\pi : N \longrightarrow W \longrightarrow PD(W)$ assigns for each agent, for each world $w \in W$, a probability distribution $\pi_i(w)$ over W, such that for all i, for all $u, v \in W$ $\pi_i(u,v) > 0$ implies $\pi_i(u) = \pi_i(v)$. We will write $\pi_i(w,v)$ in place of $\pi(i)(w)(v)$.
- $f : W \longrightarrow X \longrightarrow [0,1]$ is a probability assignment so $f_w(x)$ is the probability of the variable x being true at the world w.

A *pointed probability model*, $P_w = (W, \pi, f, w)$, specifies a world in the model as the point of evaluation.

Given pointed model P_w, the semantic interpretation of formula α is $P_w(\alpha) \in [0,1]$ which is the expectation of the formula being true in a model,

Definition 2.2. The *semantics of the dynamic aleatoric calculus* take a pointed probability model, f_w, and a proposition defined in $\mathcal{ADEL}$, α, and calculate the *expectation of* α *holding at* P_w. Given an agent i, a world w and a $\mathcal{ADEL}$ formula α, we define *i's expectation of* α *at* w as

$$E_w^i(\alpha) = \sum_{u \in W} \pi_i(w,u).P_u(\alpha).$$

Then the semantics of ADEL are as follows:

$$\begin{array}{rcl} P_w(\top) & = & 1 \qquad P_w(\bot) \;=\; 0 \qquad P_w(x) \;=\; f_w(x) \\ P_w((\alpha?\beta\!:\!\gamma)) & = & P_w(\alpha).P_w(\beta) + (1 - P_w(\alpha)).P_w(\gamma) \\ P_w((\alpha\,|\,\beta)_i) & = & \frac{E_w^i(\alpha\wedge\beta)}{E_w^i(\beta)} \text{ if } E_w^i(\beta) > 0 \text{ and } 1 \text{ otherwise} \\ P_w([\alpha]\beta) & = & P_w^\alpha(\beta) \end{array}$$

where P^α is the model (W, π', f) such that for all $u, v \in W$, $\pi_i'(u,v) = \frac{P_v(\alpha).\pi_i(u,v)}{E_u^i(\alpha)}$ if $E_u^i(\alpha) > 0$, and $\pi_i(u,v)$ otherwise.

The concept of *sampling* is intrinsic to the semantics. Suppose our agents are committed aleators, in that they use labelled coins (or sample probability distributions) for everything. If we ask "is x true" the agent will take the coin marked x, flip it and if it lands heads, reply "yes". Every formula is evaluated as a sampling process this way. To interpret $(\alpha?\beta\!:\!\gamma)$, the agent will execute the sampling procedure for α and if it returns true, the agent will proceed with the sampling procedure for β, otherwise the agent samples γ.

The conditional expectation operator $(\alpha\,|\,\beta)_i$ expresses agent i's expectation of α *conditioned* on β. This is, as in the Kolmogorov definition of conditional probability, agent i's expectation of $\alpha \wedge \beta$ divided by agent i's expectation of β.

The observation operation $[\alpha]\beta$ is the expectation of β after α is *observed by all agents* (or *publicly announced* in the terminology of dynamic epistemic logic). The interpretation of α is also stochastic, so we imagine that as before, the mental model of the universe is sampled, and α is true in that sampling. Further, we suppose that all agents are told that α was true in that sampling. Now every agent updates their mental model of the universe. P_w is their prior expectation of the universe, and we apply Bayesian conditioning to determine the new (posterior) model.

3 EXAMPLE: THE RESISTANCE

The Resistance [2] is a bluffing game for five to ten players. Approximately one third of the players are allocated as being government spies, while the rest are true members of the resistance. The spies know each others' identity, but the true members do not know who is a spy. The game consists of five rounds, each as follows:

(1) A leader is allocated (randomly, or next left to previous)
(2) The leader proposes a group of players to go on a "mission". The size of the group is given (depending on the number of players and round), and the leader may include themselves.
(3) Players vote publicly on whether the mission proceeds. If there is not a majority for, leadership moves left, and restarts. Five missions voted against in a row means the spies win.

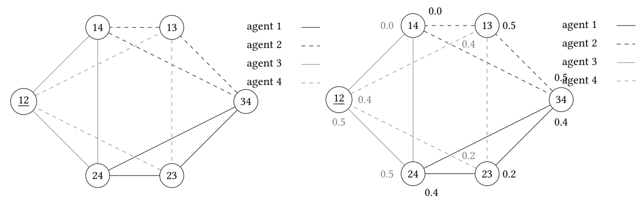

(a) Before betrayal. (b) After betrayal.

Figure 1: Epistemic probability models, before and after. In Figure 1a all connected worlds are considered equally likely.

(4) Each player on the mission plays a token (face down) to indicate whether they betray the mission. These are shuffled and then revealed to everyone. If a betrayal token was played, the mission fails, and otherwise it succeeds.

The true members want a majority of missions to succeed, whilst the spies want a majority to fail. As spies are in a minority, they hide their identity. However, spies need to influence the debate and vote so they are sent on enough missions to achieve their goal.

Here we present a simple analysis of a small version of the game with four players, $\{1,2,3,4\}$, two of them spies. This gives six possible initial configurations. Spies know the identity of all the other spies, the other players do not. The non-spies know only that they are not spies, and therefore assign equal probability to the three worlds in which they are not spies, and zero to the worlds in which they are spies. For every agent, $i \in \{1,2,3,4\}$, there is a variable x_i, which is the probability of i betraying a mission, if i is a spy on that mission. We will suppose that for all i, for all worlds where i is a spy, x_i has initial value $\frac{3}{4}$. There are also variables s_i to dictate who is a spy, so s_1 has probability 1 in worlds w_{12}, w_{13} and w_{14} and probability 0 in all other worlds. This gives the model in Figure 1a, which is a common prior for all players. The left-most world is underlined: the actual world where 1 and 2 are spies.

Suppose that 2, 3 are sent on a mission, and 2 betrays it. All agents are informed that exactly one agent betrayed the mission, which is equivalent to the announcement $(x_2 \wedge \neg x_3) \vee (\neg x_2 \wedge x_3)$. We can calculate this event has 0 probability in the world (14), since neither 2 nor 3 are spies in that world. The event has $\frac{3}{4}$ probability in worlds (12), (13), (24) and (34), and probability $\frac{97}{256}$ in (23).

Agents infer different information from this announcement. Agent 3 will know 2 is a spy, and assigns equal probability to 1 and 4 being spies. Agent 4 does not know who is a spy, but the fact that only one agent betrayed the mission makes it less likely that both 2 and 3 are spies, so 4's expectation that 1 is a spy actually increases.

The aleatoric calculus allows us to express more complex policies for agents. In the instance described above an agent simply flips a biased coin (with probability $\frac{3}{4}$ of coming up betray). However, we could also specify a policy whereby agent 2 will betray, if both non-spies think 2 is a spy, or if both non-spies think 2 is not a spy. That is, in world w_{12} of Figure 1a agent 2's likelihood of betraying a mission could be $P_{w_{12}}(E_2(E_3 s_2 \wedge E_4 s_2) \vee E_2(E_3 \neg s_2 \wedge E_4 \neg s_2)) = \frac{46}{81}$.

This demonstrates the reasoning capabilities of aleatoric dynamic epistemic logic.

REFERENCES

[1] Alexandru Baltag and Sonja Smets. 2008. Probabilistic dynamic belief revision. *Synthese* 165, 2 (2008), 179.

[2] Don Eskridge. 2010. The Resistance. (2010). http://www.indieboardsandcards.com/resistance.php.

[3] Ronald Fagin, Joseph Y Halpern, and Nimrod Megiddo. 1990. A logic for reasoning about probabilities. *Information and computation* 87, 1-2 (1990), 78–128.

[4] Yishai A Feldman and David Harel. 1982. A probabilistic dynamic logic. In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing*. ACM, 181–195.

[5] T. French, A. Gozzard, and M.Reynolds. 2019. A modal aleatoric calculus for probabilistic reasoning. In *Proceedings of the Indian Conference on Logic and its Applications (Lecture Notes in Computer Science)*. Springer, 52–63.

[6] Theodore Hailperin. 1976. *Boole's logic and probability: a critical exposition from the standpoint of contemporary algebra, logic and probability theory.* North-Holland.

[7] Joseph Y Halpern. 2017. *Reasoning about uncertainty.* MIT press.

[8] Barteld P Kooi. 2003. Probabilistic dynamic epistemic logic. *Journal of Logic, Language and Information* 12, 4 (2003), 381–408.

[9] Dexter Kozen. 1985. A probabilistic pdl. *J. Comput. System Sci.* 30, 2 (1985), 162–178.

[10] Nils J Nilsson. 1986. Probabilistic logic. *Artificial intelligence* 28, 1 (1986), 71–87.

[11] Johan Van Benthem, Jelle Gerbrandy, and Barteld Kooi. 2009. Dynamic update with probabilities. *Studia Logica* 93, 1 (2009), 67.

[12] Jon Williamson. 2013. From Bayesian epistemology to inductive logic. *Journal of Applied Logic* 11, 4 (2013), 468–486.